Readability of conventional applications within immersive virtual environments

A thesis
submitted in fulfilment
of the requirements for the Degree
of
Doctor of Philosophy in Computer Science
at
The University of Waikato
by

Cameron Grout

THE UNIVERSITY OF
WAIKATO
Te Whare Wānanga o Waikato

2022

There have been times, lately, when I dearly wished that I could change the past. Well, I can't, but I can change the present, so that when it becomes the past it will turn out to be a past worth having.

Terry Pratchett

— I Shall Wear Midnight

Abstract

This thesis explores the potential role of virtual reality in allowing users to make full use of the space around them for computing tasks, and the impact such an environment could have when using conventional application software.

As computers have become more core to our lives, the number of tasks that we perform on them have increased. Computer systems have evolved their support for multiple tasks through software tooling to aid management and layout of windows, and adoption of large- or multi-display output devices to increase digital work space. While usage of software tooling is mixed among users, a steady trend in adoption of larger and more numerous displays has been observed, suggesting that these adaptations are important to users. These developments may not be sufficient however, as users still encounter issues with layout and identification of task-related windows that impact on their performance.

Use of modern, high-quality, mass-market virtual reality devices to replace existing displays has the potential to address the shortcomings of existing solutions. These devices can present large virtual environments in which windows can be laid out, and provide novel interaction techniques for facilitating layout, organisation, and identification of windows. While widespread adoption of such technology might lead to purpose-built applications to replace existing systems; in the short-term finding ways to support existing applications in immersive virtual environments would be helpful. Translating existing tasks and windows into a new virtual environment presents challenges around usability, interaction, and performance, and the impacts of these need to be evaluated.

This thesis addresses the impact of an immersive virtual environment when working with existing, familiar applications and interactions. Reading performance is identified as a critical task through observation of users at work, and the users' reading performance is measured in an immersive virtual environment developed for the experiment; and compared to performance on conventional displays. It is shown that with careful setup, reading tasks can be performed at the same level as on conventional displays, and that under less ideal circumstances accuracy of reading can be maintained at the expense of speed. Font attributes identified as potential candidates for improving text legibility are investigated, and the presence of ClearType font smoothing is demonstrated to have beneficial effects on the distance at which text can be read within the presented environment. Finally, informal feedback from users of the developed environment suggests that the use of conventional mouse and keyboard controls translates well into a spherical virtual environment, although not being able to see the devices was of concern.

Acknowledgements

For my supervisors Mark Apperley and Bill Rogers; thank you for your support over the past years, and especially for the support over the last few months. Despite my insistence on being one of the more writing averse students of your careers, you have stuck with me and hauled me over the line. I have no doubt that without your help and enduring positive outlooks, that I would not have made it to submission.

To my friends and colleagues who have listened to me talk about my research, writing, and numerous other gripes over the years; you have helped keep me grounded and sane throughout this experience. Special thanks to Jemma König and Andreea Calude for putting up with the brunt of it, and keeping me well caffeinated. I dread to think of the volume of coffee that we have collectively consumed over the years.

Thank you to the University of Waikato for funding this research through the Waikato Doctoral Scholarship, and to the Division of Health, Engineering, Computing, and Science for additional grant funding. This support has allowed me to pursue something I never thought I would be in a position to do, and I am grateful that you have given me this opportunity. I would like to especially thank the Department of Computer Science and all of my co-workers within it. Your support in helping me find time to complete this writing even with the associated difficulties is above and beyond what I personally feel I deserve.

Finally, to my wife Olive. We met for the first time the week that I first began this research, and we have since gotten engaged, bought a house of our own, married, and welcomed our daughter Darcy to our family. This research has been a constant, and often burdensome companion to both of us these past years, but it is finally done. This has been a long road, and I could not have asked for a better companion while walking it.

Contents

Abstract iii

Acknowledgements v

Front Matter vi

List of Figures xii

List of Tables xiv

Program Listings xvi

1 Introduction 1

1.1 Research Questions 3

1.1.1 Conventional Computing 3

1.1.2 Evaluation Metric 4

1.1.3 Environment Performance 5

1.1.4 Software Configuration Changes 5

1.2 Structure of this Thesis 5

1.3 Contributions 7

1.4 Notations 7

1.5 Publications 8

1.6 Ethical Consent 9

2 Traditional Display Organisation and Window Management 10

2.1 Conventional Computing Systems 10

2.1.1 Personal Computing Systems 11

2.1.2 Typing Devices 12

2.1.3 Pointing Devices 13

2.1.4 Display Technology 14

2.2 Graphical Interfaces 15

2.2.1 Window Managers 16

2.2.2 Document Interfaces 20

2.2.3 Assisting Window Switching 23

2.2.4 Assisting Window Arrangement 26

2.3 Increasing Space 28

2.3.1 Large Displays 29
2.3.2 Multiple Displays 30
2.3.3 Task Performance Benefits 32
2.3.4 Downsides & Considerations of Increasing Space 33
2.4 Summary & Conclusions 34

3 Workplace Observation 35
3.1 Study Procedure 36
3.1.1 Participants 36
3.1.2 Preliminary Demographic & Hardware Survey 37
3.1.3 Periodic Observations 37
3.1.4 Closing Interview 38
3.1.5 Data Treatment 39
3.2 Preliminary Survey Results 40
3.2.1 Demographic 40
3.2.2 Display Number 40
3.2.3 Display Sizes 42
3.2.4 Display Resolutions 43
3.2.5 General Observations 43
3.3 Periodic Observations Results 44
3.3.1 Concurrent Tasks 44
3.3.2 Windows Count 47
3.3.3 Window Arrangements 48
3.3.4 Screen Time 50
3.4 Closing Interview Results 51
3.4.1 Daily Usage 52
3.4.2 Window Identification 54
3.4.3 Physical Environment Improvements 55
3.4.4 Digital Workflow 57
3.5 General Observations 58
3.5.1 Lifecycle Management 58
3.5.2 Window Layout 59
3.5.3 Window Identification 60
3.5.4 Content Consumption 61
3.6 Summary 62

4 Virtual Reality Display Technologies 63
4.1 Virtual Reality Environments & Systems 64
4.1.1 Virtual Environments 65
4.1.2 Virtual Reality Systems 66
4.1.3 Head Tracking Systems 72

4.1.4 System Selection 75
4.2 Virtual Environment Considerations 76
4.2.1 Content Presentation 76
4.2.2 Bounded or Free Windows 78
4.2.3 Environment Reference Frames 80
4.2.4 Environment Selection 83
4.3 Evaluating Suitability 85
4.3.1 Eye Physiology 86
4.3.2 Eye Movement 86
4.3.3 Saccades in Reading 87
4.3.4 Head Movements during Reading 89
4.3.5 Reading within HMDs 90
4.4 Summary 92

5 Virtual Environment Design 94
5.1 Design Requirements 94
5.1.1 Environment Design 95
5.1.2 Virtual Reality Device Target 95
5.1.3 Host Operating System 96
5.1.4 Other Designs Considered 96
5.2 Content Capture Exploration 97
5.2.1 BitBlt 98
5.2.2 PrintWindow 100
5.2.3 DWM Thumbnail API 101
5.2.4 DXGI Desktop Duplication API 101
5.2.5 General Optimisations 102
5.3 VR Environment Exploration 102
5.3.1 Planar Mesh 103
5.3.2 Curved Distortion Mesh 104
5.3.3 Other Meshes 108
5.3.4 Mesh Dimensions 108
5.4 Generation 1 109
5.4.1 Multi-window support 110
5.4.2 Mouse interaction 111
5.4.3 Keyboard interaction 113
5.4.4 Virtual Window Management 113
5.4.5 Hotkeys 115
5.4.6 Test Environment 117
5.5 Generation 2 117
5.5.1 Positional Tracking 118

- 5.5.2 Input Methods . . . 118
- 5.5.3 Run-time Configuration . . . 119
- 5.5.4 Testing Harness . . . 119
- 5.6 Discussion . . . 120
  - 5.6.1 Virtual Window Geometries . . . 120
  - 5.6.2 Input . . . 121
  - 5.6.3 Mouse Clamping . . . 122
  - 5.6.4 Window Capture Performance . . . 122
- 5.7 Summary . . . 123

6 Reading Study 124

- 6.1 Questions . . . 125
  - 6.1.1 Text Content . . . 125
  - 6.1.2 Virtual Window Size . . . 126
  - 6.1.3 Virtual Window Geometries . . . 126
  - 6.1.4 Presentation Distance . . . 127
  - 6.1.5 Device Preference . . . 127
- 6.2 Hypotheses . . . 128
- 6.3 Experimental Design . . . 128
  - 6.3.1 Equipment . . . 132
  - 6.3.2 Method . . . 132
- 6.4 Participants . . . 134
- 6.5 Data Processing . . . 135
- 6.6 Results . . . 137
  - 6.6.1 Preferred Distortion . . . 137
  - 6.6.2 Optimal Reading Distance . . . 138
  - 6.6.3 Inter-participant Reading Distance Consistency . . . 139
  - 6.6.4 Reading Performance Superior with Curved Distortion 141
  - 6.6.5 Conventional Display Outperforms Virtual Environment 143
  - 6.6.6 Device Preference Order . . . 146
- 6.7 Discussion . . . 147
  - 6.7.1 Display Distance . . . 147
  - 6.7.2 Distortions . . . 148
  - 6.7.3 AWPM Performance . . . 149
  - 6.7.4 Reading Accuracy . . . 150
  - 6.7.5 Legibility of Text . . . 151
  - 6.7.6 Text Haloing . . . 153
- 6.8 Summary . . . 154

7 Tuning Text 156
- 7.1 Questions . . . 156
  - 7.1.1 Text Haloing . . . 157
  - 7.1.2 Fonts . . . 158
  - 7.1.3 Engine Distortions . . . 159
  - 7.1.4 Readable Distances . . . 159
  - 7.1.5 Device Improvement . . . 160
- 7.2 Hypotheses . . . 161
- 7.3 Experimental Design . . . 161
  - 7.3.1 Equipment . . . 166
  - 7.3.2 Method . . . 166
- 7.4 Participants . . . 167
- 7.5 Data Processing . . . 169
- 7.6 Results . . . 169
  - 7.6.1 Font . . . 169
  - 7.6.2 ClearType . . . 170
  - 7.6.3 Zoom . . . 171
  - 7.6.4 Fixed vs Free Distances . . . 172
  - 7.6.5 Device Improvement . . . 173
- 7.7 Discussion . . . 175
  - 7.7.1 Fonts . . . 176
  - 7.7.2 ClearType . . . 177
  - 7.7.3 Zoom . . . 178
  - 7.7.4 Acceptable Display Distances . . . 180
  - 7.7.5 Device Improvement . . . 181
  - 7.7.6 Time per Sample . . . 184
  - 7.7.7 Outliers . . . 184

8 Summary and Conclusions 187
- 8.1 Research Questions . . . 187
  - 8.1.1 RQ1: Conventional Computing . . . 188
  - 8.1.2 RQ2: Evaluation Metric . . . 189
  - 8.1.3 RQ3: Environment Performance . . . 190
  - 8.1.4 RQ4: Software Configuration Changes . . . 191
  - 8.1.5 Primary Research Question . . . 192
- 8.2 Limitations . . . 194
- 8.3 Future Work . . . 195
  - 8.3.1 Curving & Dancing Windows . . . 195
  - 8.3.2 Intelligent Text Extraction and Presentation . . . 196
- 8.4 Closing Remarks . . . 197

References **198**

Glossary **220**

Appendices **225**

Appendix A: Workplace Observation 225
Appendix B: Publications 233
Appendix C: Reading Study 242
Appendix D: Tuning Text 258

List of Figures

2.1 Desktop, mobile and tablet device market-shares 11
2.2 Example typing devices 12
2.3 Example pointing devices 13
2.4 Dell WD19S USB-C laptop dock 15
2.5 Tiling window manager layouts 17
2.6 Windows snap assist layouts 18
2.7 Stacking window manager layouts 18
2.8 Compositing window manager examples 20
2.9 Single document interface examples 21
2.10 Multiple document interface examples 22
2.11 Tabbed document interface examples 23
2.12 Taskbar examples from different operating systems 23
2.13 Task Switcher presentations 25
2.14 Zooming overview visual implementations 26
2.15 Desk and display sizes overlaid and compared 28
2.16 Display resolution trends 30
2.17 Multiple display example configurations 31
2.18 Display layout configuration 31

3.1 Concurrent tasks over time 45
3.2 Line plot of open windows over time 47
3.3 Line plot of open window statistics over time 48
3.4 Bar chart showing window organisation methods over time . . 49
3.5 Bubble chart showing window arrangements 50
3.6 Line plot of on-screen activity statistics over time 52

4.1 Milgram’s reality–virtuality continuum 65
4.2 Mann’s reality, virtuality, mediality continuum 66
4.3 Classification of alternate realities 67
4.4 Fiducial markers used by marker systems 73
4.5 Valve VR test room 74
4.6 Perspective and movability reference frames 81
4.7 Eye, head, and gaze positions while reading 89

5.1 Visualisation of a window obscuring another 99
5.2 Windows tiled for `BitBlt` 99
5.3 Planar mesh concerns, expressed as a right-angle triangle . . . 104
5.4 Curved distortion mesh geometry visualisation 105
5.5 2D visualisation of curved distortion meshes 106
5.6 Intersection of greater-diameter curved windows 107
5.7 Oculus Rift remote control 118

6.1 Participant demographic information 135
6.2 Interval plot for final selected virtual window distances 140
6.3 Biased AWPM percentage differences for samples 142
6.4 Text detail through headset lens 152
6.5 Magnified text on different subpixel layouts 152
6.6 ClearType fringing on letter ‘m’ 154

7.1 Oculus Rift remote control 162
7.2 View of experimental text sample and response buttons 163
7.3 Magnified view of a single word within a text sample 165
7.4 Participant demographic information 168
7.5 Readable distances of text across HMDs 174
7.6 Enhanced view of lossy text presentation 177
7.7 Effect of zoom on scaled text samples at near distance 179
7.8 Effect of zoom on scaled text samples at far distance 180
7.9 Population readability ratings for display distances 181
7.10 Screen-door effect visualisation 182
7.11 Independent variable font $\mathbb{D}_n$ and $\mathbb{D}_f$ values 184
7.12 Independent variable ClearType $\mathbb{D}_n$ and $\mathbb{D}_f$ values 185
7.13 Independent variable scale $\mathbb{D}_n$ and $\mathbb{D}_f$ values 186

List of Tables

2.1 Example CRT, LED, and LCD display details 14

3.1 Participant field and computing environment 40
3.2 Participant display details 41
3.3 Display details, grouped by diagonal size 42
3.4 Concurrent tasks per participant 45
3.5 Open windows over time 46
3.6 Window arrangements against number of open windows 49
3.7 On-screen activities over time 51
3.8 Daily restart responses 52
3.9 Environment reconfiguration time 53
3.10 Participant space and productivity responses 55
3.11 Digital workflow responses 57

4.1 Classification of Virtual Reality Systems 68
4.2 Saccade duration and distances for differing tasks 88

5.1 Oculus Rift device specifications 96

6.1 Captured window size details 130
6.2 Experimental sample sequence details 130
6.3 Participant distortion preferences 137
6.4 Chosen sample distance descriptive statistics 139
6.5 Paired T-tests for distortions pairs 141
6.6 Paired t-test for planar sample performance against baseline . 144
6.7 Paired t-test for curved sample performance against baseline . 144
6.8 Paired t-test for planar sample accuracy against baseline . . . 144
6.9 Paired t-test for curved sample accuracy against baseline . . . 145
6.10 Participant device preferences 146
6.11 Angular virtual window sizes 148

7.1 Comparative font choices 158
7.2 Tested attributes 163
7.3 Experiment and sample duration details 168

7.4 Paired t-test for font attribute 170
7.5 Paired t-test for ClearType attribute 170
7.6 Paired t-test for zoom attribute 171
7.7 Coefficient of variations for comfortable reading distances . . . 173
7.8 Readable distance descriptive statistics 174
7.9 Tested IVR HMD feature comparison 175
7.10 Population IPD statistics 183

Program Listings

5.1 Tracked window updating pseudocode 110
5.2 Ray–mesh intersection pseudocode 112
5.3 Virtual window update and proxy pseudocode 113
5.4 Window attribute updating pseudocode 114
5.5 Top-most window updating pseudocode 115
7.1 Stylesheet used when rendering text samples 164
7.2 Pseudocode showing experimental sequence calculation 165
7.3 Example results file 166

Chapter 1

Introduction

When focusing on professional workers and their computer use, it is clear that the number of tasks that a user is expected to manage and perform on-screen has increased over time. This increase leads to problems when it comes to the number of documents and windows they have open at any given time, with marked decreases in productivity observed as their workspaces become more cluttered. To help address this, different assistive technologies are included in the computer window manager to aid with organising and locating items, helping computer users better lay out their tasks, and to aid in switching. In addition, it is increasingly common to see multiple displays being attached to computers, increasing the physical space available upon which tasks can be displayed and arranged.

These approaches to aiding computer users in dealing with quantity of windows do not scale up indefinitely. Assistive technologies suffer from a lack of discoverability or interest amongst users, leading to a lack of awareness or usage of techniques that could be helpful. Increasing the number of displays that a computer user has access to seems to have a more direct impact on productivity, with measurable impacts and an easier adoption. The number of additional displays that a computer user can make use of is limited both by capability of the computer system and the available space to place the displays. Adding displays and the requisite hardware to drive them is an expensive undertaking, with graphics cards capable of managing larger numbers of displays commanding a price premium. These limitations mean that some users may be limited to a fewer number of displays than desired, limiting potential productivity improvements and satisfaction.

Despite computers still maintaining the conventional look, alternative display and input technologies have been investigated. While these have not been adopted into mainstream computing, there has been significant activity in the academic space right back to the initial adoption of home computing. An alternative to conventional display monitors that has been of particular interest is the use of virtual reality (VR) devices, which can offer virtual workspaces of varying types that promise more natural organisations of items, and increased space in which to arrange them. While promising results have been presented around the use of these devices, they have not taken over as primary output devices for computers.

Lack of general availability of virtual reality devices is commonly blamed for the lack of adoption, with the only items on the market being expensive, and of quality that lags behind the conventional monitors. In recent years however, immersive virtual reality (IVR) devices have been released that address these issues, with devices offering high resolutions, low weights, and relative affordability. While the primary focus of these consumer devices is entertainment, the increased availability and desirability reawakens the idea that these could be used as general-purpose computing interfaces that could replace conventional single or multiple display computing environments, whilst reducing overall costs and space requirements. If focus were to be placed on a seated virtual reality experience, existing office environments would be able to support changing to these devices with no change to workspace requirements.

VR devices for general computing have been investigated in the past, however priority has been placed on how such an environment could be created, and how users can engage in novel interactions within them. This research has also historically been limited by the specifications and availability of devices at the time. The area has seen renewed interest commensurate with the increasing availability and specifications of modern VR devices, but the focus of research remains similar.

If virtual reality devices did find a place in the general computing market, interfaces and applications designed with the advantages and disadvantages of the devices would be developed that would be used in preference to existing software. During adoption however, many applications would not immediately have such an interface or alternative software that could reasonably replace it. During the transitional period between the conventional computing space and a hypothetical virtual reality device interaction future, providing support for displaying and interacting with conventional applications with minimal or no modifications required by the software vendor would be helpful.

1.1 Research Questions

From the hypothetical transition, from desktop to immersive virtual environment, the central research question of this thesis emerges. What this thesis seeks to address in part is:

> ***To what extent can applications designed for conventional interfaces and displays be presented and used inside a general-purpose immersive virtual reality environment with only user preference modifications?***

To support research into this question and to limit the scope of investigation, several more specific questions need to be investigated and answered.

1.1.1 Conventional Computing

We have some idea of what computer hardware and software is used by computer users worldwide from various sources. The Steam Hardware survey releases data monthly that collects information about computer hardware in use by users of the Steam game distribution platform, but this data is limited for this question. It only collects information about users of the platform, which is an entertainment platform unlikely to be found on an office machine. It collects hardware information, but presents statistics over the information collected, and doesn't give a population breakdown. We can see for example the primary and multi display resolutions and how they rank, but we cannot see the proportion that utilise multi display setups. This limits the conclusions that we can draw with regards to how widespread use of such setups are.

Websites like w3schools collect statistics that visitor's web browsers present, but this is limited to aspects such as operating system, browser, and display resolution. The captured screen resolutions are broad, capturing only common resolutions and not seemingly capturing multi-monitor states. The data presented is also skewed towards users likely to need information presented on the website - software developers and students primarily in the w3schools case. Generalisations about populations are therefore hard to make with this data.

Other usage statistics are collected by commercial research firms but are restricted access due to paywalls. Additionally, these surveys are skewed towards a United States based audience, which is not necessarily applicable to New Zealand office work-spaces. Academic surveys also exist that investigate this, but typically focus on single-discipline populations.

The number of applications running on a given user's computer, the number of displayed windows, and the number of concurrent tasks being performed is also something that needs to be captured. As this varies between disciplines and timeframes, getting an overall statistic is difficult. When presented in aggregate, this prevents identification of specific interesting workflows that could be used as the basis for future investigations.

Collecting data points about the hardware and usage of computers in a multi-disciplinary New Zealand office space will allow us to capture examples that can be used as the basis for future decision making. Recording information about screen size, resolution, and number utilised will allow for comparison with other existing solutions, indicating if a change could result in an improvement in effective space. The number of windows and tasks in combination with management techniques used with these will allow us to draw conclusions about whether management techniques are sufficient or discoverable. Interviews with users will allow us to determine the sentiment of users around screen space and management, identifying if users desire more screen space or better management techniques.

From the above discussion, the first supporting research question can be formed:

Research Question 1. *What does a conventional computing workspace look like for modern office workers in New Zealand?*

1.1.2 Evaluation Metric

Using data collected in addressing Research Question 1, evaluation of observed user behaviours and utilised applications should be undertaken to identify common activities performed by all users across workflows and application. Following identification of these commonalities, these should be evaluated to identify the activity or task would have the greatest impact on task performance if required time was increased or decreased in an alternative environment. This identified metric can then be evaluated within an alternative environment to determine the impact on performance the environment has. This can be summarised with the second supporting research question:

Research Question 2. *What performance metric should be used when comparing applications within differing environments?*

1.1.3 Environment Performance

To determine the suitability of a virtual environment for displaying conventional applications, comparison should be made between a conventional and virtual environment displaying the same applications. The metric identified in Research Question 2 should be evaluated, recording the performance of the identified metric and comparing the results between implementations to determine if a change in environment adversely affected performance with respect to the metric. Possible reasons for performance differences should be identified, with conclusions drawn about whether these differences are resolvable through software configuration changes to the conventional applications, improvements to the virtual environment, or to improvements in the virtual reality hardware. This is summarised with the third supporting research question:

Research Question 3. *With respect to the chosen metric, to what extent does changing from a conventional environment to a virtual one affect performance?*

1.1.4 Software Configuration Changes

Using possible software configuration changes identified when addressing Research Question 3, further evaluation of the virtual environment should be undertaken with modifications made to these configuration options. For each user-modifiable software configuration change, it should be determined if changing this affected some measurable aspect of performance in the environment. This is summarised in the fourth and final supporting research question:

Research Question 4. *With respect to the chosen metric, what software configuration changes can be easily made to conventional applications that may improve performance within a virtual environment?*

1.2 Structure of this Thesis

In Chapter 2 (Traditional Display Organisation and Window Management), window management is identified as an important aspect of the conventional computing experience, and techniques and tooling seen within different environments are described and contrasted. Space management is identified as a primary driver of window management development and increases in physical display space is shown to have positive effects on productivity. Constraints on increasing physical display space are described, and it is shown that developed window management techniques are attempting to address these limitations. Chapter 3 (Workplace Observation) takes these observations and describes an observational study where New Zealand office workers are observed and surveyed about their current display spaces in both digital and physical forms. Observations of their space usage is conducted and demonstrates that many users do not leverage window management techniques effectively, and experience performance impacts due to lack of available space.

Chapter 4 (Virtual Reality Display Technologies) identifies that virtual reality environments and devices may allow for an increase in perceived display space, without requiring addition of more or larger conventional displays. An overview of the term virtual reality is given in terms of environments and systems, and suitability of these systems for use office worker display replacement is considered. IVR head-mounted displays (HMDs) are identified as devices of interest due to suitability and commercial availability, and virtual environment design considerations are discussed. Evaluation metrics are identified through analysis of common tasks performed by observed users, and reading is identified as the primary evaluation metric for later studies. Mechanics of reading are briefly explored, and factors that may limit reading within an IVR HMD are outlined.

Chapter 5 (Virtual Environment Design) describes the design and implementation of a general-purpose virtual environment developed as an evaluation platform. Approaches to content capture and presentation are discussed, and different generations of environment used in studies are described. Observations made by users experiencing the environment during development are presented, identifying considerations for future developments.

Chapter 6 (Reading Study) presents a user study measuring reading performance within a virtual environment viewed through different IVR HMDs. Conventional application windows of different dimensions are evaluated when displayed with different distortions, with reading performance contrasted against baseline readings from a conventional display. Curved window distortion is shown to offer consistent reading accuracy equalling baseline performance, with reading speed trailing behind baseline at larger sizes, and planar distortion panels are shown to have acceptable performance at lower sizes. Acceptable distances for reading are shown to vary between participants, though trends in acceptable distance are shown to be consistent across size and distortion of displayed windows. Observations made by participants identifying font choice and ClearType font smoothing as possible text attributes that affect readability are discussed, identifying these as attributes to be tested for impact on reading. Chapter 7 (Tuning Text) evaluates the effect that these, and other, attributes have on readable distances, collecting readability ratings on a three-point scale at fixed display distances. ClearType is shown to have positive effects on readability at greater distances when enabled, while font is not found to impact on readable distances for chosen font families. Zoomed interfaces scaled down to increase texture density is shown to negatively impact on readable distance, and discussion is given around this. Improvements in utilised IVR HMDs are presented, and the impact on readable distances across devices is shown to not correlate directly with display density or resolution.

Finally, Chapter 8 (Summary and Conclusions) summarises the contributions of this thesis, addressing each of the supporting research questions. The summary identifies potential topics for future research direction and describes recommendations for design of general-purpose virtual reality computing environments that may assist in future designs.

1.3 Contributions

The contributions of this thesis are as follows:

- Quantification of concurrent task numbers and associated windows managed by a sample group of New Zealand office workers performing typical office tasks, and window layout and identification strategies utilised when managing these.
- Quantified reading performance values for conventional application content within a virtual environment presented within an IVR HMD using running record techniques for accuracy of reading.
- Demonstrated benefits to reading within an immersive virtual environment when reading unmodified conventional application content from curved surfaces, identifying that utilising curved surfaces for wide windows allows for reading accuracy to match that seen on conventional displays.
- Established that ClearType sub-pixel anti-aliasing improves readability of text when utilised on conventional applications displayed within the virtual environment, allowing text to be considered readable at greater distances than when ClearType was not present.
- Identified that conventional mouse and keyboard input devices function acceptably within a spherical virtual environment, with mouse movements identified as mapping particularly well to rotational movements.

1.4 Notations

In this thesis several shorthand notations are used for both statistical values and study observation measurements.

Statistical notations used are accepted statistical notations that are widely used, presented here for reference

- $\mathbb{N}$, $\mathbb{N}(\cdot)$ represents a sample size, either as the whole or as a subset defined by $\cdot$.
- $\bar{x}$ represents the sample mean, being the average of all values in a sample set.
- s represents the sample standard deviation of a sample set.
- $s_{\bar{x}}$ represents the standard error of sample mean for a sample set.
- c_v represents the coefficient of variation for a sample set.
- Q_0 represents the minimum value within a sample set.
- Q_1 represents the quartile 1 value within a sample set.
- Q_2 represents the median value within a sample set.
- Q_3 represents the quartile 3 value within a sample set.
- Q_4 represents the maximum value within a sample set.

Study observation notations are non-standard and are defined as follows.

- $\mathbb{D}$ represents a distance within a virtual environment, expressed using in-engine units (IEUs).
- $\mathbb{D}_n$ represents a minimum distance boundary within a virtual environment.
- $\mathbb{D}_f$ represents a maximum distance boundary within a virtual environment.
- $\mathcal{T}$ represents a time measurement, expressed in seconds.
- $\mathcal{T}_s$ represents a time measurement for a single sample within a larger sample set, expressed in seconds.

1.5 Publications

Results of the experiment detailed in Chapter 6 was presented at the 15th New Zealand conference on Human-Computer Interaction (CHINZ) in 2015. The full text of the paper can be found in Appendix B.

Grout, Cameron, Mark Apperley, William Rogers and Steve Jones (2015). ‘Reading text in an immersive head-mounted display: An investigation into displaying desktop interfaces in a 3D virtual environment’. In: *CHINZ 2015: Proceedings of the 15th New Zealand Conference on Human–Computer Interaction*, pp. 9–16. ISBN: 9781450336703. DOI: `10.1145/2808047.2808055`.

1.6 Ethical Consent

All studies conducted during this research followed the ethical conduct in human research and related activities regulation of the University of Waikato. Applications for ethical consent were submitted to, and approved by, the Human Research Ethics Committee of the Faculty of Computing and Mathematical Sciences at the University of Waikato prior to commencement of related studies. Copies of approval letters from this committee can be found in Appendices A.1, C.1 and D.1, which pertain to the studies described in Chapters 3, 6 and 7 respectively.

Chapter 2

Traditional Display Organisation and Window Management

This chapter provides background on conventional computing systems, which inform observations and decisions made in further chapters of this thesis. The forms and physical appearance of computers is described with an emphasis on personal computer devices, and aspects of graphical user interfaces (GUIs) displayed on these devices are explored. Space management and methods of increasing available space is discussed, highlighting large- and multi-display configurations along with window management techniques as areas of increasing interest and adoption.

2.1 Conventional Computing Systems

Since the introduction of personal computers in the late 1970s, they have become commonplace in homes and workplaces, with the expectation in modern times that individuals will have access to at least one computing device on a regular basis. As these devices have become more ubiquitous, so too has performing tasks using them, with many previously manual tasks being automated using these, and physical representations of task artefacts being supplanted by digital equivalents or replacements. Despite this massive increase in usage and evolution of computing technology, many modern computing devices available today would be easily recognisable to users of early graphical computers.

In this section, physical aspects of modern conventional computer systems will be discussed, identifying common forms, input devices, and output devices used with these systems. This will form the definition of the term “conventional computing system” that will be used throughout this thesis.

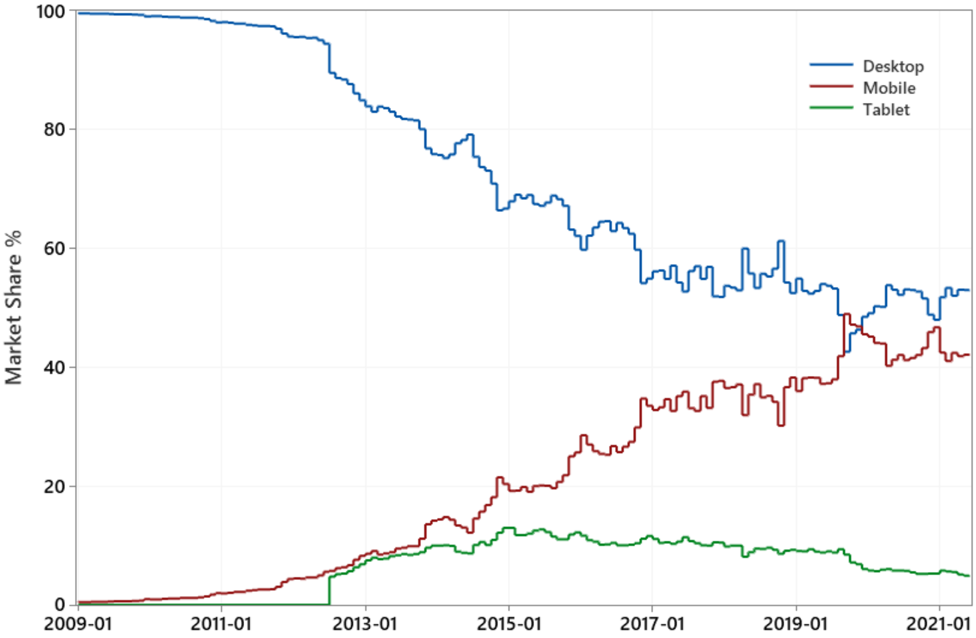

Figure 2.1: Desktop, Mobile and Tablet device market-share from 2009-01 through to 2021-06 (Statcounter Global Stats, 2021).

2.1.1 Personal Computing Systems

Modern consumer computing systems come in many different forms that serve different markets and purposes, including smartphone and tablet mobile computing devices, desktop workstations, and laptops. These devices are often grouped into two categories, those of personal computer devices which includes laptop and desktop computers, and post-pc devices which encompass mobile devices such as tablet computers (Murphy, 2011). When considering general purpose computing tasks performed by home and office users, personal computers are commonly seen, with post-pc devices becoming increasingly prevalent in the years since their mainstream introduction into the consumer computing space. While post-pc tablet and smartphone devices represent a significant market share of devices, accounting for 47% of the general computing market within New Zealand as of June 2021 (See Figure 2.1), this thesis will exclude these devices from discussion, instead focusing on desktop and laptop personal computer devices.

The laptop and desktop forms of personal computer emerged early in the evolution and mainstream adoption of personal computers (Haigh & Ceruzzi, 2021, ch. 8), serving the portable computing market with laptop devices, and desktop computers serving single location use-cases. Despite differing hardware requirements when serving these different use-cases, many aspects of these devices remain in common, with input and output devices experiencing

Figure 2.2: Different forms of keyboard used on historic and modern personal computer systems; *(a)* the keyboard from an Apple II from 1977 (Encyclopædia Britannica, Inc, 2021), *(b)* a modern 104 key desktop computer keyboard, *(c)* a typical laptop computer keyboard from a 2020 Dell Latitude laptop, *(d)* an ergonomic split "ergodox" keyboard.

minor modifications to suit the application. Both category of device features displays onto which content is output, and keyboards which bear a strong resemblance their typewriter forebears (Noyes, 1983) from which typed input is received. Pointing devices are also ubiquitous on both devices (Kar et al., 2015), with the mouse and trackpad representing the most common devices in use for this purpose.

2.1.2 Typing Devices

The hardware seen on keyboard devices has evolved as alongside improved manufacturing processes and technological advances, however the cosmetic appearance and functionality has remained consistent with early devices. Variants of keyboard designs are commonly seen across a range of devices, with common changes omitting keys to conform to a different form factor or utilising different glyph layouts to support non-latin alphabet inputs. More extreme variations on keyboards are also seen, with enthusiast and domain-specialised keyboards taking very different forms to more commonly accepted designs (See Figure 2.2), and analogue switches being investigated as possible additions to existing keyboard technologies (Razer Inc, 2021; Dietz et al., 2009).

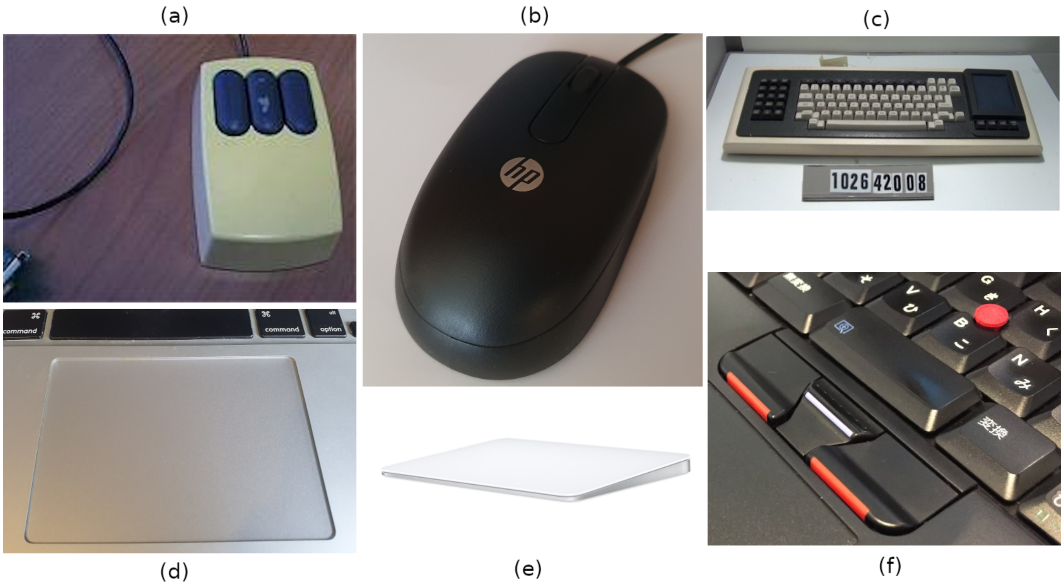

Figure 2.3: Different forms of pointing device seen on different form-factors of personal computer; *(a)* Xerox Alto I mouse from 1973 (Bordynuik, 2005), *(b)* typical 2-button mouse commonly seen with desktop computer systems, *(c)* early trackpad from a 1982 Apollo computer attached to a keyboard (Computer History Museum, 2021), *(d)* modern trackpad from a 2011 Macbook, *(e)* Apple Magic Trackpad device usable with a desktop computer, *(f)* pointing stick device found on some laptop devices.

2.1.3 Pointing Devices

Pointing devices have similarly seen many changes, with the development of trackpads and other variants of the traditional computer mouse being developed to meet different uses of portable computing as the most obvious of these changes. The primary difference between these two forms is the method through which the pointing is achieved, with a traditional mouse detecting physical movement of itself relative to a surface, and a trackpad tracking the movement of finger across its surface, with the cursor position tracking the finger movement on the screen (Kar et al., 2015). As technology has improved, capabilities and technologies used with these devices have changed, with mice moving from mechanical rolling ball sensors to optical sensors for movement detection, and trackpads introducing gesture and pressure sensitive touch capabilities. Despite these changes though, the operating principles and core appearances of these pointing devices has remained common with early devices (See Figure 2.3).

Table 2.1: Example displays, identifying the viewable diagonal size in inches, the native resolution in pixels, physical dimensions given as $W \times H \times D$ in millimetres, and weight in kilograms for selected models. Dimensions and weight include supplied stand for each monitor.

Type	Model	Size	Native Res.	Dimensions	Weight
CRT	Dell E551	15”	1024 × 768	360 × 388 × 384	12.5
CRT	Mitsubishi DV997FD	18”	1280 × 1024	447 × 440 × 449	21.3
CRT	Phillips 202P45/94	22”	1600 × 1200	501 × 501 × 466	29.6
CRT	Sony GDM-W900	24”	1920 × 1200	582 × 502 × 549	40.5
LCD	Dell E178WFP	17”	1440 × 900	404 × 338 × 137	3.6
LED	Dell U2718Q	27”	3840 × 2160	612 × 409 × 200	6.2

2.1.4 Display Technology

Display hardware has also seen changes both in the technology used, and in the forms taken. When considering the conventional desktop display, the most apparent form change is seen with the change from cathode-ray tube (CRT) displays to thinner and lighter liquid-crystal display (LCD) or light-emitting diode (LED) displays, moving from large bulky convex devices to thin light and predominantly flat displays. Laptop displays have historically utilised LCDs due to their portable form factor (Sarma & Akinwande, 1996), though early displays were of limited size and quality when compared to their desktop cousins. These have benefited from improvements in display technology, with laptop devices now sporting large high-resolution displays in a portable form-factor. Due to the switch from CRT to lighter and more compact LCD and LED technologies, the physical size and weight of displays was significantly reduced (See Table 2.1), allowing for larger display surface monitors to be placed on desk surfaces which may not have accommodated the weight or dimensions of a similarly sized CRT. Coupled with reductions in display panel production costs, the adoption of larger higher resolution displays has increased (VALVE Corporation, 2021; Robertson et al., 2005).

Increasing capabilities of graphical computing hardware within computer systems has allowed the usage of multiple displays simultaneously, which is seen in increasing numbers amongst users (Robertson et al., 2005). Modern computers are commonly able to connect 2 or more displays, allowing for a “continuous interaction space” (Pirchheim et al., 2009) to be formed through allowing the computing environment to span multiple discontinuous displays. This ability can increase the working area available to the user when spanning displays and improve the ergonomics of presenting content when used in a mirror configuration. Laptop computers are still commonly seen sporting a

Figure 2.4: A Dell WD19S USB-C laptop dock, showing connectivity options available to connected laptops. Dell Inc., 2021

singular display due to their form factor, however multiple display models are available (ASUSTeK Computer Inc., 2021), and range of portable low-power displays can be utilised to provide multiple displays in mobile computing situations (Mobile Pixels, 2021). Additionally, laptop docking stations may be used at fixed locations, allowing for the the connection of additional displays and other peripherals when connected (See Figure 2.4).

2.2 Graphical Interfaces

As with physical aspects of computing, many aspects of graphical interfaces will remain familiar to users of early graphical systems running on personal computers. From early in development of graphical interfaces, the windows, icons, menus, and pointer (WIMP) model has become the standard method of representing graphical information on a personal computer (Dam, 1997; Kyritsis et al., 2016). The windows within the WIMP paradigm are used to implement the desktop metaphor of computer interaction, where documents can be arranged within a display space in a manner similar to how physical items would be arranged on a desk surface (Stadler & Lorenz, 2007). These digital documents are augmented through menus which organise and control transforms of documents, and interaction is accomplished with pointers which are in tern controlled by physical pointing devices. Icons are used to represent unopened documents, tools, and actions that can be performed to opened or unopened documents as appropriate. While many “post-WIMP” paradigms — interfaces which forgo menus, pointers, and other WIMP features in favour of alternative operational specifications including voice command and gesture navigation (Dam, 1997) — have been discussed for personal computing (Jacob et al., 2008; Jetter et al., 2013), adoption of these has been seen primarily on post-pc devices (Jo et al., 2017), leaving WIMP and the desktop metaphor as ubiquitous in the personal computer space.

Many different approaches to supporting these paradigms have been seen over the years, and significant investigation and development has gone into improving the user experience. One area of particular interest is the representation and management of windows within a graphical computing environment, investigating and implementing many different approaches to different aspects of windows to support common and specialised workflows. In the following subsections, system window managers will be highlighted, describing different approaches to window organisation on a system level, individual application window usages will be identified, and supporting window management tooling will be described.

2.2.1 Window Managers

System window managers are components of operating system graphical environment subsystems responsible for the appearance and layout of individual windows within the graphical user interface (GUI) (Myers, 1988). While window managers can provide many facilities and perform a wide range of tasks, the primary task associated with them is that of window arrangement (Robertson et al., 2000), supporting positioning, resizing, and controlling the visibility of application windows. There are several different approaches taken with task that have seen use in major operating systems throughout their development, which each has different priorities and levels of responsibility for the user.

Tiling Window Managers

Tiling window managers arrange windows on the screen in a non-overlapping fashion, aiming to make maximum use of available display space (Goodfellow, 1986). When a single window is displayed, it occupies the full extent of the display in a “full screen” fashion, and multiple windows are displayed, the display space is subdivided into mutually non-overlapping regions into which each window is placed (Laukkanen, 2011). When subdividing the display space areas are not necessarily equally sized, allowing for smaller or less important windows to be positioned in a less obtrusive fashion, while larger

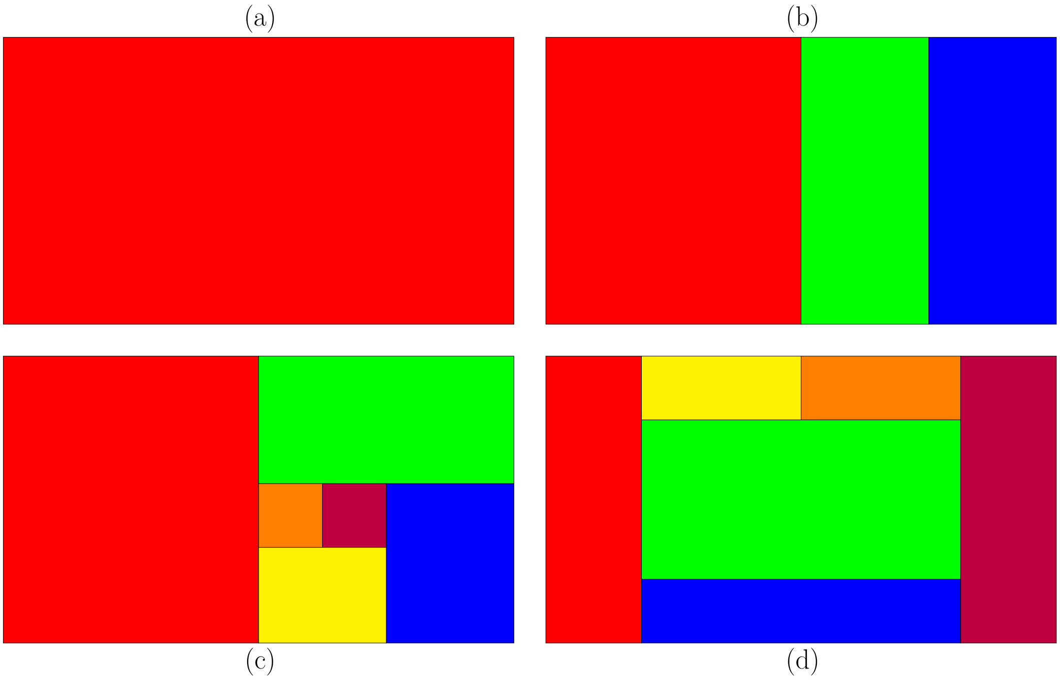

Figure 2.5: Example tiled window layouts, with each coloured region representing an individual window placement; *(a)* single full screen window, *(b)* horizontal layout, *(c)* spiral arrangement of decreasing window sizes, *(d)* theatre layout.

or more important windows can take up a greater proportion of the display (See Figure 2.5). Subdivision of display space can be performed manually or automatically (Cohen et al., 1986), and layouts can often be switched between rapidly to adjust for changing workflows.

Tiling window managers were utilised as primary window managers in several early graphical operating systems, including Microsoft Windows 1.0 (Goodfellow, 1986), and still see use in GNU/Linux systems where many tiling window managers are available (XMonad Developers, 2021; AwesomeWM Project, 2021). Tiling window managers have been noted as improving task performance under certain workloads when compared to other approaches (Bly & Rosenberg, 1986; Zeidler et al., 2013), however users have expressed preferences for the alternative stacking approach for general computing tasks (Zeidler et al., 2013). While use of exclusive tiling window managers is not common with modern computer systems (Laukkanen, 2011), most window managers support arrangement of windows in a tiling fashion either as once-off arrangements that simply resize and reposition windows, or as persistent layouts that respond to resizing events in a fashion similar to dedicated tiling window managers (See Figure 2.6).

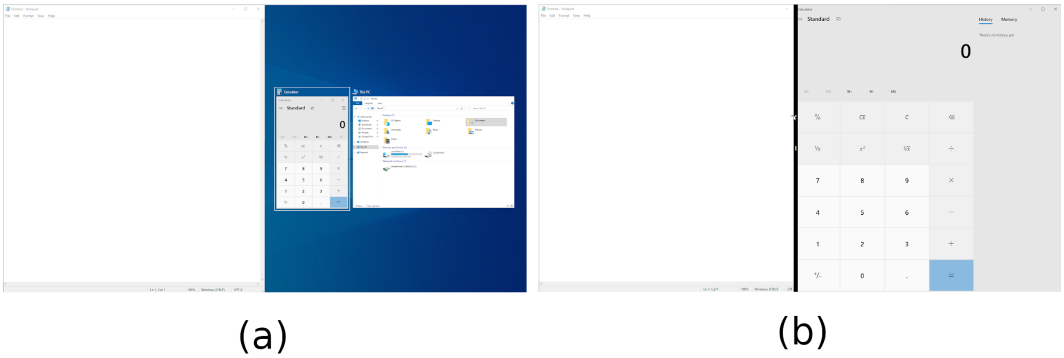

Figure 2.6: Windows snap assist tiling features in Microsoft Windows 10; *(a)* half-screen tiled window with neighbouring window selection prompt, *(b)* tiled window resizing of both tiled windows.

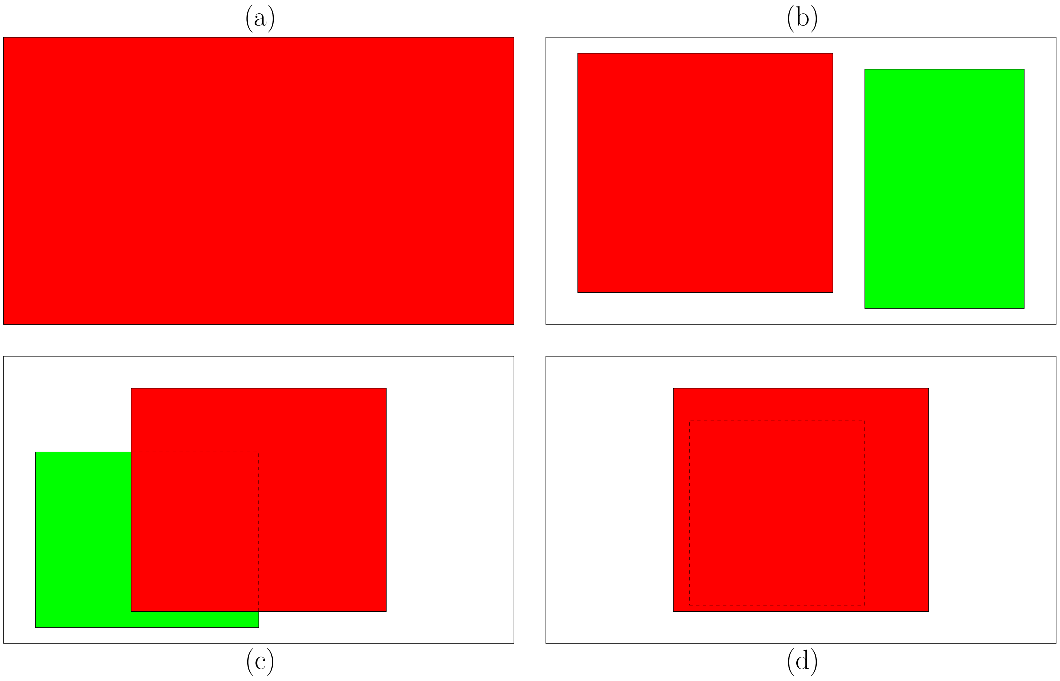

Figure 2.7: Example of stacking window layouts, with coloured regions representing individual windows, and dashed lines denoting edges of obscured windows; *(a)* single full screen window, *(b)* two mutually non-overlapping windows, *(c)* a foreground window partially obscuring another window, *(d)* a foreground window completely obscuring another window.

Stacking Window Managers

Stacking window managers support free-form arrangement of windows, allowing complete or partial overlapping of windows within the main display area (Zeidler et al., 2013), and partial placement of windows outside of the display space . This method of arrangement is analogous to arrangement of physical documents on a surface, where they can be arranged in stacks or spread out, allowing occlusion of documents to hide unneeded content (Agarawala & Balakrishnan, 2006). Windows can be displayed in a full screen fashion to fill the available display area, or as freely sizeable windows that can be arranged alongside others. Organisation of window layouts is a largely manual process requiring user pointer inputs to adjust size and position, but does allow for useful organisational layouts of windows, and effective use of space through tactical obscuring of related windows (Hutchings & Stasko, 2004a).

Windows presented within a stacking window manager are presented to the screen in a depth-sort order according to their z-order, referred to as a painter's algorithm, with the window lowest in the order drawn first, and the top-most window last (Berg et al., 2008). This approach requires cooperation of all windows when refreshing the display, and poorly performing applications can lead to undesirable visual artefacts in the presented view. Many graphical operating systems have made use of stacking window managers historically (Deurzen, 2019) with some still available and in use (fluxbox et al., 2021).

Compositing Window Managers

Compositing window managers are a class of window manager that provide off-screen buffers into which applications draw window content, with these buffers composited into an image which is then displayed (Ritger, 2006; Savidis & Maragudakis, 2013). This approach of providing each window its own dedicated graphical surface reduces the impact that poorly performing windows have on others and allows for additional transforms to be performed on window content to support alternative representations of content without the drawing application being aware of the transform (Wimmer & Hennecke, 2010). This ability to transform content allows for less traditional representations of windows in 3D environments and presents options for visual effects that support the user experience in a variety of ways (See Figure 2.8). Both tiling and stacking approaches to window management can be accomplished using compositing managers, however the description is typically applied to window managers that utilise stacking. Modern graphical operating systems have been making

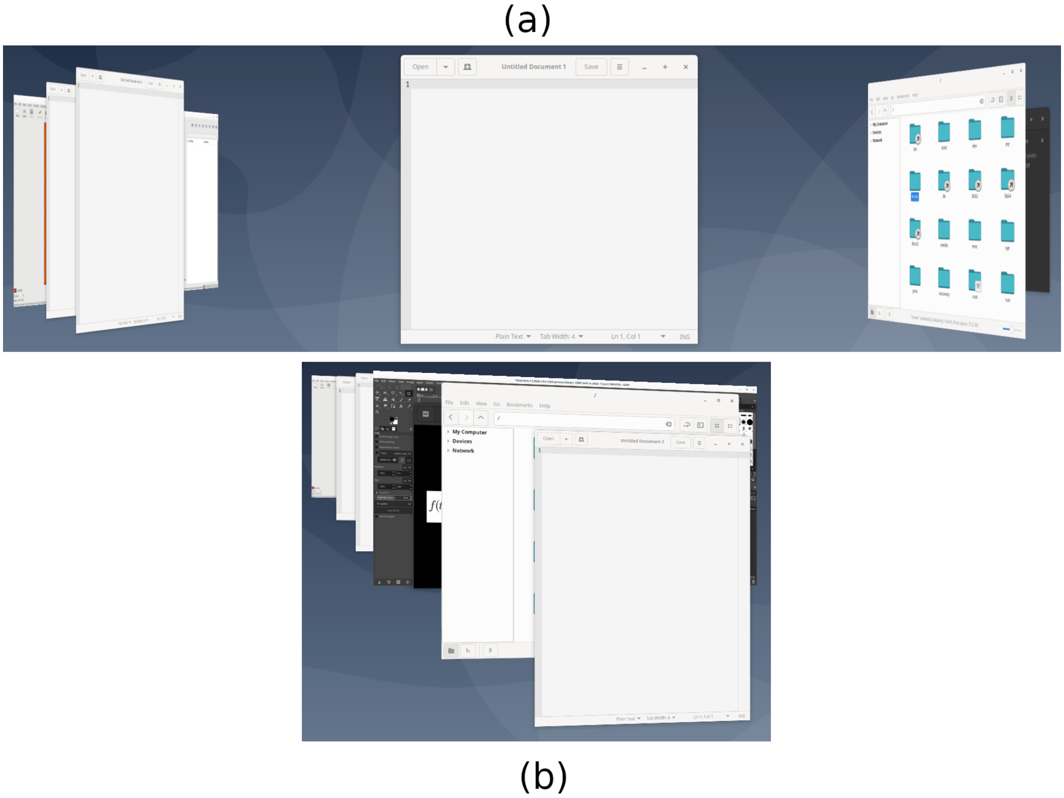

Figure 2.8: Compositing window manager feature demonstrating windows applied to 3D surfaces; *(a)* a cover-flow task switcher, *(b)* a timeline task switcher.

use of compositing window managers for some time, with Apple macOS beginning with OSX Jaguar in 2001 (The Trustees of Indiana University, 2021) and Microsoft Windows beginning with Vista in 2006 (Microsoft Corporation, 2018a).

2.2.2 Document Interfaces

Applications utilise several distinct approaches, or combinations of approaches, to managing windows that were created and controlled by them, which are categorised as document interfaces. These broadly fall into two categories, those of single document interfaces (SDIs) or multiple document interfaces (MDIs) (Sadiq & Pirhonen, 2014), with MDI encompassing several specialisations of document interface, most notable of which is the tabbed document interface (TDI).

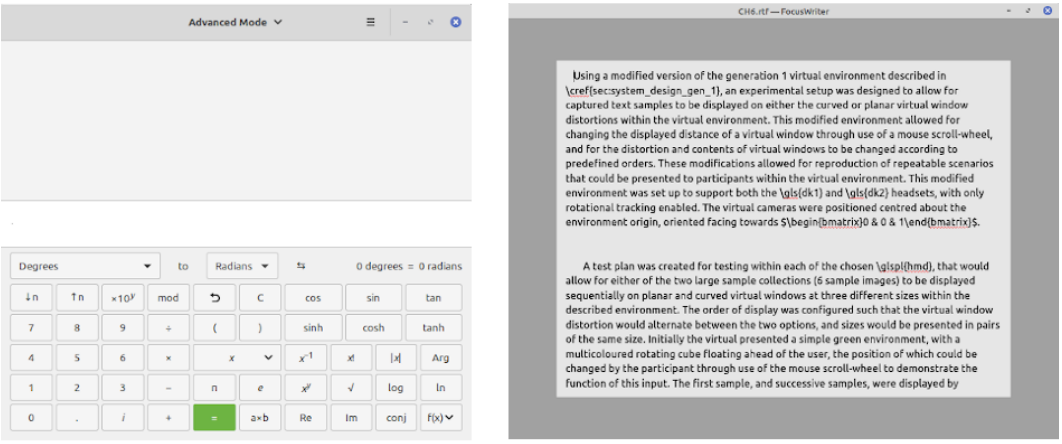

Using a modified version of the generation 1 virtual environment described in \cref{sec:system_design_gen_1}, an experimental setup was designed to allow for captured text samples to be displayed on either the curved or planar virtual window distortions within the virtual environment. This modified environment allowed for changing the displayed distance of a virtual window through use of a mouse scroll-wheel, and for the distortion and contents of virtual windows to be changed according to predefined orders. These modifications allowed for reproduction of repeatable scenarios that could be presented to participants within the virtual environment. This modified environment was set up to support both the \gls{dk1} and \gls{dk2} headsets, with only rotational tracking enabled. The virtual cameras were positioned centred about the environment origin, oriented facing towards $\begin{bmatrix}0 & 0 & 1\end{bmatrix}$.

A test plan was created for testing within each of the chosen \glspl{hmd}, that would allow for either of the two large sample collections (6 sample images) to be displayed sequentially on planar and curved virtual windows at three different sizes within the described environment. The order of display was configured such that the virtual window distortion would alternate between the two options, and sizes would be presented in pairs of the same size. Initially the virtual presented a simple green environment, with a multicoloured rotating cube floating ahead of the user, the position of which could be changed by the participant through use of the mouse scroll-wheel to demonstrate the function of this input. The first sample, and successive samples, were displayed by

(a) (b)

Figure 2.9: Single document interface example applications; *(a)* calculator application, *(b)* simple text editor.

Single Document Interface

SDIs take the form of independent and self-contained windows, with each window displaying a document which contains all information and controls within the window container. This approach to user interface ensures that when the window is made visible, both the content and controls for the window are visible without requiring additional window management operations and allows the system window manager to handle layout and display of the window in a familiar manner (Sadiq & Pirhonen, 2014). SDIs typically manage and display a single document per window and require additional application instances to be launched with their associated windows if multiple documents are to be displayed.

Multiple Document Interface

MDIs allow for application instances to own multiple windows, that may be presented to the system window manager independently or contained within a frame window (Galitz, 2007, p. 401) (See Figure 2.10). In the former case, independent windows can be managed by the system window manager, allowing them to be arranged and positioned in accordance with the window manager approach, leveraging familiar window interactions. In the latter case, sub-windows can be created which are constrained to a parent frame window (Sadiq & Pirhonen, 2014), existing only within the bounds of the frame window, and the parent window becomes responsible for providing management techniques that apply to the application domain. sub-windows are typically unaware of the position of the parent frame window, allowing for the frame

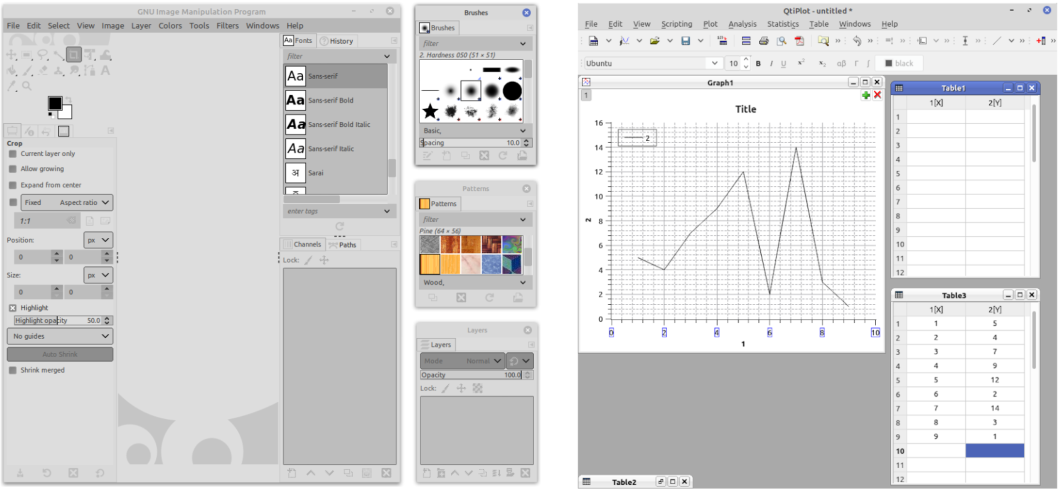

(a)

(b)

Figure 2.10: Multiple document interface example applications; *(a)* utilising independent windows, *(b)* managing windows within a parent frame.

window to be arranged within the system window manager without changing the sub-window positions or states relative to the frame window. MDIs frame windows can allow for a large virtual management space so that different workspaces of sub-windows can be managed as needed.

Tabbed Document Interface

TDIs are a specialisation of MDIs, allowing multiple sub-windows can be managed by a single container window, but restricting display of a single sub-window at a time in a full screen manner rather than allowing free-form arrangement (See Figure 2.11). As with other MDIs, management of sub-windows is performed by the parent frame window, with a typical interface providing a list of tabs, often temporally ordered (Chang et al., 2021), which can be interacted with to reorder content, and allow for the displayed document to be switched (Jakobsen & Hornbæk, 2010). This style of interface has found popularity in web browsers and text editors, where having many documents open simultaneously is desirable, but only one document needs to be visible at any given time. A major drawback of this form of interface is observed when switching documents, with “tab clutter” identified as an issue amongst the majority of browser users (Kulkarni et al., 2019) leading to many users struggling to find tabs they need (Chang et al., 2021).

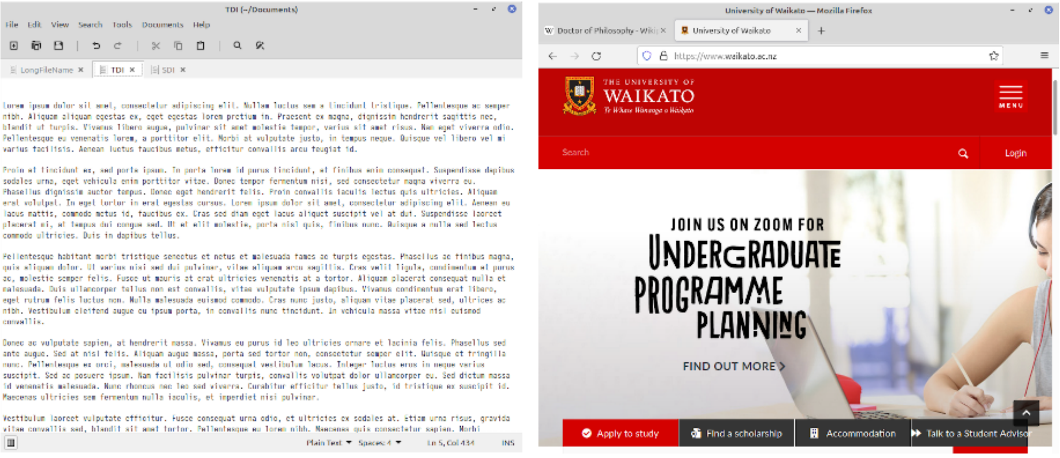

(a) (b)

Figure 2.11: Tabbed document interface example applications; *(a)* programmer text editor, *(b)* Mozilla Firefox web browser.

(a)

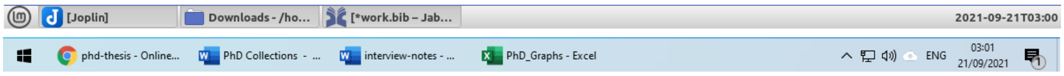

(b)

Figure 2.12: Taskbars from different operating systems showing from left-to-right, application launchers, program lists, and widgets; *(a)* a GNU/Linux Cinnamon DE taskbar, *(b)* a Microsoft Windows 10 taskbar.

2.2.3 Assisting Window Switching

Switching between windows is a common task, allowing for users to bring forward obscured content or to change the active window that receives input. In a tiled window manager, the desired window can simply be visually identified and selected, but with stacking window managers that may have partially or completely obscured windows this may require shuffling through windows to identify the target (Warr et al., 2016). While this could be accomplished by moving windows around the screen analogous to moving papers on a desk, many different assistive tools have emerged to make the process quicker and easier. Of the available options available tools, taskbars, task switchers, and zooming overviews are now present on all modern operating systems in some form, and present useful assistive approaches to window switching.

Taskbar

A taskbar is a control displayed within desktop environments that displays many different elements, including applets for time or weather, shortcut icons for rapidly launching applications, and program lists from which documents and applications can be accessed (See Figure 2.12). A primary element displayed within a taskbar is a window list, which displays a list of windows present on the system (Hutchings et al., 2004), represented as icons and optionally textual window information . Entries in this list can be presented individually or grouped by application depending on configuration (Microsoft Corporation, 2021c). When searching for a target window, users can visually scan the list of entries to identify possible candidates without seeing the windows themselves and bring each to the foreground by selecting the entries. Maintaining a list of windows in this fashion allows windows to be minimised, removing them from view, without losing easy access to them when they need to be accessed in the future. It also allows for window management controls, such as toggling full screen view or closing the window, to be managed directly from this list, often through a context menu, rather than requiring the window to be activated prior to this taking place.

Taskbar entries can be augmented with additional features to improve the performance of window identification, which may become difficult with many open windows in the list. Displaying thumbnail previews of window contents when hovering over a taskbar entry is a common feature provided by taskbar providers, allowing for visual inspection of windows near the taskbar entry prior to selection (Hutchings et al., 2004). An alternative to small thumbnail images which may lack sufficient detail is to temporarily activate the window represented by the window entry being hovered over, allowing for a full-sized view of the window in its original location.

Task Switcher

A task switcher is a common window switching feature found in modern operating systems, often known by one of its more well-known shortcut key combinations, “Alt-Tab”. Task switchers allows for switching between displayed windows ordered by their window manager z-order, starting with the window immediately below the currently active window. When used with a shortcut key, repeated toggling of this key combination will select each window in order, wrapping around to the top-most window once all options have been seen, until the desired window has been located. When a desired window has been located and selected, the target window moves to the top of the z-order, displaying above other windows and receiving input from devices. Visual representations

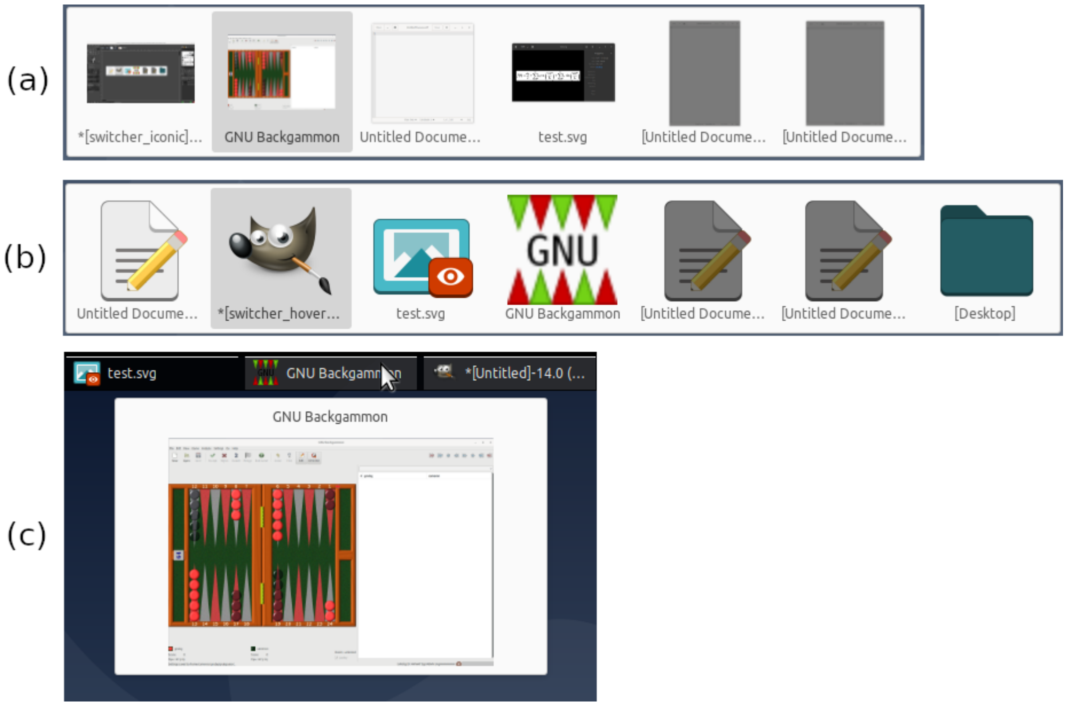

Figure 2.13: Different task switcher presentations; *(a)* ordered list of windows showing thumbnail previews and titles, *(b)* ordered list showing application icons and titles, *(c)* hover preview showing a thumbnail for an item in a window list.

of this window cycling vary with implementation (See Figure 2.13), with some systems temporarily activating windows as they are selected to display each window full-sized, while others display an icon or thumbnail list of windows with visual indicators for selected window. 3D versions are also seen, with examples such as Windows Flip-3D and the GNU/Linux Cinnamon DE task switcher options allowing for presentation of window previews in a timeline or album art cover-flow styles (See Figure 2.8), allowing windows that will be selected with subsequent activations to be seen behind the currently active window.

Zooming Overview

Zooming overviews utilise zooming user interface (ZUI) elements to obtain an overview of all open windows on a system simultaneously, allowing for 2-dimensional visual search rather than requiring sequential presentation of windows (Warr et al., 2016). This approach to window switching displays all open windows as thumbnails in an interactive view, ensuring that each thumbnail is presented without overlap, where the user may select the desired window to activate. This removes the requirement to repeatedly toggle a shortcut key to view each window, instead allowing for pointer input to select the target window following a single shortcut invocation. Exact visual implementations vary

Figure 2.14: Different zooming overview visual implementations, showing relative location and size thumbnails alongside a regular grid arrangement. Figure extracted from Warr et al. (2016).

for this feature (See Figure 2.14), with some approaches displaying thumbnails arranged and sized based on their relative location and size, and others using regularly sized thumbnails aligned to a grid ordered by window z-order (Warr et al., 2016).

2.2.4 Assisting Window Arrangement

When arranging windows within stacking window manager, the user has a large degree of control over the size, position, and visibility of each window. Users can optimise their display to emphasise window groupings, display multiple windows simultaneously, or focusing attention on a single window. Some of these optimisations are simple to accomplish, with full screening an application for singular focus often accomplished with a button on the window itself being a prime example. Other optimisations however may be tedious to accomplish with basic window controls, so additional assistive tools are present to provide simpler approaches to achieve the desired result. Of note are two technologies, window snapping and virtual desktops, which help address the optimisations of displaying multiple windows simultaneously and grouping windows respectively.

Window Snapping

Window snapping support is present in most modern operating systems and broadly describes two related utilities, the ability for windows to be aware of other window edges when resizing or moving, and the formation of tiled layouts with windows. Both of these utilities aim to simplify the process of displaying multiple windows simultaneously in a manner similar to that seen in tiling window managers.

For the former utility, windows undergoing resize or movement operations are made aware of other window edges, allowing for snapping behaviour to align moving edges with stationary ones. When moving or resizing a window, edges that intersect those from other windows briefly exhibit resistance at the point that the edges align, allowing for feedback on window alignment. This simple utility greatly simplifies edge-to-edge alignment of windows, allowing for manually tiled layouts of windows to be easily accomplished.

The latter utility extends that of edge snapping to provide automatic tiled window layouts of varying complexity when gestures, shortcut keys, or alignment options are invoked. Simple side-by-side window layouts, with two windows displayed each occupying half of the display space, were introduced as a native feature on the Microsoft Windows platform in 2009 and on the Apple macOS platform in 2015, though 3rd party solutions were available prior to these implementations (Mizage, LLC, 2019; Czarny, 2010). Newer implementations allow for quadrant or custom (Microsoft Corporation, 2021d) layouts, present prompts for windows to be added to the layout (See Figure 2.6), and link layouts that adjust neighbouring window sizes when others are resized to retain the layout at different sizes to provide increased layout flexibility.

Virtual Desktops

Virtual desktops are extensions to the window management system that "expand the space available for application windows by allowing users to switch between different workspaces" (Ringel, 2003), and were first introduced by Henderson and Card (1986). Each workspace represents a virtual display onto which windows can be arranged, and one virtual display at a time can be shown on the physical display attached to the system (Tomitsch, 2004). Windows can be assigned to one or more virtual desktop workspaces and will only appear within those to which they have been assigned. Utilising virtual desktops, windows can be grouped together with different partition strategies in contained regions where unrelated windows do not interfere (Tomitsch, 2004). Different partition strategies for allocation of windows to virtual desktops were observed by Ringel (2003), including grouping by task or sub-task, and utilising a primary/secondary display approach to avoid distraction with less important activities.

2.3 Increasing Space

A common factor that emerges when evaluating graphical interfaces is usage of finite space available on conventional displays, with window managers aiming to either completely fill this space with one or more windows or utilising the space to present organisational groupings analogous to physical paper arrangements. Tooling provided by window managers seek to improve the ease of arranging windows to fill this space and to locate windows that are hard to find due to layout constraints.

The space-constrained nature of the desktop metaphor has long been observed (Henderson & Card, 1986; Amir et al., 2021), with the difference in size between even large displays and their physical desktop counterparts being stark (See Figure 2.15). Henderson and Card (1986) used this difference as impetus for development of virtual desktop technologies that sought to use software tooling to increase available space but observed that increasing space virtually comes with its own drawbacks, specifically with navigation between workspaces and “simultaneous access to separated information” (Henderson & Card, 1986) when windows are distributed between spaces.

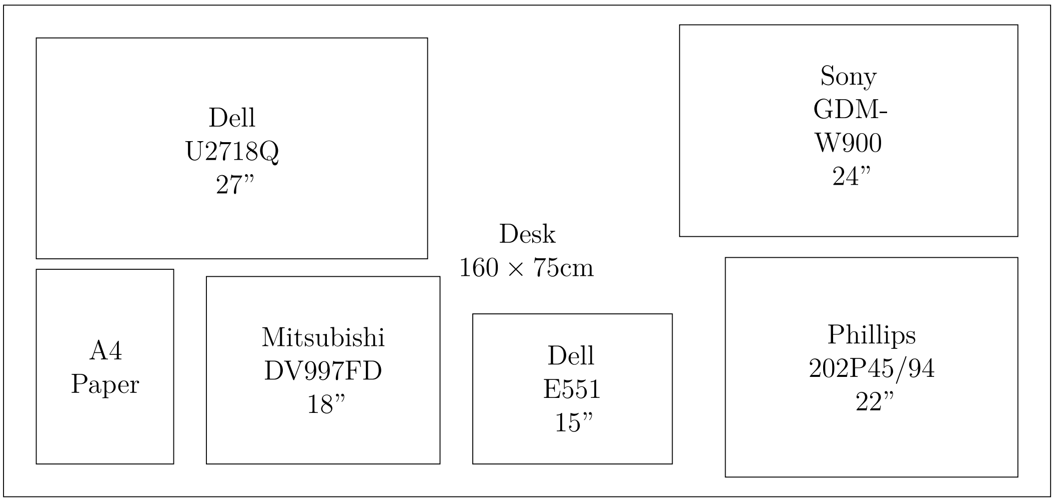

Figure 2.15: Superimposed outlines of displays given in Table 2.1, an A4 paper sheet, and a physical office desktop. It can be seen that even large displays have notably less surface area than a desk.

Utilising the software features of modern computing systems, numerous windows can be effectively managed by many users without significant task performance degradation, however it has been noted that “most windows management systems do not scale well to a large number of windows” (Tan et al., 2004). As the number of windows managed increases within the limited available space, the impact of window and task switching becomes more apparent, introducing significant time overheads over the course of a working day (Jeuris et al., 2018). This effect can be further pronounced due to accumulation of "window clutter" over time, caused by duplicate or no-longer needed application windows being left open, that interfere with window identification and lead to a "sense of information overload" (Wagner et al., 2012). This issue can be seen as a result of using insufficiently large displays, and can be summarised by the quote:

> "The impoverished environment of the single monitor forces users to make explicit context switches on the introduction of new information, frequently in the form of a new window overlaying the previous one. This severely affects the user's ability to make comparisons and requires the user to expend valuable mental resources on the minutiae of managing views rather than on the problem at hand" (Andrews et al., 2010)

A logical alternative to addressing the concerns of space through software is to modify hardware, physically increasing the space available for a user to leverage when arranging windows. When examining the state of conventional computing as it stands today, this is accomplished using larger displays, or multiple displays connected to a single system. In this section, these two approaches for increasing usable space are outlined, with benefits and drawbacks of each approach presented, and task performance benefits of increasing space are discussed.

2.3.1 Large Displays

As discussed in Section 2.1.4, large displays have become more available to users due to both reduced price (Klinke et al., 2014), and the smaller size and weight requirements of LED and LCD panels. Physically large and high-resolution displays can now be easily purchased, or constructed from arrangements of multiple smaller tiled displays (Phenix, 2017), allowing for significantly increased physical viewing, and digital rendering space.

Adoption of physically larger displays has trended upwards over the past years, with the default size of desktop displays moving from 15" to 21" in the 2000–2010 period (Bi & Balakrishnan, 2009). Usage of higher resolution displays can also be seen trending upwards, with 1920×1080 emerging as the most seen resolution since 2018 (See Figure 2.16). While this upward trend shows adoption of larger displays, the rate of change is low, suggesting barriers may be present that prevent users from adopting these newer devices. The most apparent of these may be cost, with physically smaller and lower resolution

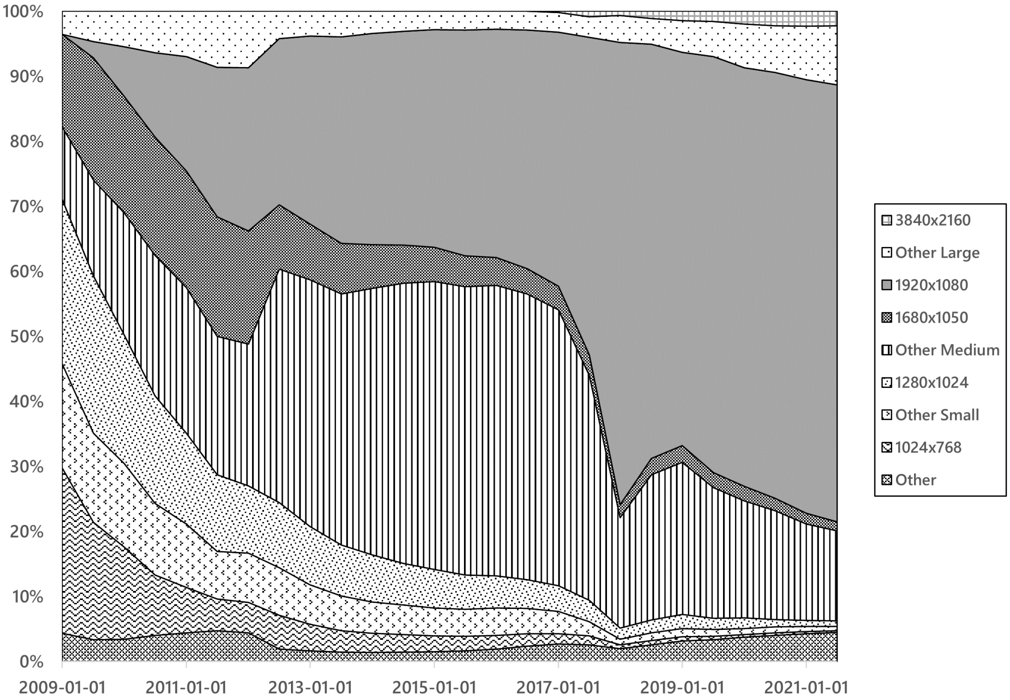

Figure 2.16: Trends in primary display resolutions utilised by users of the Steam game distribution platform (VALVE Corporation, 2021)

devices commonly being cheaper than larger higher resolution alternatives. A related explanation for slow adoption is the longevity of computer displays, with users utilising older still functional displays rather than purchasing new improved devices.

2.3.2 Multiple Displays

Modern computer systems are generally capable of connecting more than a single display for output, with the display area spanning across these displays, or mirroring content between them (Microsoft Corporation, 2021b). Many different multiple display configurations can be utilised, from single external displays, through to large format display alternatives formed from tiled monitors (See Figure 2.17). The popularity and utility of multiple display configurations has been noted over the years, with observations from researchers indicating that users prefer a larger number of larger displays if available (Owens et al., 2012), and suggestions that "multiple monitor systems [are] likely to become commonplace in the near future" (Hutchings & Stasko, 2004b) being raised.

When multiple displays are connected to a system, users are able specify the position and orientation of each display relative to other connected displays (See Figure 2.18), enabling the system to map movements of graphical elements between displays in a manner that aligns with the physical arrangement of the

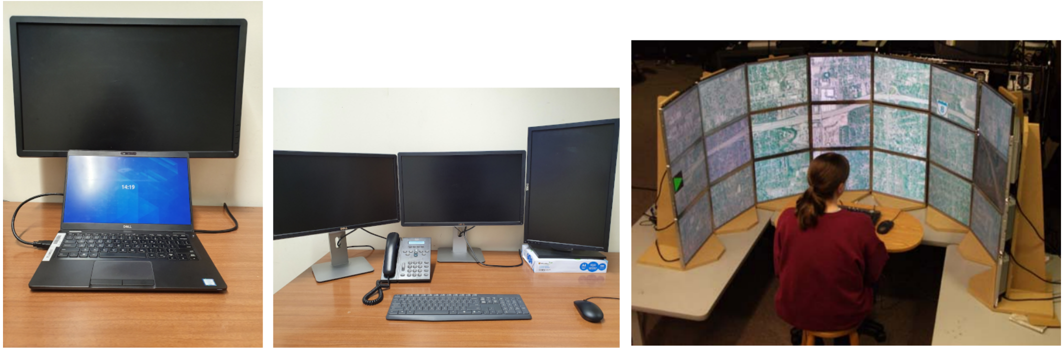

(a) (b) (c)

Figure 2.17: Multiple display configuration examples; *(a)* laptop device connected to a single display, *(b)* horizontal three-display configuration demonstrating different display orientations, *(c)* 24 display grid of displays in a curved configuration Shupp et al. (2006).

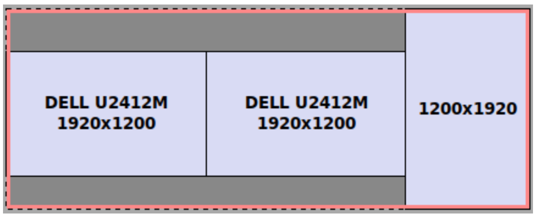

Figure 2.18: Display layout configuration presented by the NVIDIA X Server Settings on a Debian Linux system. Three displays are connected to the system arranged horizontally, with one display rotated to a portrait configuration. The pink border shows the rectangular frame buffer used by the system.

displays. Display subsystems record the positions of displays within a larger rectangular frame buffer, sized using a minimum-area enclosing rectangle to contain all display areas, which can be utilised by the windowing system and associated tools when performing layout operations. This presents an advantage for multiple display configurations over the single large display alternative, as window management operations can be performed within physically subdivided zones within a larger frame buffer, rather than utilising a single zone within a singular or tiled display.

There are several disadvantages observed with the use of multiple display configurations, some of which are common to those seen with large displays. Price of displays is of note when many are required, and ongoing costs through power consumption (McGill et al., 2020) can limit the accessibility of these solutions. Space requirements for multiple display configurations are also greater than single display configurations, requiring a larger workspace than might otherwise be used. When utilising displays arranged horizontally, a curved arrangement of displays may be required to effectively view content (Shupp et al., 2006), utilising user rotation in preference to translation to address each display. This may again impact on workspace size and arrangement or restrict the number of displays that can be effectively used. Curving arrangements of displays present their own drawbacks, with observations from many sources (Burruss et al., 2021; Gallagher et al., 2019) identifying required neck movements and postures as potential sources of discomfort and injury to users when addressing displays.

2.3.3 Task Performance Benefits

The use of large displays and multiple display configurations have been widely observed to improve task performance across a wide range of tasks when compared to smaller display setups. These benefits are seen in both single window and multi-window tasks, and in conventional 2D applications as well as 3D information-rich virtual environments (IRVEs).

When considering singular application tasks, Shupp et al. (2006) observed that larger viewport sizes provided by large or multiple display setups improve performance time and reduce user frustration when performing geospatial tasks when compared to single monitors, and Raptis et al. (2013) noted improvements in task efficiency for involved information retrieval tasks when performed on larger displays in a study utilising small-screen mobile devices. Experiments on the impact in performance larger displays have on IRVE search and comparison tasks have similarly observed benefits of larger displays, noting users became less reliant on way-finding aids and became more proficient with navigation tasks when using large displays (Ni et al., 2006). Spatial task performance has been shown to benefit from large displays, even if the display is positioned further away from the user, with Tan et al. (2003) demonstrating a large projected display with the same resolution as a desktop display showed 26% performance gains in tasks despite being positioned 110" further away from the user than the reference display.

When multi-window tasks are considered, multi-display configurations are widely acknowledged to convey significant benefits for users across a range of tasks, with users "unanimously prefer[ing]" (Bi & Balakrishnan, 2009) large display areas especially when working on multi-window tasks. These benefits are observed on large singular displays, large tiled displays, and multi-monitor configurations, indicating that the space provided by these larger environments conveys the benefits, rather than the specific approach to providing it. Increasing awareness of peripheral content through placing windows on secondary displays is noted as a positive benefit even when these secondary displays are significantly smaller than the primary display, with Grudin (2001) noting benefits of displaying calendar content on a small palmtop display in the user's peripheral vision. When comparing numbers of window switching actions and number of mouse clicks, users of multi-monitor configurations exhibited fewer actions than on single-monitor alternatives, which was preferred by study participants (Ling et al., 2016). For overall performance, Czerwinski et al. (2003) noted significant productivity and user satisfaction benefits when performing complex multi-window tasks on a large display, with 300 fewer window activation events recorded when compared to a smaller single display setup.

2.3.4 Downsides & Considerations of Increasing Space

While many benefits of using large and multi-monitor display configurations can be seen, there are observed disadvantages with certain configurations that influence on productivity, performance, and user comfort that should be considered. Of particular note is the presence of display bezels within the display area, and distractions that can emerge when greater space is available. Additional considerations such as required space and cost may also play a role in adoption of such solutions, with users limited by workspace or funds not being able to access these more beneficial configurations.

When making use of tiled large displays or multi-monitor configurations, the presence of bezels surrounding individual displays introduces visual separation between display areas. This has been observed to negatively impact on task performance in some cases, most noticeably in visual search tasks when targets are split across distinct displays (Forlines et al., 2006; Bi et al., 2010).

This observation does have counterpoints, with Wallace et al. (2014) recording negligible impact on visual search performance, noting that bezels may act as visual anchors supporting two-phase searching across the total area and within each bezel-bounded region. Truemper et al. (2008) observed that the multi-display configurations where the centre of the display area coincided with a bezel, such as 2 × 2 arrangements of displays, presented usability problems resulting in users preferring to utilise a horizontal or vertical subset of displays rather than all available space.

The effect of bezels may also impact on the usage of multi- or tiled-display configurations, with users reporting that they rarely position windows such that they straddle the bezels between two or more distinct displays (Grudin, 2001). Users instead prefer to constrain windows to individual display bounds, often identifying one display as the “primary” display on which the majority of work occurs (Dostal et al., 2013; Grudin, 2001) and identifying configurations that lack an identifiable primary display as uncomfortable (Truemper et al., 2008).

There is also evidence to suggest that benefits of increasing the number of displays used in multi-display configurations does reflect matching increases in task performance, with increases in available space providing diminishing returns (McGill et al., 2020) in performance gain and utilisation as the number of displays increase. Truemper et al. (2008) identified that users felt that four displays was too many, with feedback suggesting the increase in available space lead to greater attempted multi-tasking leading to distraction. Multi-tasking itself is widely recognised as having a negative impact on user performance (González & Mark, 2004; Hembrooke & Gay, 2003; Rubinstein et al., 2001), and distraction or interruption of tasks has been linked with higher stress and user frustration (Mark et al., 2008).

2.4 Summary & Conclusions

In this chapter, a background of conventional computing systems, input and output devices, and GUIs has been given, restricting discussion to laptop and desktop personal computers. Window managers and associated assistive tooling have been outlined, identifying techniques available to users for arranging and identifying windows. Approaches to physically increasing available space has been discussed, and it is show that increasing space has benefits for task performance and user satisfaction. Downsides of increased space have also been identified, suggesting that while benefits to larger computing spaces exist, there may be trade-offs involved if they are adopted.

Chapter 3

Workplace Observation

In Chapter 2, an overview of conventional computing was given, identifying common hardware and software features of computer systems used in homes and offices. Management of digital space was identified as a common concern, with window managers and assistive technologies seeking to make the process of window layout and identification simpler. Despite the addition of these software features, physical changes to computing hardware in the form of larger and more numerous displays are noted, suggesting a dissatisfaction with the performance or benefit of software solutions when dealing with modern workloads. Understanding the display configurations, window numbers, task loads, and usage of window manager assistive tooling may provide insight to this observation and indicate further investigative pathways for improving user experience.

This chapter describes a workplace observation study, in which users within a large New Zealand business were observed over the course of a working day to identify the quantity of windows they managed on their computing systems, and how they managed window arrangements. Details are given around the hardware used by users, and interview results are presented in which users were questioned about their digital working process, usage of assistive window management technologies, and perceived benefits of multiple displays for increasing space. This addresses Research Question 1, which stated:

> *What does a conventional computing workspace look like for modern office workers in New Zealand?*

It is identified that most observed participants engage in multiple simultaneous digital tasks throughout their working day, with each task involving multiple windows. Observations of daily work tasks and interview results identify that most participants struggle to manage the number of windows present on their system, with many not making use of window management techniques or available assistive window management tooling effectively. Interview results reveal that participants desire more physical screen space upon which to work in the form of larger or more numerous displays, despite the majority already making use of more than one display, and that they believe their productivity would increase if these desires were met. Finally, activities performed by all participants are discussed, identifying common application-independent tasks performed by all participants, providing a basis for addressing Research Question 2.

Supporting material for this chapter can be found in Appendix A.

3.1 Study Procedure

In order to establish a model of computer usage amongst New Zealand office workers, a workplace survey study was conducted with the assistance of participating office workers from a large New Zealand business. This study was composed of three parts, a preliminary demographic and hardware survey, periodic observations, and a closing interview. In this section, the design and process for each of these parts is described.

3.1.1 Participants

Participants were sourced from a large[1] multi-discipline regional infrastructure and services company in New Zealand, with 120–150 employees. Employees of this company were informed ahead of time that the researcher would be on site via an email message sent by a senior staff member and were asked to consider participating in the study if they felt comfortable doing so and felt that it would not impact unduly on their work.

The researcher had previously been an employee of this company, working in several different divisions over several years. The researcher had not been employed by the company for 2 years prior to the commencement of this study but was known to many of the staff.

---

[1] Large business as defined by Financial Reporting Act 2013, §45.1(a–b) (The Parliament of New Zealand, 2013), either having total assets valued in excess of \$60 million *or* total revenue exceeding \$30 million at the balance date for each of the 2 preceding accounting periods.

On the first day of the study, staff working from a central office space were approached at their desks and given a brief outline of the experiment. They were then asked if they would participate in the study, and if they expressed interest a more in-depth description of requirements and data confidentiality was given. Of the approached staff, 15 agreed to be participants in the study, with 5 participants identifying as female and 10 as male. Each participant was assigned a randomised ID number, against which observations were recorded.

3.1.2 Preliminary Demographic & Hardware Survey

Following study participant identification, each participant was asked about their job, collecting their job title, and information about their daily work tasks and computer usage. While exact details of each participant's job were collected, participants were also asked to provide a broad field into which their job could be classified to allow for anonymity. Each participant was also asked how long they had worked in their current field independent of their current employer, rounded to the nearest whole year.

Details about each participant's computing environment was collected, recording the number of distinct monitors in use at their computer system, and the diagonal size and configured resolution of each of these. Additional observations were made about equipment used within each working group within the office to help provide context for observation.

3.1.3 Periodic Observations

On a day following the initial demographic and hardware survey, participants were visited at their place of work at hourly intervals, and the current state of their computer screen and workload were be collected. Observations collected the number of open windows on the system, the number of concurrent tasks being performed, how many windows were associated with each task, how windows were arranged on the display, and how much of the time since last observation had been spent on on-screen activities.

The number of open windows was recorded by manually counting entries on the system taskbar and recording the results. Only windows managed by the operating system window manager appearing in the system taskbar were counted to minimise disturbance to participants, and to maintain a focus on the system window management techniques. Applications with sub-windows managed within a MDI or TDI window – such as browser or document editor tabs – providing application-specific management techniques for arranging content were not recorded separately. Two participants using Microsoft Internet Explorer had configured this to display each open tab as a distinct taskbar entry, rather than only displaying a single instance. In both cases, the browser was restored so that the number of distinct windows could be observed, and this value recorded.

Arrangements of windows were recorded as the predominant arrangement of windows at each observation, recording “full screen” if most windows were maximised or otherwise expanded to fill the majority of the display, “tiled” if the majority of windows were laid in a tiled fashion regardless of if assistive tooling was used, and “split” if an equal mix of arrangements were used.

Each observation attempted to minimise the distraction to participants and took 1–2 minutes per participant. If a participant was absent from their workspace for more than 45 minutes of the 60minute evaluation block, they were considered to be out-of-office for that block, and no data was recorded. Results of observations were recorded in a notebook during the study and were later transcribed into a digital form for processing.

3.1.4 Closing Interview

On a day following periodic collection, participants were given a face-to-face interview where they were asked several questions about their computing environment. These interviews were audio recorded where permission was granted to allow for later analysis. Where permission was not given, paper notes were made during the interview. Questions were presented in a manner that allowed for direct factual responses but were followed up with questions that encouraged elaboration on the response.

Questions were grouped into collections that related to different aspects of the participants workflow, experiences, and thoughts regarding their computing environment. These are presented below.

Daily Usage

- Does the participant shutdown, restart or log-off their computer at the end of the day?
- How long does it take after logging into the system at the start of the day before the participant is able to work on a task?
  - If starting from a cold shutdown, this is time required to open daily-use applications
  - If starting from a suspended environment, this is the time required to reacquaint with open applications, windows, and files

Window Identification

- Does the participant have difficulty identifying or locating application windows when working on a task?
  - Is this difficulty related to the number of open windows?
  - Are there tools or system features that assist in identifying windows?

Physical Environment Improvements

- Does the participant believe that they would be more productive if provided with more working screen space?
  - An increase in working screen space would encompass more screens, larger screens, higher density screens
- Does the participant feel that they would enjoy having more screen space to work with?
- Following their last screen upgrade, did the participant notice an increase in productivity, or find it easier to work with the new configuration?
- Does the participant believe that their window managing behaviours would change if extra screen space was available to them?
  - If given more displays, would they use more full screen windows
  - Would windows be grouped differently?

Digital Workflows

- Would the participant prefer if their work was entirely on-screen, with no physical papers or references on their desk?
- Does the participant believe that it would be possible to move to an entirely digital workflow for their working field?
  - Is there any aspect of the participant's job that cannot be performed purely on screen?

3.1.5 Data Treatment

Following completion of the study and processing of data, participant IDs were normalised to sequential IDs for use in data presentation. Interview recording were processed, recording core answers for each asked question, and noting specific examples provided with jobspecific applications and processes simplified for anonymity.

Table 3.1: Participant field of work and computer system summary details. Experience is given in years, rounded up to the nearest whole number. Count specifies the number of distinct displays available to the participant at their desk. Layout describes how available displays were laid out on the desk with respect to each-other. Absent values are given for participants with only a singular display.

ID	Field (Experience)	Computer Type	Count	Layout
1	Administration (20)	Workstation	2	Horizontal
2	Information Management (8)	Workstation	2	Horizontal
3	Regulatory Compliance (30)	Workstation	2	Horizontal
4	Accounting (35)	Workstation	1	
5	Regulatory Compliance (6)	Workstation	2	Horizontal
6	Project Management (15)	Docked Laptop	2	Vertical
7	Administration (2)	Workstation	3	Horizontal
8	Information Technology (4)	Workstation	3	Horizontal
9	Systems Monitoring (23)	Workstation	5	Horizontal†
10	Information Management (1)	Workstation	2	Horizontal
11	Project Management (1)	Workstation	2	Horizontal
12	Project Management (7)	Workstation	2	Horizontal
13	Engineering (3)	Workstation	2	Horizontal
14	Accounting (4)	Workstation	1	
15	Information Technology (10)	Workstation	2	Horizontal

3.2 Preliminary Survey Results

In this section, results collected from the preliminary survey of participants is presented.

3.2.1 Demographic

Participants who chose to take part in this study came from a variety of fields and departments within the surveyed company and represented a range of experience levels within their identified fields.

3.2.2 Display Number

From examining display data presented in Tables 3.1 and 3.2, the majority of participants had more than one display on their desks. Only 13% of participants, all of whom identified their fields as “Accounting”, had access to only a single display computing system. 66% of participants had access to two displays, while 13% had three. A single participant had access to 5 displays, which was an outlier for the entire company as observed.

Table 3.2: Screen sizes in diagonal inches for each participant. Displays are presented in the order in which they were arranged on the participant’s workspace in left-to-right order, excepting participant 6 for whom the order represents bottom-to-top. Bold-faced values indicate the primary display as configured within the operating system.

ID	Display Size & Order				
1	**24”**	17”			
2	**24”**†	24”†			
3	**24”**	24”			
4	**24”**				
5	**24”**	17”			
6	**15”**	24”			
7	**24”**	24”	17”		
8	24”†	**24”**†	17”		
9	22”	22”	**22”**	22”	22”
10	**24”**†	24”†			
11	**24”**	24”			
12	**24”**	17”			
13	**19”**	19”			
14	**24”**				
15	**22”**	22”			

Table 3.3: Number of observed displays, grouped by diagonal display size. All displays for a given size were of the same manufacturer model, with the resolution set to the display native resolution through group policy system management. Aspect ratio and native resolution of each display size is provided.

Size	ℕ	Aspect Ratio	Resolution
15”	1	16:9	1366 × 768
17”	5	4:3	1024 × 768
19”	2	5:4	1280 × 1024
22”	7	16:9	1920 × 1080
24”	12	16:9	1920 × 1080
24”†	6	16:10	1920 × 1200

Arrangements of these displays was noted as an additional point of data during the survey data collection. All workstation users with multiple displays chose to arrange them horizontally across their desk spaces, with the edges of the displays touching. All of these horizontally arranged displays, with the exception of participant 9, had positioned their displays such that each display faced the position of the chair at the desk so that when facing each display, their gaze would fall approximately perpendicular to the display surface. Participant 9’s displays were arranged in a straight line, with each display face parallel to the front edge of the desk. It was noted by this participant that the reason for this arrangement was due to the common usage of the workstation, where multiple people would regularly view different display contents, rather than the participant consuming content from these exclusively. Participant 6 arranged their displays vertically, with the docked laptop centred below the secondary display.

3.2.3 Display Sizes

When inspecting Table 3.3, the most common size of display seen used by participants was 24”, representing 54% of observed displays. 21% of displays were 22” models, though these were only seen with two participants, with the 5display configuration of participant 9 accounting for most of these. 17” displays accounted for 15% of observed displays, and 19” displays accounted for 6%, all of which were utilised by a single participant. The smallest display observed was 15” in size with only a single recorded observation and was the laptop display used by participant 6.

3.2.4 Display Resolutions

Display resolutions varied within observed displays, and within the 24” displays. When considering the displays 17” or larger — representative of the stand-alone displays — for each increasing step in display size, a corresponding increase is seen in resolution for this display. 24” displays were noted in two different aspect ratio models, with the 24”[†] 16:10 displays being primarily used by participants who made regular use of computer-aided design (CAD), geographic information system (GIS), and graphics editing software. This less common 16:10 aspect ratio display presents an increase in vertical resolution when compared to the more common 16:9 24” displays observed.

The smallest display, the 15” display that made up a part of the laptop computer used by participant 6, bucked the trend of larger displays having a greater resolution. This smaller 15” display exhibited a resolution greater than that seen with the larger 17” stand-alone display models. As this was the only laptop display amongst the surveyed participants, no trend can be established between laptop display resolutions and sizes for the surveyed cohort.

3.2.5 General Observations

Within the company all computer systems were supplied by Hewlett-Packard Enterprise, and as a result a high level of consistency and homogeneity was observed within the surveyed computer systems and peripherals. Workstation computer systems within company departments were all the same model and specification, though variation in system models were seen between departments. Laptop systems were less common than their workstation counterparts, with the majority of these seen used by managerial staff. Variation was seen in laptop system models, with clear demarcation between office and field usage models apparent, but with models varying within these two usages.

Peripherals including displays also varied between and within departments, but clear generations of equipment could be observed distributed amongst staff. Displays of equal size were all noted to be of the same model, with distinctive features allowing for simple identification. All computer systems were equipped with wired optical mice and keyboards, with no alternative pointing or input devices observed.

Each office worker was provided with a spacious individual workspace, approximately 8–9m^2 in area surrounded on 2–3 sides with walls or workspace dividers. The workspace of participant 9 was an exception, with their desk and associated equipment positioned in the centre of a large room with no dividing walls. Office workers were provided with large straight or corner desks approximately 0.75m deep, with other furniture varying between departments and requirements.

3.3 Periodic Observations Results

The periodic observations component of this study aimed to record aspects of each participant's daily working process, and to record their usage of different tools available to them in managing digital aspects of tasks. In this section, each recorded aspect is presented, highlighting the range and extremes of observed behaviours.

3.3.1 Concurrent Tasks

The number of concurrent tasks being performed by each study participant was of significant interest for this study, and evaluation of the results has shown that it is common for users to be performing multiple tasks concurrently on screen. Observing Figure 3.1, a collection of results is observed in the 1-2 tasks band, with a smaller number observed in the 2-3 tasks band. A pattern of alternation can be observed from this chart, in that participants switch the number of concurrent tasks quite regularly. This alternation pattern indicates that users do start and stop tasks throughout the day, though the exact reason for this pattern is not clear. Possible explanations for this pattern are that new smaller tasks are brought to the attention of users, and they manage these as they come up; or alternatively these smaller tasks could be completed as a break from working on larger tasks.

Two major outliers are observed, peaking on the extreme end of the scale, with the job categories of Information Technology and Information Management respectively. This could be construed that this indicates a higher loading of tasks for these job categories, but due to the limited number of participants from each category, it could easily be attributed to these participants merely multi-tasking more heavily over the course of the day. Should this study be conducted again, inclusion of higher numbers of participants from each job category would allow more useful conclusions to be drawn.

Table 3.4: Number of concurrent tasks were being worked during each observation time block for each participant. Blank values indicate that the observed participant was absent for more than $\frac{3}{4}$ of the time block.

ID	08:00	09:00	10:00	11:00	12:00	13:00	14:00	15:00	16:00
1	1	2	2		1	1		1	1
2	1	1	3	3		3		6	2
3	1	1	1	1		1		1	1
4	2	1	1	2		2	1	2	2
5	3	3	3	3		3	2	1	2
6	1	1	1	1		1			1
7		2	1	2		1	2		
8	2	5		3		2	2		3
9	3	3	3	3	3	3	3	3	3
10	1	1	1	1		1	1	1	
11	1	1	1	2		1	1	2	2
12	0	1	1		1	1		1	1
13								1	1
14	1	2	2	2	1	1	1	1	1
15	2	2	2				1		

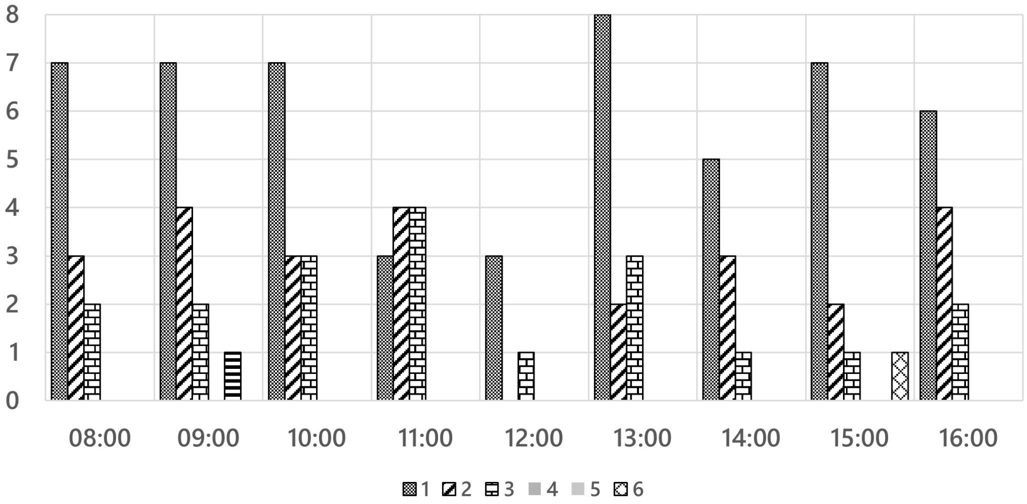

Figure 3.1: Number of concurrent tasks reported by all reporting participants at each observation interval.

Table 3.5: Numbers indicate how many program windows were open during each observation time block for each participant. Sub-windows such as browser tabs, or paged documents in an application are not counted as individual windows. Blank values indicate that the observed participant was absent for more than $\frac{3}{4}$ of the time block, and as such was unavailable for data collection.

ID	08:00	09:00	10:00	11:00	12:00	13:00	14:00	15:00	16:00
1	1	5	3		5	6		5	5
2	3	4	11	11		11		13	9
3	7	5	5	7		12		13	10
4	2	2	2	2		2	1	3	3
5	7	7	7	7		7	6	11	15
6	4	4	4	6		2			1
7		4	1	5		4	4		
8	5	10		15		12	10		12
9	7	7	7	7	7	7	7	7	7
10	5	5	5	7		8	8	8	
11	7	8	8	9		8	8	13	13
12	5	5	5		7	5		5	5
13								1	3
14	2	3	3	3	4	4	4	4	4
15	4	9	3				5		

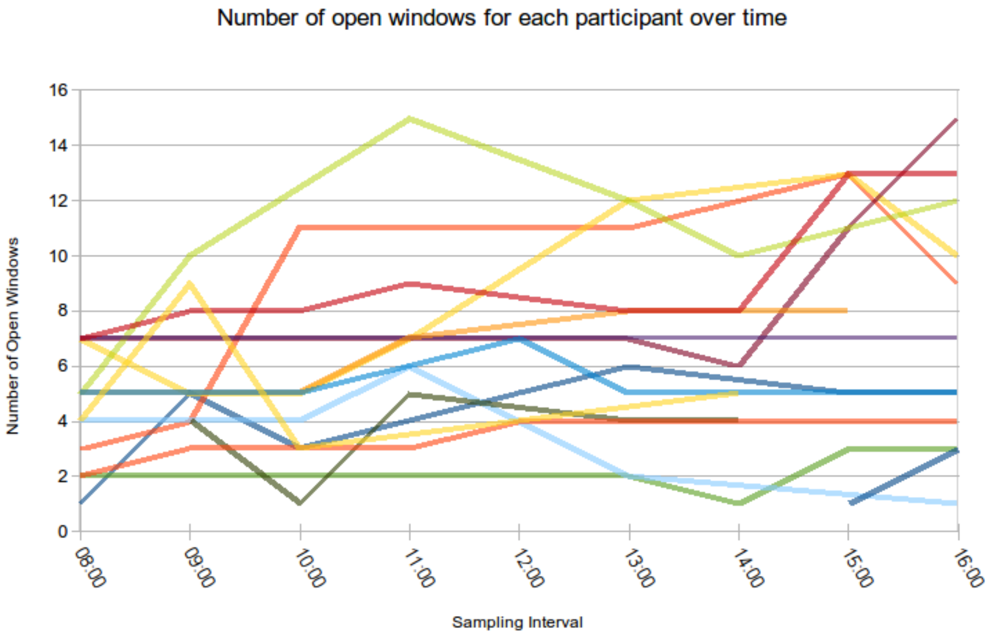

Figure 3.2: Line plot showing the number of open windows for each participant at each sampling time. Missing values are omitted, with the line being plotted between known points.

3.3.2 Windows Count

From observing the average numbers of open windows presented in Figure 3.3, a steady upward trend in the number of open windows can be seen as the day goes by. This suggests an accumulation of windows over the course of the working day, where windows that were not considered important enough to be opened at the start of the day are kept around once they have been opened. From observing Figure 3.3 it could be construed that a part of this apparent upward trend can be attributed to a subset of users with larger numbers of open windows, which results in larger swings of the mean and median values. Though this could be an explanation, there is an apparent upward trend visible in Figure 3.2 which indicates that these users with a profusion of open windows are not solely to blame for this observed pattern.

From observing the mode statistic in Figure 3.3, a peak and trough cycle can be clearly seen — specifically troughing around 10:00 and 13:00; Morning tea and lunch break times respectively — indicating that a proportion of users shed windows before leaving their desks. This could be attributed to a desire to complete tasks before leaving the office for breaks, and prioritising smaller tasks for these morning blocks because of this. The opposite pattern can be seen around 15:00 — traditionally afternoon tea — in which the mode peaks; indicating that the participants were working on more complex tasks at that point, that extended over their break.

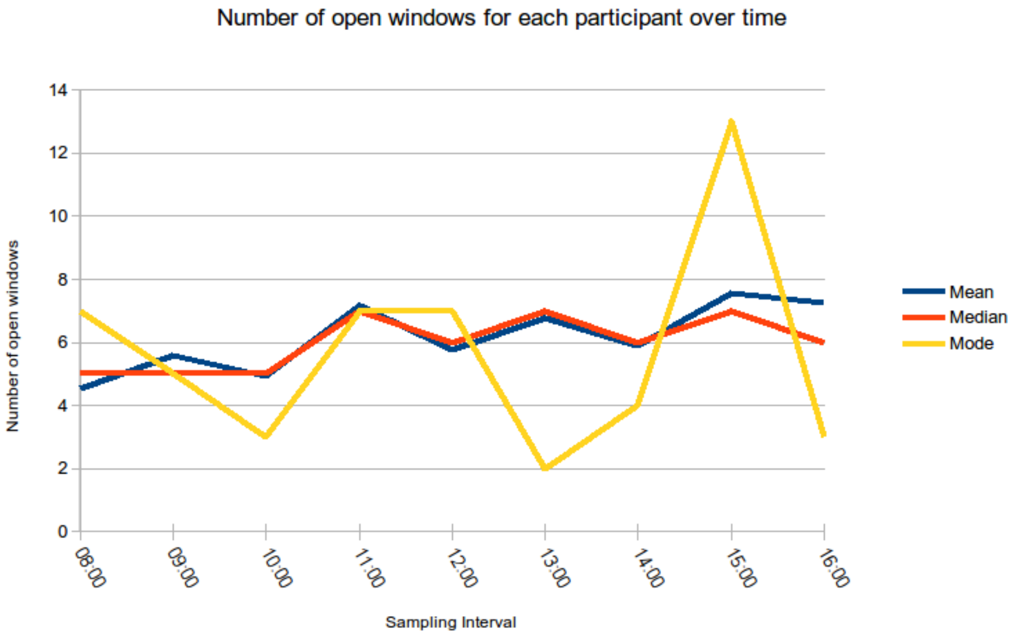

Figure 3.3: Line plot showing the number of open windows for each participant at each sampling time, represented as statistical measures.

3.3.3 Window Arrangements

From examining Figure 3.4, it can be immediately seen that the most common window arrangement method for the participants was making windows full screen. This method ranked the highest in all but one observation period, with this period being one of low reporting for participants. Both tiled (in which windows are freely sized and allowed to overlap) and split (in which windows were both tiled and full screen) ranked relatively equally throughout the study.

From the raw data shown in Table 3.6, it can be observed that participants using larger numbers of windows tend towards using tiled or split arrangements as their arrangement methods, whilst participants with fewer windows to manage tended towards full screen arrangements. This relationship can be visualised by observing Figure 3.5, where while the full screen arrangement is by far the most dominant arrangement method with low numbers of open windows, it falls entirely out of favour once 8 windows are open; with the preference shifting towards the tiled and split arrangements.

Table 3.6: Number of window arrangement method incidences against the number of windows open at the time of observation.

ℕ	Tiled	Full screen	Split
1		5	
2	1	5	
3	5	1	2
4	8	5	
5	2	12	4
6	1	2	
7	1	15	2
8		1	6
9	1	1	1
10	1		2
11	4	1	
12	2		1
13	1		3
14			
15	1		1

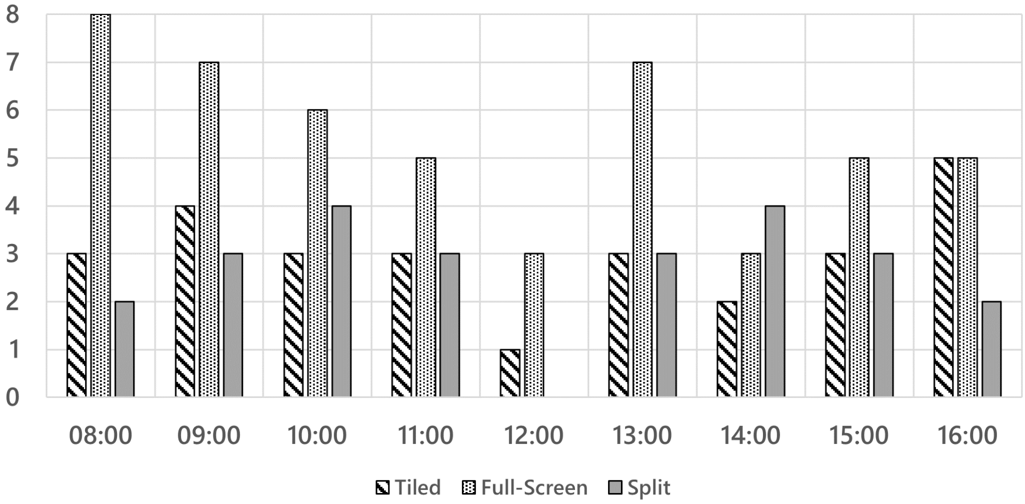

Figure 3.4: Bar graph showing the frequencies of window organisation methods during each observation period, for participants that were in office during the specified time period.

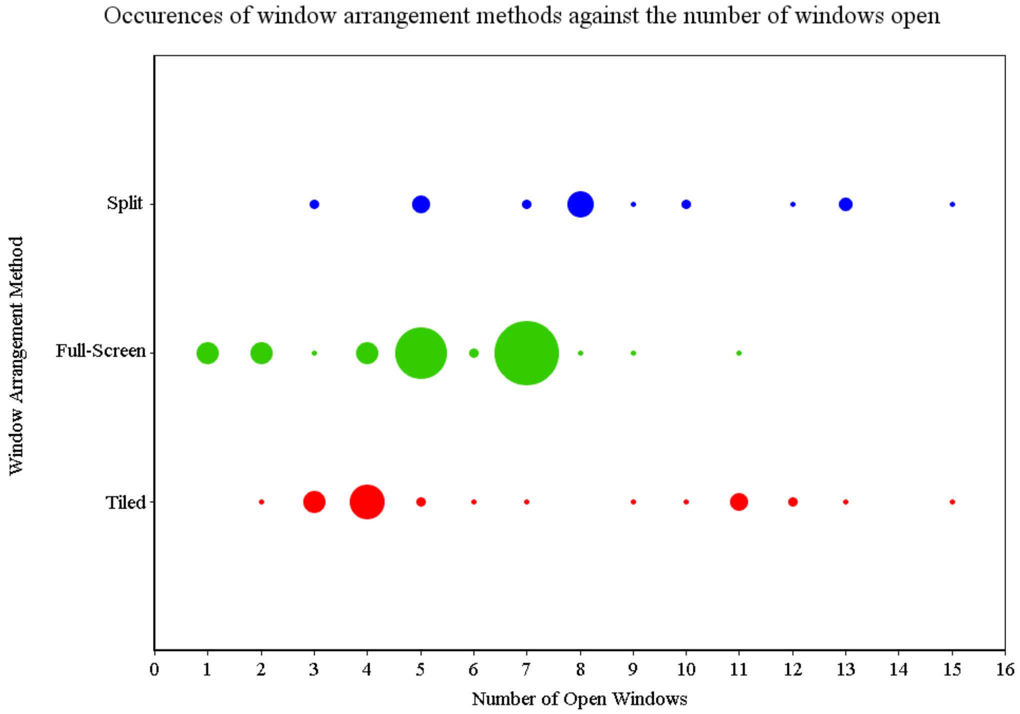

Figure 3.5: Bubble chart showing the number of instances an arrangement method was observed, against the number of open windows at the time of observation.

3.3.4 Screen Time

Observing Figure 3.6, for over 50% of the observed participants, the entirety of the time they spend at their desk consists of time spent working on their computer screens. A large trough in which the median drops to 0% can be observed occurring at 12:00, which can be attributed to the lunch break of users - as reflected in the absent values in Table 3.7, as well as in the window drop discussed in Section 3.3.2. The mean percentage fluctuates between 50% and 80% screen time - with the obvious outlier at 12:00 being excluded - indicating that while some participants were required to work using off-screen, a proportion of work was performed on-screen. Off-screen resources utilised by participants varied from hand-drawn diagrams, documents and notes; through to reference manuals, books and regulated hard-copies depending on the participants job - with some job categories requiring significantly more of these than others.

A notable outlier observed in Table 3.7 is participant 13 - an engineer - who for most of the day reported a 0% screen time. This result is due to the participant being outside of the office performing fieldwork over the course of the day. While the results from this participant does affect the mean result graphed in Figure 3.6, it is a valuable result, demonstrating that some users

Table 3.7: Percentages indicate the division of work time devoted to activities on-screen during each observation time block for each participant. Participants absent for more than $\frac{3}{4}$ of the time block are recorded as having a 0% time-on-screen value due to their unavailability for data collection.

ID	08:00	09:00	10:00	11:00	12:00	13:00	14:00	15:00	16:00
1	100%	100%	33%	0%	50%	20%	0%	100%	100%
2	100%	100%	100%	100%	0%	100%	0%	50%	100%
3	100%	100%	100%	100%	0%	100%	0%	100%	100%
4	100%	100%	33%	100%	0%	100%	100%	100%	100%
5	50%	50%	50%	50%	0%	50%	25%	100%	100%
6	100%	100%	100%	100%	0%	100%	0%	0%	100%
7	0%	33%	33%	66%	0%	66%	100%	0%	0%
8	100%	100%	0%	100%	0%	100%	100%	0%	100%
9	100%	100%	100%	100%	100%	100%	100%	100%	100%
10	100%	100%	100%	50%	0%	100%	100%	100%	0%
11	100%	100%	100%	100%	0%	100%	100%	100%	100%
12	25%	50%	50%	0%	100%	50%	0%	50%	50%
13	0%	0%	0%	0%	0%	0%	0%	100%	100%
14	100%	100%	100%	100%	100%	100%	100%	100%	100%
15	100%	100%	100%	0%	0%	0%	100%	0%	0%

- while classified as primarily office workers - can and do spend time away from their office working. This observation can be used to provide support for a recommendation for providing or supporting portable computing solutions that provide facilities of the office environment, but the exact methodology of this is outside of scope for this study and discussion.

3.4 Closing Interview Results

On a day following the conclusion of periodic observations, each participant participants in an interview with the researcher. These interviews covered the topics outlined in Section 3.1.4 and took no longer than 15 minutes. 14 participants consented to having their interviews recorded, with 1 declining. Written notes were taken for the non-recorded interview. Recorded interviews and written notes were processed collectively following the study, with the results presented here.

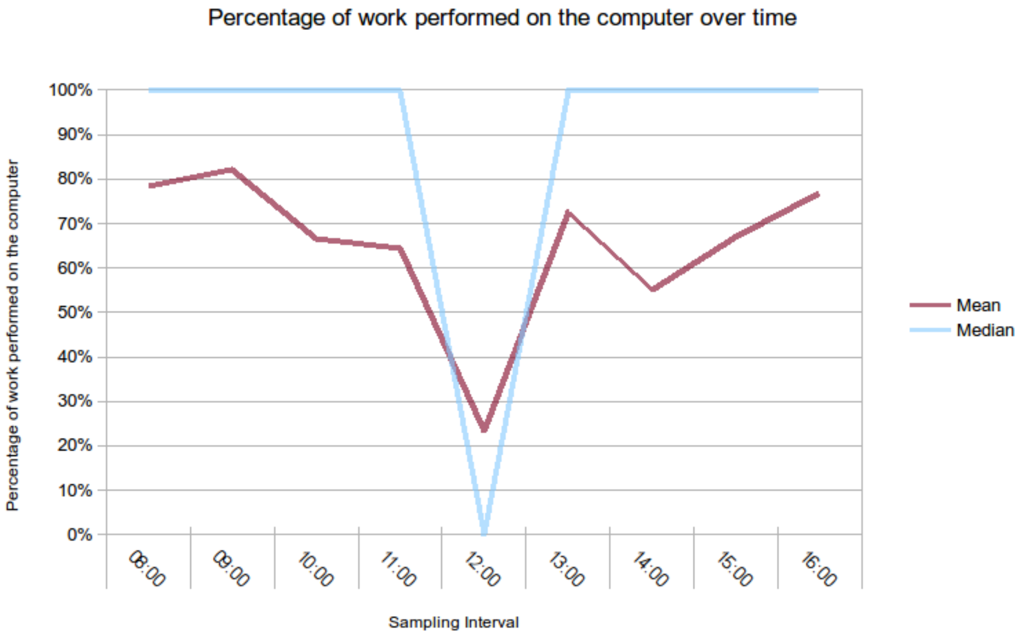

Figure 3.6: Line plot showing the percentage of on-screen activity at each sampling period. Periods with no data are assumed to have a 0% value.

Table 3.8: Participant responses for whether they shutdown or restart their computer at the end of each day.

Yes	No
9	6

3.4.1 Daily Usage

Participants were asked if they shut down or restarted their computer at the end of a working day, leaving them with a blank environment at the start of the next working day. 9 participants indicated that they did this at the end of every day without fail. One reason given for this was IT policy, requiring that computers be restarted regularly to receive updates. It was noted by several participants that on occasions where they had left applications open at the end of day, their computers had automatically restarted to apply updates which resulted in lost work. Following these occasions, they had begun shutting down daily to prevent this possibly occurring again. Three participants indicated that shutting down their computers gave them a "sense of closure" for the day, allowing them to mentally "down tools" for the evening.

Table 3.9: Estimated time required to reacquaint with or open windows for daily working environments.

Required (minutes)	ℕ
0	3
1	1
2	3
3–4	2
3–5	2
4	1
5	2
5–10	1

Participants who indicated that they left their computers running had several reasons for doing so. Participant 9 indicated that the monitoring systems were monitored “24/7/365” and restarting these was an event managed by IT only after a standby system was activated. 2 participants indicated that their work often involved tasks that would take longer than a day to complete, and that shutting down applications mid-task would cause significant delays the following day, so tasks were often left running overnight so they would be available the following day. The remaining participants indicated that they required remote access to their computers during the evenings and weekends, and the computers were left running to ensure they would be able to access them when required.

When queried on how long it took to open required applications or reacquaint themselves with open applications, 3 participants indicated zerotime requirements. Participant 9 was among these, which is unsurprising given the constantly running nature of the computer system. The other two indicated that this lack of time requirement was due to the lack of consistency of tasks performed, with each task and day requiring different applications and documents to be opened. This meant that these participants had no “standard set” of windows to be opened each day. 1 participant indicated that the required time to open their standard set of applications at the start of each day ranged from 5–10 minutes, which was noted as being due to a requirement to log into several distinct remote systems. This process was indicated to be mostly automated, but minimal work could be done during this process.

The remaining participants indicated a requirement of $1 \leq \mathcal{T} \leq 5$ minutes to get all required applications open at the start of the day. This varied based on the number of applications required to be opened, but also on the speed of the computers that were used. Several participants indicated that their computer "sits there, chattering away to itself" for extended periods when opening applications, suggesting that a mechanical hard drive is at least partly to blame for launching delays.

3.4.2 Window Identification

During the individual interviews, study participants were asked if they had difficulties locating open windows running on their systems during their daily work. In response to this question, $53.\dot{3}\%$ of the responses indicated that they did have issues finding windows of interest. Several interviewees offered indications as to what provided issues for them, and in some cases how they worked around the issue.

It was indicated by several participants, that due to the nature of their jobs, they would often have several windows of the same application open — i.e., multiple Excel, Word, CAD or Accounting software windows open — but displaying different files or information in each window. This leads to confusion when switching between windows, as each window instance of the target application would need to be considered prior to activating a window.

The hover-view feature of the Windows 7 taskbar — where hovering the mouse over a taskbar item will bring the associated window to the front of the screen — was praised by a number of participants when dealing with a few similar windows but was indicated to not be particularly effective when dealing with larger numbers. This indication was attributed in part, to the time needed to activate the hover-view feature for each window; where the mouse is required to hover over the taskbar item for a brief period before the functionality is triggered. When dealing with two windows, the user only needs to check one of the windows before deciding which is the correct one — requiring just one hover action — whereas with larger numbers of windows, they need to check several to identify the correct one, requiring a noticeable amount of time to view each window.

Table 3.10: Participant responses to physical environment improvement questions.

Aspect	Yes	No
More productive with more space	10	5
Desire more space	11	4
Recent upgrade improved productivity	10	5
Window management change with more space	11	4

A number of participants indicated that they made use of their multiple displays to assist in window identification. Several participants explained that when working from similar windows — such as referencing information from one spreadsheet while working in another — they would position each of the documents on separate screens so that window switching would not be required. Other participants indicated that specific windows would always be placed on a specific screen, such that they knew approximately where it would be located.

3.4.3 Physical Environment Improvements

When interviewed, most participants indicated that more space was both desirable and linked with increases in perceived productivity. When queried about their beliefs on the impact on productivity providing larger or more numerous displays would have on them, 10 participants indicating they felt they would be more productive. Of the 5 participants that indicated they did not believe this would improve their productivity, 4 indicated that they believed that their current number of displays was sufficient for the number of windows they managed, and that additional displays would be unused while larger displays would just display their current windows larger which was unnecessary. The remaining participant indicated they had previously had a larger number of displays, and that the overhead of managing windows over this space began to overwhelm any perceived benefits of the increase. They did indicate that this could perhaps be overcome with better tooling to group and identify windows.

When asked if they desired more space independent of possible productivity benefits, 11 participants responded in the affirmative. A number of these participants indicated this desire related to an element of prestige and value, where larger “fancier” displays in larger numbers would make them feel important and would further make them feel that the business valued them by providing high quality equipment. Others indicated that while they did not believe this would significantly improve their productivity, it would improve comfort. Those that responded negatively suggested that they believed their desk space was already being impinged by displays, and that increasing space requirements would make their desk environment feel more cramped. They indicated that any possible benefits gained by increasing space would negatively impact their other activities and would thus be undesirable.

All participants indicated that they had experienced display improvements since working with the business, receiving additional displays or larger displays during the upgrade. 10 participants indicated that they believed that their productivity was improved by the upgrade. 8 of these participants indicated the upgrade involved increasing their number of displays, with the remaining 2 receiving larger displays to replace existing display configurations. Those who received an additional display for the first time indicated that their productivity was significantly increased, while those who received displays additional to their existing multi-display configuration suggested that the increase in productivity was only felt with certain workflows. The 2 participants receiving larger displays both moved from “square screens to the nice wide ones” during their upgrade, and indicated they felt more productive as two windows could be comfortable arranged side-by-side when required, which was “not a nice experience” on previous displays.

Of the 5 responses indicating that their recent display improvement did not improve their productivity, all noted that their upgrade involved replacing their existing configurations with larger displays. All indicated that their window management strategies did not change as a result of the upgrade, so no productivity benefits were experienced. 4 of these participants did indicate that they still appreciated the upgrade, suggesting that the displays felt nicer to use than before. The remaining participant indicated the upgrade was given due to the previous displays failing, and that the displays were “not a significant upgrade” in any regard on what was used prior.

11 participants indicated that if they were given additional or larger displays that their window management strategies would change. Most of these responses focused on additional display behaviours, with most indicating that larger displays would have a lesser impact on their window management behaviours. With additional displays, participants who made use of manually tiled layouts suggested they would worry less about detailed arrangement and would distribute windows to different displays as full screen or split-view arrangements instead. Those who made use of primarily full screen layouts indicated more displays would allow for reduced switching and hunting of windows by allowing more full screen windows on different displays. Those who

Table 3.11: Participant responses for digital workflow questions.

Aspect	Yes	No
Desire digitisation of work	11	4
Possible to work digitally	12	3

responded in the negative all indicated that they felt their current window management strategy suited their workflows, and that an increase in space would at best allow them to apply the same strategy to more displays, or at worst remain unused for primary layouts.

3.4.4 Digital Workflow

When asked about digitisation of workflows, participants provided responses around their preference for digitisation, and for whether their job could currently be performed digitally if desired. Of the 15 participants, 11 interviewees indicated that they would prefer a working environment in which all their information, work and reference material was digital. A primary reason for this preference related to ease of access to information, with in-document search and document library search identified as a major advantage of digital representations over paper. Digital submission of documents to avoid illegible writing or soiled paper documents was another strong reason for digitisation given by those who regularly engaged with external contractors.

Those that did not indicate a preference for digitisation gave two reasons — a preference for physically writing or sketching notes, and a belief that digital processes replacing paper models were slower and less efficient than staying with the paper model. Two of these participants indicated that if a digital solution could be found that allowed for rapid free-form note taking and sketching — such as a drawing tablet and stylus dedicated to the task — was provided, then this could convince them to switch to digitising this portion of their workflow, however the other two participants indicated that they were "too stuck in [their] ways" to change. All of these participants indicated that if improvements were made to the software and processes replacing the paper models, that this would encourage them to use these. All noted that they had nothing against the digital approach inherently and acknowledged that the gradual digitisation of processes had mostly been positive, making their jobs easier.

12 participants indicated that it is entirely possible for their job to be performed in an entirely digital manner, meshing well with their preferences. Several participants indicated that they already attempted to keep their workflow as digital as possible, digitising physical documents on receipt before beginning processing the digital versions. These participants noted that this was reflected in their desk arrangements, with paper contained to in-out trays with minimal backlog. One participant who indicated that their job could be done digitally but expressed a preference to sticking with paper felt that due to this possibility, the option for sticking with paper may soon be taken away from them, and indicated they were not happy with this idea but were resigned to it as a reality.

The three participants who indicated that their work could not be entirely digitised gave different reasons for this. One participant indicated that due to regulatory requirements, physical papers were required to be signed and kept on hand during moving time windows for audit, and that digitisation of this task would leave them in breach of regulations. A second participant indicated that for health and safety reasons, a “lock out, tag out” system was maintained in-office to represents services being maintained in the field, and physical interaction with this system was required before digital actions could take place. The third participant explained that due to the nature of field work undertaken, no computing equipment could be taken on site, requiring paper be used for observations and notes. It was not indicated in this case if restrictions on equipment were due to environmental, security, or regulatory reasons.

3.5 General Observations

During periodic collections of data, several observations were made regarding actions performed by all participants that were of note. While participant backgrounds were varied, all participants demonstrated behaviours that could be collectively categorised.

3.5.1 Lifecycle Management

Study participants all performed application life-cycle management tasks, creating new window instances, and destroying others. There were several different approaches to each of these demonstrated by participants, which did appear to vary based on applications in use, and familiarity with computing systems.

The most common method observed for creating a new window instance was selecting and opening a file or shortcut from the system file explorer or desktop. This action opened an application instance directly, most often with content from the selected file already displayed. This action took advantage of the operating system's file associations, reducing a twostep task of opening an application then selecting a file to open, into a single step of simply locating the file of interest. A less frequently used method observed was using the system application menu to view and select from a list of applications, creating a new instance of that application when selected. A final method noted was creation of new windows from already open MDI application windows. Specific observations were dragging browser tabs free of the main container to form a new container and detaching tool-pane components from a CAD application.

Closing applications was accomplished similarly in different ways. The most common method observed was to select the "close" button in the top-right of application windows and acknowledging any resulting prompts. A subset of users was observed closing application by right clicking the taskbar entry of a window, and selecting the close option from the resulting context menu. A single participant was observed closing windows using the Alt-F4 hotkey.

3.5.2 Window Layout

While window layouts were of particular interest for the observations, observing participants performing these layout actions was rare. This suggests that while organisation of windows is of interest to many participants, that it is not a frequently performed task. The only occasions where this was observed was when new application window instances were created, where the observed participant would perform a layout action before working with the window.

Methods of layout varied based on the layout strategy preferred by the participant. Those observed that preferred full screen applications simply dragged newly created windows to their display of choice, then maximised them using the button in the top-right of the application window. Those who utilised split-view functionality used either the keyboard shortcuts Super-Left or Super-Right to position the window to the appropriate split area or dragged the window to the edge of a display to toggle the splitting behaviour.

Only a single layout observation was made for a participant preferring a tiled style of layout, which was a more involved process. Several existing windows were re-positioned and resized to create a space for the newly created window, which was then positioned and sized to fill the gap. Despite the extra required steps involving other application windows, these actions were performed quickly and confidently, suggesting the observed participant had a layout planned prior to beginning the organisation process. This participant was not queried on this behaviour at the time of observation to avoid disruption to work, and when questioned during the closing interview was unable to recall the circumstances described, instead suggesting experience with a repeated workflow as the likely reason behind this.

3.5.3 Window Identification

Participants with many open windows were often seen hunting for specific windows, utilising three primary strategies to do so.

The most frequently used approach was the use of the taskbar and taskbar peeking, looking through the taskbar list of windows to identify the window of interest, or hovering over taskbar entries to display thumbnail previews of windows for visual identification. Some participants were observed positioning their cursor at one end of the taskbar and moving it along each entry in the window list, inspecting each thumbnail as it was displayed until the desired window was identified, then activating it from the same taskbar entry. This approach seems to suggest that the application name and icon representation in the taskbar held little meaning to these participants, requiring a visual identification of the desired document to associate with it. Other participants used a different strategy, hovering only over taskbar entries for applications that could possibly contain the document of interest. These participants possibly had a stronger association of general document types to applications (i.e., spreadsheets will be opened in an Excel window), or specific documents to specific applications (i.e., this process document is always shown in a help viewer window).

A popular alternative to taskbar search was the use of window cycling with the Alt-Tab hotkey. Some participants were observed rapidly swapping between two windows using this key combination when the windows overlapped. Others were seen cycling through the window list to identify the window of interest. While some participants noted that they were aware of the Vista style 3D flip approach to window cycling, no participants were observed making use of this. No reason for this was forthcoming, suggesting the preference of the older Alt-Tab cycling over the newer Super-Tab cycling is due to familiarity with the keybinding rather than a direct preference.

A final method observed with a minority of participants was to show and hide windows. Participants utilising this approach would repeatedly minimise the current topmost window until the desired window was found, or no further windows were displayed. In the latter case, these participants would then move to the taskbar, restoring each window in sequence until they found the target window. This was a noticeably slower method for window identification than with the other approaches and was mostly observed in older participants.

A notable example who experienced difficulties with window identification was participant 3. For their workflow they often had many visually similar spreadsheets open simultaneously and required frequent reference between different documents. This participant made use of both taskbar peeking and Alt-Tab cycling though windows depending on the number of open windows. When many similar windows were open, they noted that the taskbar peeking method was less efficient than with Alt-Tab, as the number of thumbnails presented would often exceed the width of the display, preventing identification of some windows. Alt-Tab cycling was noted as generally less efficient for their purposes as every window need be displayed sequentially, but was useful in these larger window number scenarios. To aid with window identification, this participant introduced temporary visual elements to each document to help with at-a-glance identification. The primary method observed with documents and spreadsheets was to set the background colour of a cell range or paragraph to a distinct colour, to allow for document association in this way. They noted that they often applied ordering strategies with these additions, giving an example of using the natural order of colours to visually identify where a in a time sequence each document lay. While an exceptional example, this participant did highlight a shortcoming with window identification tooling, in that for large numbers of windows having similar appearances, that consideration of each document is required to identify the target in the absence of visually distinguishing features.

3.5.4 Content Consumption

By a large margin, the most common content consumption task observed with participants was that of reading text from the screen. While this took several different forms — including reading values from specialised systems, tabular data from spreadsheets, and bulk text from documents — all participants engaged in reading tasks for most of the day. When initially surveyed about their fields and what was involved with their jobs, several participants joked that most of their job was “reading things that other people have put in all the effort of writing”.

Camera and map images were commonly seen as reference content for other tasks amongst many participants. These were rarely present on observed participant's screens for extended periods, seemingly opened to provide context to written notes and observations. The information management group presented an exception to this, with these items present in larger numbers for extended periods while CAD and GIS systems were updated to reflect content presented on the items.

Consumption of video media was minimal during observations, with a single participant observed making notes while watching training videos. Other participants were noted as consuming video media indirectly through background videos unrelated to their work, however this content was rarely in the foreground.

3.6 Summary

In this chapter, the results of an in-office exploratory user study were presented. From the results collected it was identified that most study participants make use of more than one display, and that displays used have been trending towards larger 16:9 aspect ratio displays during upgrades. Participants indicated that upgrades to their display sizes and number largely improved perceived productivity, with those moving from a single display to multi-display configurations indicating the strongest sense of performance and productivity benefits. This provides qualitative evidence to support results reporting on productivity benefits of multiple displays presented in Chapter 2.

Desire for increase in display space was clear amongst participants, with belief that this would allow for further increases in productivity expressed by many. Reasons given for this belief were primarily connected to the number of windows managed, which were observed to be numerous and tied to multiple concurrent tasks during observation. Window layouts observed show a strong preference for use of full screen windows for arrangement, suggesting that positioning and sizing of windows is not at the forefront of most user's thoughts when performing tasks, at that they prefer expediency of layout over efficiency of space usage. This further suggests that assistive window management techniques are difficult to learn or use for users, and that further pursuit of these to improve utilisation of space will be less impactful than simply providing more space in which to place windows.

Chapter 4

Virtual Reality Display Technologies

In the previous chapter, it was identified that users preferred increasing physically available display space for increasing productivity, with assistive window management techniques not being sufficient to support increasing numbers of tasks and windows displayed on limited space. While simply adding larger or more numerous displays to computing systems may address this desire for many, there are constraints on the number of displays that can be effectively used with conventional computers. An alternative approach would be to consider use of virtual reality (VR) hardware and environments to replace the current conventional environments, making use of a potentially infinite amount of virtual space without physical restrictions imposed when using conventional displays.

In this chapter, an overview of the term virtual reality in terms of environments and hardware is given, and the possible applicability to office worker computing use-cases for each is discussed. Immersive virtual reality (IVR) head-mounted displays (HMDs) are identified as device candidates for investigation, and virtual environment design considerations are presented. Methods of evaluating virtual environment performance are outlined, with reading performance selected as the metric for evaluation, addressing Research Question 2, which stated:

> *What performance metric should be used when comparing applications within differing environments?*

Finally, a brief overview of the physiological mechanisms associated with reading is presented, and possible impacts on reading within IVR HMDs devices is discussed.

4.1 Virtual Reality Environments & Systems

After evaluation of the state of current conventional computing systems in Chapter 2 and the results of the user survey and observations in Chapter 3, it is apparent that there is a desire amongst users for more space in which to organise an increasing number of application windows. While significant effort has been put into creation of software space management tooling, the observed participants of the user study did not make use of many of these, for reasons that included unfamiliarity with the techniques, and lack of desire to change their existing workflow.

To provide more working space, these same participants made use of larger and more numerous conventional displays, which was noted as providing self-reported benefits to their productivity, verifying observations made in synthetic workflows (Czerwinski et al., 2003; Anderson et al., 2007). Most participants of the study also indicated that they felt that more displays would further increase their productivity, with a minority indicating that they felt there was an upper limit to what could be achieved through addition of more displays. These observations are reinforced by the literature, which has identified limitations on the practicality of increasing the number of displays, with costs, space requirements (McGill et al., 2020), and physiological constraints (Nimbarte et al., 2013) presenting practical limitations to this approach of increasing space.

As it appears that both software and current hardware approaches to increasing and managing working space are not sufficient for managing the number of application windows, alternative options should be considered. While development of further software space management tooling could be pursued, the apparent resistance to adoption of existing tooling suggests that these would likely go unused. This leaves alternative hardware as a potential avenue for investigation, moving from the long-established conventional displays to alternative or supplementary devices to address the problem.

While Section 2.1.4 discussed the use of conventional displays as the usual method of visualising outputs from a computer system, several alternative display technologies have been seen over the evolution of home and office computing, though few have seen widespread adoption. One of these technologies of note is that of virtual reality devices, which at present have seen adoption for speciality applications such as CAD (Autodesk Inc, 2021), building information modelling (BIM) (BIM Holoview Ltd, 2021), training (Larsen et al., 2009), conferencing (Gunkel et al., 2018), manufacturing and assembly (Evans et al., 2017), and telemedicine (Riva & Gamberini, 2000). While general-purpose computing environments have been developed for these devices (DesktopPlus,

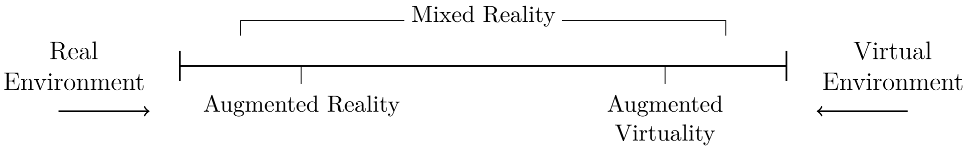

Figure 4.1: Milgram's reality–virtuality continuum.

2021; Virtual Desktop, Inc., 2021; Immersed Inc., 2021), they have not yet seen significant uptake with consumers. In this section, aspects of virtual environment experiences, devices, and environment implementations will be discussed.

4.1.1 Virtual Environments

Virtual reality as a term can be used to describe a variety of different environments and experiences, independent of the hardware used. When considering the term in the context of the environment or experience presented to the user, different classifications have been used to describe these environments in terms of different aspects of the experience, providing more manageable terms for use in discussing these. One of the more well-known classifications is that given by Milgram et al. (1995), which proposed a continuum ranging from a completely real environment through to a completely virtual environment, with experiences falling somewhere along this axis. This reality–virtuality continuum (See Figure 4.1) describes any mix of the extremes of this continuum as "mixed reality (MR)". While this continuum broadly allows for all VR experiences and environments to be placed upon it, it has a primary focus on the visual experience, and does not explicitly account for aspects such as audio and haptic feedback (Skarbez et al., 2021). It also fails to provide sufficient granularity for separation of experiences that may appear or be experienced in radically different ways due to implementation, which has led to development of alternate or supplemental classification.

One of the earlier extensions to this continuum was the mediated reality continuum (See Figure 4.2), which added an additional mediation dimension to classification (Mann, 2002). This allowed for more granular classification of similar experiences, such as an augmented reality experience presented using an optical see-through headset and the same experience presented on an opaque computer display utilising a camera to mediate the experience. This addition of dimensions has proven popular, with many more classifications emerging to further extend this idea (Mann et al., 2018).

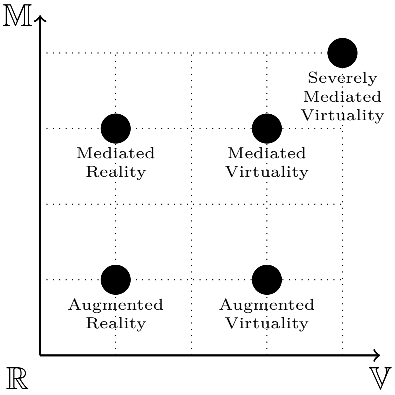

Figure 4.2: Reality, Virtuality, Mediality continuum as presented by Mann (2002). ℝ represents the real world and 𝕍 signifies virtuality, with the shared axis forming the reality–virtuality continuum presented by Milgram et al. (1995). 𝕄 represented mediality, the degree to which the experience is mediated or modified intentionally or accidently.

When restricting classification purely to the visual experience of an environment, the classification of alternate realities given by Want (2009) — visualised in Figure 4.3 — provides a simple set of terms that can be used to describe virtual environments. These terms are specified at intersections of the reality or virtuality of the primary experience, and the reality or virtuality of secondary overlaid data displayed. With this classification, the term “Virtual reality (VR)” describes an experience that is at the extreme of the virtuality scales for both primary and secondary displayed information, with no real elements included. The terms “Cross reality (CR)” and “Augmented reality (AR)” describe alternate and opposing extremes of these scales, displaying a combination of real and virtual content displayed for primary or supplemental consumption. “Mixed reality (MR)” continues to describe experiences that deviate from the extremes of each classification, following the usage of the term given by Milgram et al. (1995).

The classifications and terms given by Want (2009) will be used to describe environments throughout this thesis.

4.1.2 Virtual Reality Systems

Virtual reality as a term is also used to describe hardware utilised for presentation and interaction with virtual environments. While virtual reality is often associated with IVR HMDs The breadth of systems that can be described with this single term is vast, covering the “Goggles ’n’ Gloves” (Steuer, 1992) systems of body tracking and immersion, sophisticated and accurate vehicle simulators, and video games displayed on a conventional display. While some-

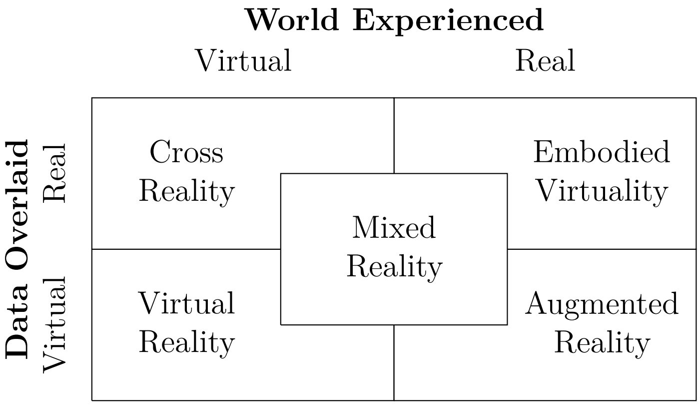

Figure 4.3: A classification of alternate realities, given by Want (2009), classifying alternate realities in terms of the source of the primary experience, and the source of overlaid data.

what dated, the classification of virtual reality systems given by Biocca and Delaney (1995) (See Table 4.1) provides a good overview of virtual reality systems upon which virtual environments can be displayed. In this classification 6 distinct virtual reality systems are outlined, with these grouped as free-standing or HMD systems.

Within this classification, four systems have seen use in literature as potential general computing environments, with the details of these systems expanded on below.

Window Systems

Window systems as described by Biocca and Delaney (1995) present a view into a virtual environment through a portal, which can take the form of conventional displays, hand-held displays, or other display technologies. These systems present an affordable approach to presenting virtual environments due to the use of existing and ubiquitous display technologies.

Window systems can present virtual reality environments directly to perceivers, and other forms of virtual environments are made possible when combining the display with additional peripheral devices. Use of external cameras allows for mediated augmented or cross reality environments to be displayed, with tablet systems including multiple cameras presenting simple self-contained window system devices for this purpose. Smartphone and tablet devices can be seen supporting augmented reality applications for consumers, with games such as Pokemon Go, and design applications like IKEA Place providing use-cases for this technology. When considering pure virtual reality environments, 3D videogames such as Quake and Second Life are common examples that can be seen presented on conventional displays.

Table 4.1: Classification of Virtual Reality Systems. Reproduced with adaptations from table published by Biocca and Delaney (1995, p. 59).

System Type	Description
Window	A computer screen provides a window or portal onto an interactive, 3-D virtual world. Desktop computers are often used and users sometimes wear 3-D glasses for stereoscopic effects.
Mirror	The users look at a projection screen and see an image of themselves moving in a virtual world. Video equipment is used to record the user's body. A computer superimposes a cut-out image on a computer graphic background. The cut-out images of themselves on the screen mirrors their movements, hence the name *mirror systems*.
Vehicle-based	The users enter what appears to be a vehicle (e.g., tank, plane, car, space ship, etc.) and operate controls that simulate movement in the virtual world. The world is most often projected on screens. The vehicles may include motion platforms to simulate physical movement.
Cave	Users enter a room or enclosure where they are surrounded by large screens that project a nearly continuous virtual scene. 3-D glasses are sometimes used to enhance the sense of space.
Immersive virtual reality (IVR)	Users wear displays that fully immerse a number of the senses in computer generated stimuli. The stereoscopic HMD are a distinctive feature of such systems.
Augmented reality (AR)	Users wear a visual display (e.g., transmissive HMD) that superimposes 3-D virtual objects on real-world scenes.

Virtual environments presented through window systems have been utilised for many different purposes outside of the entertainment applications mentioned above. Training applications are popular with these systems, with examples of social skills training for people with autism spectrum disorder (Bahiss, 2021), laparoscopic surgery training (Larsen et al., 2009), assembly and maintenance tasks (Gavish et al., 2013), and general education (Inman et al., 2010) showing the breadth of possible applications these systems support.

In the context of general-purpose computing, window systems have been heavily utilised when developing modern windowing systems and assistive technologies to support general computer usage. Simple virtual environments implementing the desktop metaphor have been augmented with 3D elements to provide improved and natural physical interactions (Agarawala & Balakrishnan, 2006) or representations of elements such as documents and docks. Existing assistive technologies such as virtual desktops and window switching have been improved with 3D features to allow for better awareness of the location and composition of the components they arrange (Henderson & Card, 1986; Agarawala & Balakrishnan, 2006; Warr et al., 2016). Complex 3D environments have been developed to investigate the impact of spatial memory and organisation on search tasks (Robertson et al., 1998; Cockburn & McKenzie, 2002), and complete environments utilising navigation of virtual structures have been considered for task management (Robertson et al., 2000).

CAVE System

CAVE automatic virtual environment (CAVE) systems are immersive systems in which a user is surrounded on one or more sides with display surfaces, presenting an immersive environment in which the user can move, and experience content shown on the surfaces. These were first introduced by Cruz-Neira et al. (1993b), with the initial prototype taking the form of “three rear-projection screens for walls and a down-projection screen for the floor” (Cruz-Neira et al., 1993b) forming a 7 cubic foot virtual environment contained in a $30 \times 20 \times 13$ foot room. Modern CAVE systems provide systems with single wall and floor configurations through to fully enclosed 6-sided immersive systems (Mechdyne Corporation, 2021a), using combinations of rear or front projectors (Visbox Inc, 2021) and large format displays at a variety of sizes (Mechdyne Corporation, 2021b). Mini-CAVE systems, describing systems small enough to place on a desk or other surface, are also available, though see less use than their larger counterparts (Carvalho et al., 2011). CAVE systems typically require large, dedicated spaces due to their size, and are expensive to purchase and install, with commercial systems ranging from 160,000–630,000€ in 2005, and low cost DIY solutions sitting around 20,000€ in the same time period (Juarez et al., 2010).

CAVE systems are primarily capable of displaying virtual reality environments, with support for including physical elements within the environment to act as input methods or targets for augmentation. Sensors are used within CAVE systems to track user and object position and orientation within the space, and this information is used to transform the content displayed on the display surfaces, and to provide input to the system. 3D stereo effects are accomplished using active or passive 3D glasses in conjunction with projected content to allow for a sense of depth, and to allow for objects to be displayed inside the space.

CAVE systems are extremely flexible, capable of supporting a wide variety of input methods, and displaying many kinds of content. They have found significant use in information visualisation (Svidt & Bjerg, 2002; Cruz-Neira et al., 1993a), visual prototyping (Choi & Cheung, 2008), and educational applications. These systems have demonstrating benefits for collaboration and improved learning outcomes (Back et al., 2020) when used for these purposes.

Immersive Virtual Reality HMD

IVR HMDs are headset systems that obscure the user's vision of the real world, instead presenting content viewed on displays within the headset. Within the headset, one or more display panels are arranged ahead of the user upon which content is displayed. Optical elements are often present between the perceiver and the displays, distorting the presented view to provide a sense of curvature that may not be physically present in the display arrangement. The position and orientation of the HMD is typically tracked, with this information used to recalculate the pose of the perceiver within the virtual environment, and from there the view that is presented on the displays. The amount of space required when using an IVR HMDs varies, with some devices supporting smaller seated experience tracking areas that require no more space than a current computing system, through to room-scale environments that utilise entire rooms.

IVR HMDs can present a range of virtual reality environments, assuming external sensors are made available. Virtual reality environments are easily presented within these HMDs, providing immersive experiences that isolate the perceiver from the external world. If equipped with external cameras, cross reality and augmented reality environments can be presented.

IVR HMDs have seen popularity with home consumers, providing entertainment applications including games, movie theatre experiences, and collaborative environments. These have also seen use in design, information visualisation, and education applications. general-purpose computing environments have been implemented using IVR HMDs as the display technology both in experimental environments, and in consumer available products. Multiple approaches have been considered for these general-purpose environments, from effective simulations of conventional computing environments (Virtual Desktop, Inc., 2021; Immersed Inc., 2021) with addition of novel input devices, through to environments that significantly deviate from convention and even expected physical realities.

Augmented Reality HMD

augmented reality (AR) HMDs are headset systems that can take similar forms to IVR HMDs, covering a large portion of the user's face, or as monocular or binocular "smart glasses". These devices project or directly display content on an optically see-through panel, allowing the perceiver to directly observe the real world through the display whilst having virtual elements overlaid or inserted into the view. A system that presents information as an overlay to the real world without giving the sense of the content being inserted to the world is often referred to as a head-up display (HUD). Many AR HMDs track position and orientation of the device as with IVR HMDs to update the view of the virtual environment projected on the display. Space requirements for AR HMDs vary as with IVR HMDs, though they are often less limiting in the case of AR, allowing for their use freely anywhere in the world or limited to a building or room.

AR HMDs primarily display augmented reality environments due to their additive overlay nature, but can present mixed reality or cross reality environments in some cases. These devices have seen use in many professional fields, with medical, engineering, product design, and education applications showing some of the strongest uptakes and benefits. general-purpose computing environments are available for some AR HMDs, with the Microsoft Hololens presenting a functionally complete stand-alone computing experience that is commercially available, and other systems that have been demonstrated primarily in academic literature (Reichlen, 1993; DiVerdi et al., 2003).

4.1.3 Head Tracking Systems

CAVE systems, IVR HMDs, and AR HMDs as discussed above make use of orientation and positional tracking solutions to provide 3 degrees of freedom (3DoF) or 6 degrees of freedom (6DoF) pose information for the perceiver's head, allowing for the virtual environment to be updated to maintain a consistent or expected view. There are many different methods for performing head tracking, with different trade-offs that make each suitable for specific applications. In the context of the discussed CAVE and HMD systems, head tracking systems are broadly categorised as inertial or visual systems (Gourlay & Held, 2017), with visual systems further subdivided as "outside-in" or "inside-out" systems.

Inertial Tracking

Inertial tracking is typically accomplished using an inertial measurement unit (IMU), which is attached to the perceiver's head as part of an HMD, or as an independent device. IMUs utilise a combination of accelerometer, gyroscope, and magnetometer sensors (Ahmad et al., 2013) to measure changes in position and orientation of the unit. Calculations are performed on the data outputs of each internal measurement device, producing absolute or relative rotation and position information from moment-to-moment measurements (Zhao & Wang, 2012). IMUs report data at a high frequency, in excess of 1000 Hz in some cases (Bosch Sensortec GmbH, 2021), allowing for small movements, or oscillations to be detected and measured and provided in a timely fashion.

Low cost IMU devices suffer from measurement drift due to accumulated error from many frequent and noisy sensor readings over time (Zhao & Wang, 2012). While drift in some tracking axes can be corrected using additional sensors (Wittmann et al., 2019), appropriate sensors are not always practical or usable where IMUs are in use (Bai et al., 2020). As such, inertial tracking is generally used primarily for orientation tracking only, with any positional information being used in greater sensor fusion applications.

Outside-in Tracking

Outside-in tracking describes solutions where the perceiver's head position and orientation is determined by an external system, with the any equipment attached to the head performing minimal or no tracking of it's own. These solutions typically use fixed external visual sensors (Marks et al., 2014), though other non-visual solutions exist, and observe a defined region within which the head and other elements can be tracked. The head, and any other tracked elements, may be fitted with fiducial markers to aid the external sensors in

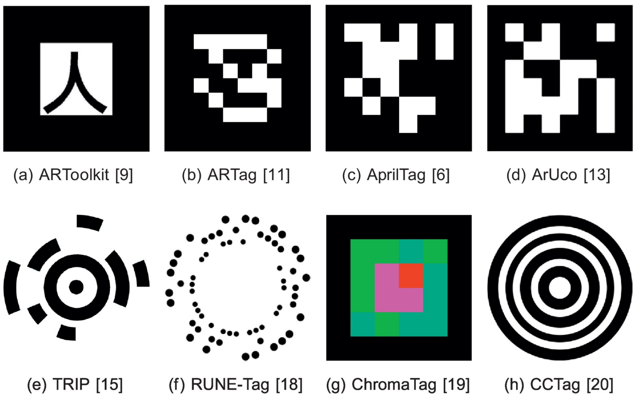

Figure 4.4: Fiducial markers used by marker systems (Benligiray et al., 2019, p. 2).

accurate and low latency tracking of the items in the region. Markers may be active, such as infrared (IR) emitters, or passive, such as retro-reflective points and visually distinct symbols (See Figure 4.4). The use of fiducial markers is not required, and marker-less tracking systems are available, though a performance impact in precision and processing time is noted for these solutions.

These systems are suited for fixed installations with constrained tracking regions, such as seated experiences, where the degree of required movement is limited. Several examples of consumer IVR HMDs make use of outside-in tracking systems, with the Oculus Rift CV1 (CV1) and Sony Playstation VR being notable examples of active marker-based solutions, and the Microsoft Kinect being an example of marker-less solution. Commercial motion capture systems, as well as CAVE systems, make use of both active and passive fiducial markers to accurately track many points at high speeds (Moeslund & Granum, 2001, p. 239; Kida et al., 2000, §4.1).

Inside-out Tracking

Inside-out tracking solutions utilise sensors attached to the perceiver's head or to a work device, where these worn sensors determine the position and orientation of the head either completely independent of external systems, or utilising information broadcast by external stations. As with outside-in tracking, the utilised sensors are typically visual, but alternatives solutions exist. Inside-out solutions utilising supporting external stations may be restricted to defined tracking regions where the appropriate stations are present, limiting these solutions to similar tracking environments and areas seen with outside-in tracking. Inside-out tracking solutions implementing forms of simultaneous localisation and mapping (SLAM) are capable of tracking the device in unknown areas (Chekhlov et al., 2007), whilst others may be limited to previously encountered

Figure 4.5: Valve VR test room, showing fiducial markers used for positioning early prototypes of IVR HMDs (Kraft & Make Community LLC, 2016).

and mapped environments (Gourlay & Held, 2017). As with outside-in tracking solutions, fiducial markers may be used to improve tracking performance, identifying reference points for environmental positioning (See Figure 4.5), or identifying elements that can be interacted with (DiVerdi et al., 2003).

Inside-out tracking solutions are suited for both scenarios where the head pose needs to be determinable in many different unbounded locations, and for tracking inside a defined region. Utilising independent systems with no external support system may increase the complexity of processing, and accuracy of determined location, and as such may need to be used in a sensor fusion arrangement to reduce the impact of this if required. This tracking approach can be seen in many commercial virtual reality systems, with the HTC Vive utilising a “Lighthouse” IR laser-based tracking approach that utilises external broadcasters with infrared broadcast markers, and the Microsoft Hololens demonstrating a independent marker-less solution.

Sensor Fusion

Sensor fusion refers to the combining of data from multiple sensors “such that the resulting information is in some sense better than would be possible when these sources [are] used individually” (Elmenreich, 2002). In the context of VR tracking technologies, this most commonly involves the combination of high-speed inertial tracking data, with the slower visual sensor output. This approach is often used by HMD devices to address shortfalls of each of these technologies, such as obscuring of visual tracking fiducial markers or slow update speeds for visual systems being addressed by inertial data, and loss of fast movement results and accumulated error or drift of inertial data being corrected by visual systems (Hogue et al., 2004).

4.1.4 System Selection

When examining the above discussed virtual systems, the suitability of these for general computing applications in the context of the office workers observed in Chapter 3 can be considered. In the first instance, the use of window systems can be excluded, as these represent the existing system in use by the surveyed office workers. As evidenced by the lack of assistive tooling usage by these users, with a strong preference for increased physical space, development of further assistive tooling or navigable environments would likely result in a system that would function in a user study situation but would not be utilised in normal work.

In the Chapter 3 study, office workers were seen to be working in an office space that provided limited scope for expansion, meaning that the space requirements of a candidate system is of significant importance. With this restriction in mind, CAVE systems are not ideally suited as candidate systems due to their large, dedicated space requirements that would interfere with a variety of office tasks. When also considering the purchase cost for these systems, it can be imagined that these would not be made available to all office workers simply due to the required financial outlay, suggesting that this is not a viable consideration for further investigation.

Both IVR and AR HMDs show promise as virtual reality system candidates when considering these constraints. Both categories of system could be used in the bounds of existing desk or office space, and due to a resurgence in interest in these devices many affordable and publicly purchasable options exist. This reduces possible friction in adoption, as minimal modification will need to be made to the existing workspace and outfitting an entire office population with devices is not an outlandish possibility. As both options are viable, the selection is ultimately a consideration of availability and access to devices for development and research, with IVR HMDs presenting the greatest range of options for this researcher. For this reason, IVR HMDs are selected as the virtual system candidates for further general-purpose computing investigation.

While many modern commercially available IVR HMDs provide external cameras, usable for supporting cross reality (CR) or AR experiences, not all do. Older and more affordable HMD do not provide these cameras as part of the core device, instead requiring additional 3rd party devices and modifications to provide this functionality. To present a fair environment that may be utilised on any of these devices, a purely VR environment will be developed to run within these HMDs.

4.2 Virtual Environment Considerations

With IVR HMDs identified as the chosen research target, consideration of the design of a general-purpose virtual reality computing environment is required. Virtual reality environments presented within IVR HMDs have advantages over conventional displays, most notable of which is that elements represented within them are not constrained by physicality of an environment. Elements can be freely scaled, positioned, and distorted within the virtual environment, with the virtual environment and content presented within not required to follow the usual rules and expectations of real environments. This enables capabilities such as allowing rotations within the environment to transform the contents and space or presenting content that follows the perceiver in different ways.

While the freedom to develop such environments affords many opportunities, presenting an experience that follows natural laws and expectations of the user is important to enhance the feeling of presence within the environment (Slater et al., 1995), and to reduce potential disorienting and motion sickness effects (Berger & Wolf, 2018). In this section, the presentation of content within a virtual environment is discussed, with emphasis on presenting conventional applications in usable fashions. Discussion of prior work is given, and a target virtual environment design is presented.

4.2.1 Content Presentation

Presenting content within a virtual environment can be accomplished in several different ways, with decision being driven by the nature of content that is intended for display. Independent of the exact form, aspects such as restrictions of movement, orientation, location, and use of depth for elements within the environment should be considered (McGill et al., 2020). When restricting discussion to the display of conventional computing application content, additional considerations such as the shape of representations and the general workflow should also be kept in mind.

The shape and form of content can easily be varied when the material to be presented has been designed with display in a free-form virtual environment as a primary consideration, however it has been shown that moving away from 2D representations of content into 3D negatively affects the efficiency of conveying information (McKenzie et al., 2017). This suggests that use of existing conventional application window forms, typically 2D rectangular frames, will be suitable for conveying information within a virtual environment assuming that the content contained within is usable. Effort should not therefore be spent in modifying the overall shape of the content frame, instead making appropriate changes to content displayed upon them, and the position and appearance of these frames within the virtual environment. This leads to the simplest form of content presentation within the virtual environment, that of a rectangular window floating within the world, upon which conventional application content can be displayed. While simple and seeming not to take advantage of many of the available freedoms afforded by a virtual environment, this representation appears to be popular with designers of virtual environments (Ens et al., 2014b; Pavanatto et al., 2021; DesktopPlus, 2021).

While representing a conventional application window as a virtual window mirrors it's conventional computing counterpart, there are modifications that can and should be made to the virtual window representations to ensure they remain usable. McKenzie et al. (2017) noted that when designing user interfaces, planar representations of content that appeared on the outside of the perceiver's view appeared to skew, which felt "unfair" to the content displayed upon it. The suggested modification to the presentation to address this was to distort planar content such that it curved around the perceiver in a cylindrical form, with the centre-point of the cylindrical surface being at the perceiver's position, or some other point within the environment (McKenzie et al., 2017). This design decision is reinforced by others, with indications that curved surfaces were considered preferable when considering accessibility of content and reduction of fatigue (Shupp et al., 2009), with "the degree of the curvature [being] somewhat up to personal preference" (Endert et al., 2012), but it was suggested that utilising a curvature where each point of the displayed item is of equal distance from the perceiver would be best. Benefits of curved content on comfort, legibility of content, and speed for visual search tasks have been identified (Shupp et al., 2009; Park et al., 2017), further suggesting that curved surfaces should be preferred over planar.

While distorting planar content can address the overall appearance of content, the size of content to be distorted is also of concern, especially when noting the preferred strategy of window management shown in Chapter 6 and the size of virtual window this would imply. Currently available IVR HMDs have limits to the field of view (FOV) presented to the perceiver, limiting the amount of content that can be displayed in a static view of the environment. Microsoft Mixed Reality guidelines suggest restricting content to a size that can comfortably fit within the headset FOV (Microsoft Corporation, 2020), with Ens et al. (2014b) identifying benefits for task times when content occupies less than 75% of the headset FOV to reduce head movements. This consideration of head movements is also identified as a concern for comfortable viewing of displayed elements, with Dingler et al. (2018) identifying a user preference for a slight negative angle of vertical head movement when consuming content, and Kim and Shin (2018) further emphasising minimising head movements to aid with comfort and task performance.

The use of depth is an additional consideration within virtual environments, with a 3rd dimension for placement presenting many opportunities for novel interactions. It is noted however that there are concerns when utilising this dimension, with eye vergence differences when focusing on elements within the environment miss-matching with the ideal vergence distance for the HMD leading to visual discomfort (Shibata et al., 2011). It is further noted that adding additional visual information to process through parallax effects and other interactions between elements of differing depth further increasing visual fatigue (McKenzie et al., 2017). Use of mixed depth content opens many different presentation possibilities, such as moving lesser used reference material further from the perceiver, however Tan and Czerwinski (2003) identified that mixing display depth of elements may negatively affect performance. This suggests that while depth could be utilised, it should be used with caution so as to avoid these potential pitfalls.

4.2.2 Bounded or Free Windows

When representing conventional application windows as virtual windows discussed above, a consideration to be made is how these should be placed, arranged, and organised within a virtual environment. When displayed within a conventional computing environment window system, windows are bound by the physical constraints of the display, with assistive tooling and optimisations made to better align and arrange windows within these constraints. The physical constraints presented by conventional monitor frames are not necessarily present within a virtual environment, so alternative options exist.

When considering general-purpose virtual environments, three primary options are seen in commercial and experimental systems that present as possible candidates for investigation. These options are treating the virtual environment as a continuous surface with no “edges” to impose constraints on placement, subdividing the space into virtual “pigeon holes” into which windows can be placed, and that of presenting virtual display “containers” containing virtual windows constrained in a way similar to conventional computing environments.

Continuous Display Space

The idea of a continuous display space presents as independent items that can be freely positioned within a large or potentially infinite surface or space. When considering this idea, several potential benefits and drawbacks can be identified. Presenting free-form placement and sizing of items within a continuous space allows users to develop novel organisational strategies, grouping tasks visually to collect ideas or related content (McGill et al., 2020). When considering collaborative virtual environment, using free form placement could allow for collecting items around real or virtual landmarks, supporting visualisations being anchored to relevant objects or persons (Sharma et al., 2019). Obvious drawbacks of such an approach can be seen when considering the arranging items into organisational units in a fast and efficient manner. With existing bounded systems, natural organisational areas can be identified as individual frames, and items can be arranged easily within these to occupy space efficiently and easily. In an unbounded space, aligning items such that they did not obscure relevant content from other applications when positioning or resizing others would be difficult (McGill et al., 2020), which may drive user demand for assistive tooling to aid in these operations.

Grid Placement

Compartmentalising the available virtual space into a “pigeon hole” grid within which virtual windows can be placed whilst avoiding occlusion in a manner similar to a widget placement grid, could be used to cleanly arrange elements (Pygmal Technologies, 2017). This form of arrangement may place limits on the size of content that can be displayed if windows are required to exist in a singular cell, or to span a rectangular range of cells which may remove some advantages of large space arrangements.

Alternative grid arrangements could also be considered, such as presenting strict environment-fixed grids, or ad-hoc grids supporting grouping of elements around key windows. Using such grid arrangements would allow for free window arrangements, allowing for emergent task arrangement strategies, whilst applying some order and manageability to the environment.

Containerised Windows

Introducing larger frames within which windows can be grouped and arranged, analogous to a display and desktop frame or virtual desktop view, which can be freely arranged in a continuous space presents a solution with many advantages. By applying existing familiar window management techniques and assistive tooling within each container, existing user behaviours and knowledge can be replicated and leveraged such as utilising the container edges for layout (Wallace et al., 2006), or taskbar usage for window identification. This interaction and constraining of windows could be re-implemented within the virtual environment or could be captured directly from an existing desktop environment representation through desktop mirroring APIs, handing off the primary management to the extant desktop environment. This approach has seen usage in both academic (McGill et al., 2020; Dominic et al., 2020) and commercial (DesktopPlus, 2021; Immersed Inc., 2021; Virtual Desktop, Inc., 2021) virtual environments, in part due to the simplicity of implementation of such mirroring. Hybrid approaches could be considered for layout, using this form of containerised windows initially with the option of removing windows from the container to act as independent items within the virtual space, a feature that is now seen in commercial offerings (DesktopPlus, 2021).

4.2.3 Environment Reference Frames

Independent of the choice of bounded or free windows, consideration of how presented elements are positioned within a virtual environment is critical. Considering how windows will be accessed, addressed, and switched between will be necessary when designing a virtual environment to ensure that these are approachable, and fit with the use-case presented by users of the system. Ens et al. presented the idea of "Ethereal planes", used to describe 2D planar elements within a virtual world, and presented terminologies and approaches to use when considering the placement of these.

Four different reference frames were presented by Ens et al., shown in Figure 4.6. These reference frames describe how virtual windows are placed with respect to the user or to the world and positioned with respect to the user FOV.

Exocentric Frame

When considering virtual window placement relative to the world, the term "exocentric" or grounded/world-locked (Lee, 2017) is used. Exocentric placement of content has mappings to real world object placement, such as the position of conventional displays upon a desk, or a book on a shelf. Exocentric content can be placed in a position within a real or virtual world, where it will remain independent of the perceiver's position within the same. This positioning has been shown to have benefits in leveraging spatial memory (Montello et al., 2004; McGill et al., 2020), aiding recall of object locations within an environment. Virtual content may be placed next to real or virtual objects that are associated with this content to take advantage of this, such as posi-

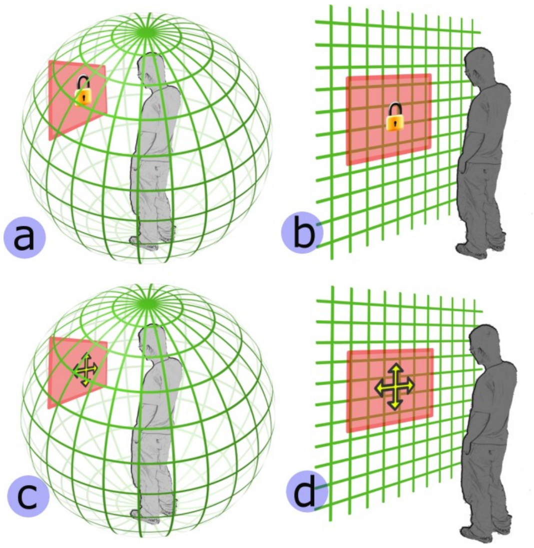

Figure 4.6: Reference frames used for classifying perspective and movability aspects of ethereal planes. Diagram sourced from Ens et al., 2014a, p. 4.

tioning a virtual phonebook beside a physical desk phone, or a digital booking system positioned beside or on top of a calendar (DiVerdi et al., 2003). Overlapping virtual content over physical representations may also be utilised to provide elements of tactile feedback (Wirtz et al., 2014), or to enable greater collaborative spaces through common placement of shared elements.

Exocentric content does present some downsides, especially with regards to location of lesser-used virtual elements and transitioning between workspaces. As elements are positioned within a given environment, moving from that environment introduces problems when the user wishes to access that content, requiring the user to return to the space where the content is positioned — a close relation to placing physical objects within a given room — or to transform a representation of that content for display within a new environment. With augmented reality environments, moving between workspaces may involve physical movement to other locations in the real world, while virtual reality environments may require navigation inputs within the environment to move the virtual camera position. AR solutions such as the Microsoft HoloLens, utilise sensors on the HMD to identify locations and restore workspaces when reentering a space (Microsoft Corporation, 2021a), whilst other systems rely of fiducial markers to establish the space and location of restored elements (DiVerdi et al., 2003). While this behaviour is beneficial in many ways, allowing natural associations between tasks and locations, it does require that workspaces be recreated if the appropriate space is not conveniently accessible. Ens et al. provided investigation into transforming content laid out in one space for display in another, comparing user performance in spatial location tasks when the layout algorithm prioritises visual constancy over visual saliency (Ens et al., 2016). It was noted that maintaining relationship of window positions with respect to each other was more important than respecting visual saliency when performing these transforms, however users did identify that context was lost when moving to a different space. Specific examples of this context loss were provided, with an example of items positioned on a computer desk due to this arrangement “fit[ting] the office paradigm” (Ens et al., 2016) losing context when moving to a different space that lacked a corresponding desk space presents a relevant scenario when considering design for office work.

Egocentric Frame

The alternative positioning approach, that of egocentricity or body-locked/tag-along behaviour (Lee, 2017), describes elements that are positioned with respect to the user position. Egocentric windows can be placed around the user, and will follow the user position, maintaining relative distance and orientation to the user independent of motions made. This style of positioning allows for a spatially consistent and familiar layout of content to be presented to the user regardless of the user’s location. Egocentric display of content has been shown by Hinckley et al. (1994) as being beneficial when dealing with free space interactions, with egocentric coordinate systems being easier to comprehend and manipulate than exocentric systems.

Downsides are present with egocentric frames as with exocentric, with the primary issue relating to collaborative constraints in multi-user environments. As egocentric content is inherently laid out for personal consumption due to their positioning relative to the perceiver, sharing content with others may require transitioning egocentric windows into an exocentric frame (Ens et al., 2014a).

Movability

Ens et al. (2014a) further classifies each of these reference frames based on the movability of content, classifying items as fixed or movable within the egocentric or exocentric frames. Fixed position elements do not move with respect to their coordinate system, with a fixed-egocentric element always being visible at the same position within a user’s FOV regardless of the user movements, and a fixed-exocentric element remaining fixed to a designated surface or object within the environment. Movable elements by contrast can be moved with respect to their reference frame, allowing for relocation of elements in response to actions or movements, for example a movable-egocentric window could describe a virtual watch face that is always shown oriented according to the user's wrist position in space (McGill et al., 2020).

Different combinations of reference frames and movability are seen to be suitable or unsuitable for different kinds of interface, for a variety of physiological and interaction reasons. Fixed egocentric content, HUD style content that remains at the same position within the user's FOV, is typically regarded as poor choice due to impacts on immersion within presented environments (Lee, 2017), and causing a sense of user discomfort or simulator sickness due to a miss-match in perceptions (Microsoft Corporation, 2020). Use of movable egocentric content can address the issue of user discomfort but can impact on immersion within the presented environment if not correctly integrated into the experience (Lee, 2017).

Adaptive Frames

Consideration of reference frames is important when designing a virtual environment, as this choice will dictate many aspects of the environment presentation and interaction, and how this meshes with the environmental use-case will be important for adoption. If choosing an exocentric reference frame, then users of a system may be required to physically move around a space to address different interface components, while with an egocentric frame the information will always be attached to the user. It is possible to consider a combination of the two reference frame styles, where the environment is fixed to the position of the user's head, while being accessed and interacted with in an egocentric manner. This approach was shown as optimal for the design of the personal cockpit (Ens et al., 2014b), suggesting that for mobile use the environment should attach to the user's head becoming egocentric when movement is identified, and becoming exocentric based around the position of the user's head position when stationary.

4.2.4 Environment Selection

When considering the design of a virtual environment for office workers observed in Chapter 3, several of the above discussed approaches can be considered as possible routes for development, while others can be discarded due to interaction with the IVR HMD devices, and purely VR environments chosen as virtual systems targets.

Considering first the environment reference frames, it is apparent that egocentric reference frames are ideal for applications where the user will be moving while interacting with content, whereas exocentric frames are more suitable for situations where the user will remain in one room. User movement while utilising the environment developed for this research is not considered a requirement, as the pure VR environment will isolate the user from the surroundings, necessitating removal of the device prior to leaving their workspace. Additionally, commercial IVR HMD solutions are primarily tethered to a computer system with a cable, restricting the movements of the user to the length of cable provided, and supported tracking region. This suggests that an exocentric reference frame will be suitable for the office worker application, and that an adaptive frame need not be considered due to the limited movement expected. While exocentric environments are often seen as fixed to real environmental landmarks, a pure VR environment displayed on a device without external sensors to detect environment features will be unable to establish this relationship. Anchoring the environment to the neutral position of the user's head when seated at their desk presents a world that can establish rough position and orientations relative to their desk, and items placed upon it.

Having chosen an exocentic frame and world fixation point, how elements within this reference frame need be decided. Referring above to the content presentation discussion, several aspects of presentation in this chosen frame can be chosen. The discussion identifies that presenting content at similar depths to avoid shift in vergence distance is preferred for both performance and comfort, which suggests that presenting content along a single surface would provide an acceptable experience for the perceiver. Utilising a planar surface would not be advisable, as it is identified that this negatively impacts perception of content displayed on it away from the user's forward position and would be unfair to that content. While the user could physically reposition themselves to address this content such that it appeared directly ahead of them, it has been identified that movement is not a priority in this scenario due to device and environment choices. A better solution evident from the above discussion is to distort the presentation surface around the user, forming a cylindrical or spherical surface that presents all content at equal distances from the perceiver's initial position within the environment.

Examining the state of existing general-purpose virtual reality computing environments, the use of containerised windows through mirroring of conventional desktops is a popular choice. This popularity as discussed appears to be due to the simplicity of implementation, and the ability to leverage existing use knowledge and behaviour with window management tools. While following this trend could be justifiable, observing that the studied office workers did not leverage many of the existing window management functionality indicates that the familiarity aspect of this approach may not benefit them significantly. Consideration is then given to the approach of a continuous display space and grid placement approach, where it is seen that the grid placement approach is simply a layered constraint over a continuous display space. Introducing additional constraints to an environment before usage patterns have been established suggests premature optimisation, while leaving the environment unconstrained would allow for emergent usage strategies to be observed. Therefore, the choice is made to utilise a continuous display space for window placement initially, with the option of introducing additional grid-like constraints to support interactions observed by user usage of the system.

As a continuous display space is to be implemented, this requires that content presentation be more sophisticated than simple conventional display mirroring onto the environment. To allow for freedom of use of the environment, each running conventional application window should be treated as independent items within the environment, with the environment acting as a window manager for these, rather than as an environment to contain virtual displays.

4.3 Evaluating Suitability

When considering suitability for IVR environments, many different metrics have been used, considering different elements depending on the use case of the system. If the usage is isolated to a system capable of displaying conventional application windows as described above, then usability concerns for managing these application windows should be considered as metrics. When considering the activities performed by office workers presented in Section 3.5, identification of target windows, application lifecycle management, window arrangements, and content consumption were the primary interactions that the observed office workers had with their environment.

Each of these observed activities serve as potential evaluation metrics, with impacts on each of these when moving to a virtual environment playing a part in the fitness for purpose of the virtual environment replacing the conventional computing environment currently used. When examining these activities, that of content consumption stands out as the most impactful task on user time, given that this was seen to occupy most of the office worker's time, and would thus present the potential performance impact if it were made easier or harder in a new environment. Within the content consumption task category, reading of various forms of text was identified as the most common task performed during the day, which is supported by White et al. who identified that on average, 312 minutes per working day (65% of an 8-hour working day) is spent on reading tasks, with a mean time per activity of 31 minutes (White et al., 2010, p. 291).

Given the apparent importance of reading in daily tasks, the performance and accuracy of reading within the virtual environment when compared to a conventional computing environment will serve as a primary metric on environmental suitability for this task. In this section, an overview of the physiology of reading is given, and potential impacts of introducing an IVR HMD into the reading process are presented with relation to existing text presentation guidelines when designing for VR consumption.

4.3.1 Eye Physiology

The human eye is a complex structure with meany discrete elements. The primary purpose of much of this structure is to focus light onto the retina in a manner similar to that of a camera focusing an image onto film (Tortora & Derrickson, 2009, p. 613). The retina contains two varieties of photoreceptors — rods and cones — which convert light rays into nerve impulses (Tortora & Derrickson, 2009, p. 610) which are eventually interpreted as sight. Rods and cones have different purposes, with rods allowing for vision in low light situations and resolving black and white images, and cones operating in better lit situations resolving colours through specialised red, green, and blue sensitive cones (Tortora & Derrickson, 2009, p. 610).

The retina contains a region called the macula lutea, which in turn contains three concentric regions known as the fovea, parafovea, and perifovea, which are responsible for high acuity vision (Purves et al., 2018, p. 237). The central fovea contains the highest density of cones within the eye, and light focused on this region will be seen with the greatest detail. The surrounding parafovea and perifovea have decreasing densities of cones, with light focused on these regions being resolved with less precise detail than is seen on the fovea. In order to precisely resolve target objects, the eye needs to move to bring the high-resolution fovea to bear on the target (Findlay, 2008).

4.3.2 Eye Movement

The human eye is actuated by six fast-twitch oculorotary muscles (Demer, 2006) which allow for rapid rotations of the eye with minimal translation (Carpenter, 1988). Eyes can move independently (duction) in different directions, or simultaneously as a pair in the same direction (version) or opposite directions (vergence) (Parks, 2006). Duction eye movements allow for broadly symmetric horizontal movements (44.9 ± 7.2° in adduction, 44.2 ± 6.8° in abduction), with vertical movements demonstrating significantly greater elevation range than depression (27.9 ± 7.6° in supraduction, 47.1 ± 8.0° in infraduction) (Shin et al., 2016).

These oculorotary muscles allow the eye to make five kinds of movements, categorised as either *shifting* or *stabilising* movements (Purves et al., 2018). *Stabilising* movements, which are vestibulo-ocular and optokinetic movements, allow the eye to maintain focus on a target when the head moves, or when other forces and actions change the visual field. *Shifting* movements, which are smooth pursuit, saccade, and vergence movements, support the ability for the eye to follow the movements of a currently focused object, or for focusing on new targets. Saccades, which take the form of "rapid, ballistic movements of the eyes that abruptly change the point of fixation" (Purves et al., 2018, p. 450), are the primary shifting movements associated with reading.

Saccades movements are rapid, moving at speeds of up to 500° per second, accelerating from the starting point reaching a maximum speed prior to the midpoint of the movement, before decelerating to the target (Rayner, 1998). These movements require planning, considering the distance between current fixation and the target, which is then passed to the oculorotary muscles to execute (Purves et al., 2018). The movements are described as "ballistic", as they do not adjust for target movements after the saccade movement has begun, meaning additional saccades must be performed to adjust for target movements during this time (Purves et al., 2018, p. 451).

4.3.3 Saccades in Reading

When reading, the eye continually makes saccade movements covering 7–10 letter spaces (Inhoff et al., 2005), followed by periods of relative stillness known as fixations (Rayner, 1998). Content is seen only during fixations, with vision suppressed during saccades due to the phenomenon of saccadic suppression, which inhibits visual processing during the saccade (Duyck et al., 2016). The duration of fixations and distances of saccades varies on a range of factors including task type (See Table 4.2), narrative perspective shifts (Ballenghein & Baccino, 2018), and plausibility of read text (Rayner et al., 2004). Saccades movements made when reading English text are predominantly left-to-right, with right-to-left saccades occurring during regressions and return sweeps.

Table 4.2: Approximate mean fixation duration and saccade distance on different tasks. Table reproduced from Rayner (1998).

Task	Fixation Duration (ms)	Saccade Size (degrees)
Silent reading	225	2 (≈ 8 letters)
Oral reading	275	1.5 (≈ 6 letters)
Visual search	275	3
Scene perception	330	4
Music reading	375	1
Typing	400	1 (≈ 4 letters)

Regressions

Regressions account for 15–25% of eye movements during reading and are defined as movements against the reading direction of the text (Booth & Weger, 2012). Most regressions are short-range, occurring at the word immediately prior to the previously fixated word (Vitu & McConkie, 2000), while others are longer-range, returning to earlier segments of text (Booth & Weger, 2012). While the exact function of regressions is largely unknown, larger numbers are seen when reading complex text. This suggests that they serve, in part, to allow for re-reading of text to aid with understanding (Blanchard & Iran-Nejad, 1987), or to resolve ambiguities in read text (Inhoff et al., 2009). Long-range regressions have been noted as very accurate, reaching their target in a single saccade (Inhoff et al., 2009).

Return Sweep

Return sweeps are saccade movements against the direction of reading that direct the gaze from the end of one line to the beginning of the next. Return sweeps are initiated prior to reaching the end of the line, and land inwards from the start of the next line, with Parker and Slattery (2019) observing the majority of return sweeps launching 8 characters from the end of the line and landing approximately 5 characters from the start of the next line in adults, with children launching return sweeps closer to the end of the line and landing closer to the start of the next (Parker & Slattery, 2019). These saccades are necessarily longer than saccades normally observed during reading, and error is often noted with the saccade undershooting the target requiring corrective saccades (Parker et al., 2019). The frequency of these errors seems to be affected in part by line length, with shorter line lengths resulting in fewer errors (Beymer et al., 2005).

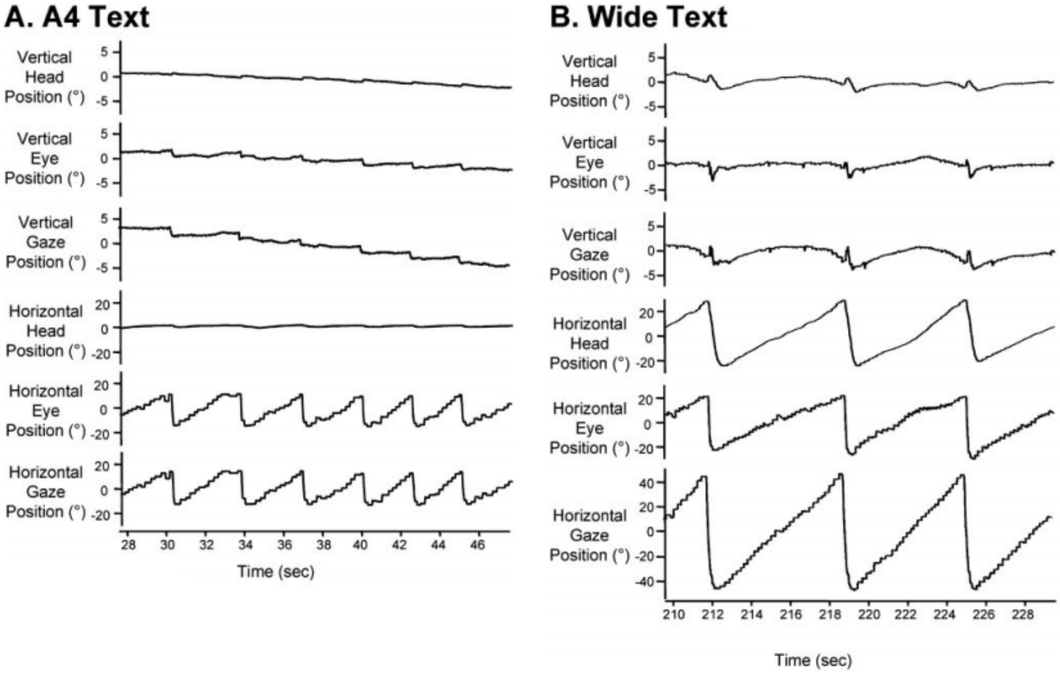

Figure 4.7: Horizontal and vertical eye, head, and gaze positions observed during reading tasks from an A4 and Wide card. Graphs from Proudlock et al. (2003).

4.3.4 Head Movements during Reading

Head movements are also associated with eye movements during reading, with wider text samples demonstrating greater amplitudes of head movements than is seen with narrow samples. Proudlock et al. (2003) conducted a study measuring head and eye movements during reading tasks from two different sized samples. One sample was read from an A4 page of text positioned 33cm from the eyes (covering $\approx 28.6°$ of visual angle), while another was read from a wider card occupying 90° of visual angle, horizontally curved at radius of 33cm. Horizontal head movements when reading from these samples accounted for 4.7% of gaze shift on the smaller A4 card and 40.3% of gaze shift for the wider card. These horizontal head movements were noted as linear when tracking left-to-right reading the content, with no observed periods of no movement (Proudlock et al., 2003). This linear movement was noted as a possible movement control strategy specific to reading. Vertical head and eye movements varied when the gaze performed a return sweep, with significant deviation noticeable for the large sample (See Figure 4.7), suggesting additional effort was required to identify the next line of text when performing a large return sweep.

Proudlock et al. (2003) noted that results presented reinforced the idea of the head and eye movements are highly coupled and adaptable, able to be activated to aid with reading or suppressed when support is not required. In the extreme, this can be seen in a study by Gilchrist et al. (1997) which observed reading performance in a participant suffering from extraocular muscle fibrosis, who has been unable to move their eye since birth. This reader was noted as making saccade movements while reading that matched those of a normal reader, including regressions and return sweeps, using only their head and neck to perform gaze shifting. This participant was able to achieve slow, but not abnormal reading speeds despite this restriction in movement.

4.3.5 Reading within HMDs

Reading within HMDs is seen frequently in literature, with many experiments or applications requiring reading tasks to greater or lesser degrees. When designing experiences for display in a virtual environment for display in an HMD, there are many guidelines that have been established to ensure that reading is achievable. Several studies have investigated the effects of different background and foreground colours of text (Dingler et al., 2018; Jankowski et al., 2010; Leykin & Tuceryan, 2004), as well as line width of paragraph text (Dingler et al., 2018), and geometry of the surface upon which the text is displayed (Wei et al., 2020). When these guidelines are followed, it has been shown that text is readable, and users can comprehend the text at comparable rates to those observed on conventional displays (Rau et al., 2018). Alternative presentations of text, such as rapid serial visual presentation (RSVP) have also been demonstrated to have promising effects when used for short texts (Rzayev et al., 2021). Guidelines also emphasise minimising head movements when reading to prevent physical discomfort, and simulator sickness symptoms (Kim & Shin, 2018).

Reading within a virtual environment from text presented in a manner that was not designed for this presentation however is a lesser explored area. Text content designed for presentation on conventional displays utilise a longer line length than is recommended for virtual environment presentation, and often in higher densities than would be considered ideal. When presented at a readable distance from the perceiver whilst retaining the expected display size, this may result in content displayed much like is seen in the wide card sample demonstrated by Proudlock et al. While Proudlock et al. (2003) showed that head movements can be made when reading wide content without major effect on the reading performance, HMDs can introduce additional limitations to peripheral vision that is not present when reading from a conventional display, or physical paper. This limitation may affect performance and accuracy of regressions and return sweeps when reading within these devices, affecting reading performance in measurable ways.

Many conventional applications allow for changes to be made to the text presentation within them, with items such as background and foreground colour, font choice, and content width being adjustable within applications, or at a system level. For supported applications, this may allow for simple user configurable changes to be made to applications to allow them to better conform to presentation guidelines for VR. There are however many legacy applications for which this is not possible, or is unreliable, so determining the impact of non-ideal text presentation on reading performance should still be determined.

An further complicating factor of reading within an IVR HMD is their utilisation of lenses between the eye and displays, which make the image displayed on the near display appear distant enough for the eye to comfortably focus on (Robinett & Rolland, 1993) amongst other purposes. While lenses carefully manufactured and positioned for the user would introduce minimal issues when viewing content through them, commercial HMDs utilise cheaper and more general optics designed to accommodate a range of users rather than specific individuals (Luckey, 2019). Aberrations in these lenses can introduce a multitude of effects that may influence clarity of vision when gaze shifting to different points in the lens, with these effects possibly influencing the performance of reading. Three of these aberrations are of note when considering commercially available HMDs and the lenses utilised within them.

Spherical Aberration

Spherical aberrations are caused by light passing through different areas of a lens converging at different focal points (Thibos et al., 2013). With an ideal lens, light passing through any region of the lens would converge to a single point, presenting a clear image when looking through the lens at this focal point, however with non-ideal lenses light passing through the lens further from the optical axis of the lens can converge at nearer or further focal points, leading to an image that is blurred (Thibos et al., 2013). When this effect is encountered, content viewed directly along the optical axis may appear clear, but as the eye moves from the neutral forward position the content will appear to blur.

If present, this aberration may influence the clarity of text read outside of a central area of the lens, possibly reducing reading speed and accuracy. Readers may adjust their reading process to include greater proportions of head movement, limiting saccades to small angular ranges to reduce this effect, or may attempt to read normally with the impacted text at a slower or less accurate rate than would be expected.

Defocus Aberration

Defocus aberrations are caused by non-ideal positioning of the eye relative to the lens, resulting in an image that appears out of focus (Burge & Geisler, 2011). This miss-match can be related to the distance of the lens from the eye, and the eye's position relative to the centre of the lens in two dimensions. Out of focus text may be difficult to read, resulting in a reduced reading speed when blur exceeds a threshold (Chung et al., 2007; Jainta et al., 2011), requiring text to be moved closer to the perceiver to present it in a legible fashion, and may benefit from making changes to the utilised font to one less sensitive to loss of clarity.

Many commercially available HMDs provide adjustments to help address this aberration, utilising both hardware and software solutions. Hardware inter-pupillary distance (IPD) adjustments that adjust lens position horizontally, eye-relief mechanisms that changes the distance from the eye to the lens, and movements of the device on the user's head represent common adjustments found on many HMD (Luckey, 2019). Software correction for IPD and relief are also available in many cases, allowing for simpler hardware construction when the properties of the lens is known. While these may help with the symptoms of this aberration issue, the available adjustments may not be sufficient for some users, lacking granularity or range of adjustment to fully eliminate the issue.

Chromatic Aberration

Chromatic aberrations are caused by the failure of a lens to focus all visible wavelengths of light onto the same point (Johnson & Farid, 2006), which can produce fringing of colours between light and dark boundaries within images. As text is typically contrasted from the background on which it is displayed, effects of chromatic aberration would be apparent if this aberration was present. It has been noted that colour fringing of words shown to the left or right of a subject fixation point has a significant effect of the "latency and accuracy of word recognition" (Yang et al., 2011) outside of HMD applications, suggesting that this may be a possible observation with HMD applications also.

4.4 Summary

In this chapter, an outline of the term virtual reality (VR) has been given, identifying both environments and hardware for which the term can be used. Suitability of different VR hardware systems for general-purpose computing is presented, with IVR HMDs identified as the device category of interest for further research. Aspects of virtual environment design has been considered, and a moveable-exocentric design allowing for seated usage and minimal novel interactions was identified as that most suitable for the office worker use-case seen in Chapter 3.

Possible evaluation metrics for determining suitability of a novel virtual environment as a replacement for conventional systems have been discussed, with reading performance identified as the primary metric due to the volume of text consumed by observed office workers. An overview of the physical processes involved in reading has been given, and possible factors affecting reading when conducted within an IVR HMD are presented.

Chapter 5

Virtual Environment Design

In the previous chapter, immersive virtual reality (IVR) head-mounted displays (HMDs) were identified as potential candidates for display devices that may address the concern of space limitations seen with conventional display configurations. Virtual environments that could be presented on these devices were discussed in the context of the office worker usages and office environments seen in Chapter 3, and a continuous display space holding free windows within a movable-exocentric reference frame was chosen as suitable for their existing physical environments and use cases.

This chapter describes the development and initial informal evaluation of a virtual environment meeting these specifications, and supporting software developed to assist with evaluation of the same. Initial exploration of window content capture is presented, and presentation of this content within the virtual environment is described. Specific attributes of two environment versions are explained, highlighting the differences in behaviour and features. Finally, observations made about the performance and design of the developed environment provided by users is discussed, identifying aspects to be evaluated and potential changes that could be considered.

The virtual environments and supporting software described in this chapter are used as the testing platforms in the studies presented in Chapters 6 and 7.

5.1 Design Requirements

In order to allow for performing a variety of potential experiments with immersive virtual environments, an environment was developed that allowed for conventional applications to be both displayed and interacted with. The construction of this environment was bound by several design requirements that were identified within previous chapters. In this section, these different requirements and constraints are detailed.

5.1.1 Environment Design

Section 4.2.4 identified different features of the developed environment, specifying a target reference frame of moveable-exocentric, utilising a continuous display space upon which independent free windows could be positioned and displayed. Windows were further specified to be distorted, such that they were distorted around the perceiver, with each point on the window surface equidistant from the environmental origin.

To enable potential testing of real-world workflows observed in the workplace study discussed in Chapter 3, it was decided that the developed environment should be functional for users. This requires that the state of displayed virtual windows remained consistent with the real application running on the host system, that the virtual window could be interacted with in the expected manner with a mouse cursor and keyboard, and that multiple applications be supported simultaneously to allow for performing tasks that span multiple application windows.

Section 4.3 further identified reading as the primary metric for evaluation of environmental suitability, so the developed system should enable experimental control of presented content and appearance to allow for controlled and repeatable testing. While reading is identified as the primary metric for evaluation, experimental controls should not be restricted to only enabling tests for this metric and should allow for granular control of the environment and robust logging to support possible alternative evaluations.

5.1.2 Virtual Reality Device Target

Section 4.1.4 identified a pure IVR HMD as the target device, due in part to availability of devices and as this solution could be utilised within the existing seated workspace constraints utilised by observed office workers in Chapter 3. The Oculus Rift family of device was selected as the target for development, as access to an Oculus Rift DK1 (DK1) device was immediately available. The decision was made at different stages of system development to continue utilising this family of devices as new and alternate options became available, due to the Oculus Rift family offering comparable devices to alternatives as they were released, and to minimise development requirements to support different device APIs simultaneously.

Devices utilised were restricted to the Oculus Rift DK1 (DK1), Oculus Rift DK2 (DK2), and Oculus Rift CV1 (CV1), each of which represented competitive solutions at their time of release. As a seated experience was targeted, only tethered HMDs were considered for support, with untethered and mobile solutions excluded. A comparison of specifications for these devices is shown in Table 5.1.

Table 5.1: Comparative specifications for Oculus Rift IVR HMDs supported as development targets.

	DK1	DK2	CV1
Rotational Tracking	Yes	Yes	Yes
Positional Tracking	No	Yes	Yes
Weight	380g	440g	470g
Display	RGB LCD	PenTile OLED	PenTile OLED
Resolution (per-eye)	640 × 800	960 × 1080	1080 × 1200
Refresh Rate	60Hz	75Hz	90Hz
Lens Type	Aspheric	Aspheric	Fresnel
Lens Spacing	63.5mm	63.5mm	58–72mm
FOV (Horizontal)	90°	94°	88°
FOV (Vertical)	90°	99°	88°

5.1.3 Host Operating System

In the study described in Chapter 3, participants were observed to be exclusively using the Microsoft Windows family of operating systems, with Microsoft Windows 7 being the predominant version in use. To ensure that observed workflows could be replicated as precisely as possible, Microsoft Windows 7 was selected as the target operating system for development, utilising only features available on this platform when implementing the virtual environment.

5.1.4 Other Designs Considered

During preliminary research and prior to conducting the observation study described in Chapter 3, initial investigatory work on content capture was conducted using a GNU/Linux system as the host operating system in addition to the Microsoft Windows system described further in this chapter. This work considered two possible routes of capturing application windows on this GNU/Linux system, the first utilising X11 Composite Extension calls, and the second using X11 window forwarding over SSH to send drawing commands to a remote host. Both approaches were successful both in capturing application window data and displaying the captured data within a simple virtual environment, with the X11 Composite approach demonstrating performance that exceeded that approach described within this chapter.

While results of this preliminary investigation were promising, further development was halted following the observation study results analysis. This decision was made to allow for the possibility of pursuing investigations involving replication of observed workflows using the same applications seen. As described above, all observed users made use of Microsoft Windows systems, and the software used in their workflows was found to be unsuitable for execution or emulation on the tested GNU/Linux system. As this would result in the inability to replicate the observed workflows without substitution of applications which may have introduced confounding factors when performing comparisons, the decision was made to focus further development effort only on Microsoft Windows systems.

5.2 Content Capture Exploration

As capturing and displaying individual windows inside the virtual environment was of primary concern, this was chosen as the first development target. Several options for capturing the screen were given by Microsoft Corporation (2013), and several of these approaches were investigated for suitability.

With each investigated method, the first target was a single frame capture of a single foreground window. Once this was accomplished, a single frame of a single obscured background window was captured. After establishing the methods for capturing single frames, the investigation progressed to capturing a series of frames for a single foreground then background window as rapidly as possible over 10 seconds. Performance bottlenecks were identified for these captures, and where possible were addressed then the tests re-run. Finally, a combination of foreground and background windows were captured simultaneously with each capture happening in a dedicated thread, and any resulting bottlenecks were addressed if possible.

In this section, the primary candidates for capturing conventional application windows from a Microsoft Windows 7 platform are discussed. For approaches that yielded success, details on considerations made when utilising the method outlined. Finally, general optimisations are discussed.

5.2.1 BitBlt

As suggested by Microsoft Corporation (2013), the `GDI` function `BitBlt` was investigated as a possible candidate for capturing content. This function allows for sections of the visible screen to be captured as it would be shown on a traditional display, by specifying the position and dimensions of a section of the displayed screen to capture. This function does not allow for capturing arbitrary window content if the target window is obscured — attempts would result in a capture containing the obscuring windows overlaid on the target window as seen in Figure 5.1 — but does allow for foreground window content to be captured easily.

With this limitation in mind, an approach considered was to manipulate the position of the target window within the host system z-order, first bringing the target window topmost in the display order with no other windows obscuring it, capturing the window contents, then returning the window to its original position in the Z-stack. This approach showed promise when capturing individual windows, however when attempting to capture multiple windows simultaneously, it was observed that windows need be captured in series rather than parallel to avoid windows being moved to the foreground while another capture was in progress. This limited the speed at which captures could take place using this approach.

An alternative solution experimented with involved carefully scheduling captures so that large windows that could not share the foreground area were captured serially, while smaller windows that could be moved on screen such that they did not overlap with others and be captured in parallel (See Figure 5.2). While this addressed the issue in circumstances where many small windows were in use it still exhibited the worst-case performance of serial capture, and added further computation to determine appropriate tiling arrangements which proved inefficient.

While in the general case `BitBlt` proved generally reliable in its reproductions of windows and contents in captures, multimedia applications proved to be variable in quality. Some multimedia applications when captured showed window decoration and controls correctly but had blank content areas shown either in black or fully transparent when hardware acceleration was enabled, or when digital rights management (DRM) was involved in the playback. This issue is noted as a possible limitation in the function documentation, but it was found that affected applications could in many cases be configured to avoid the issue or could be replaced with alternatives that were unaffected.

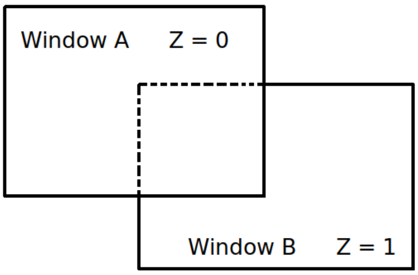

Figure 5.1: Visualisation of a foreground window obscuring a background window. The foreground window A obscures the content of window B in the area delimited by dashed lines.

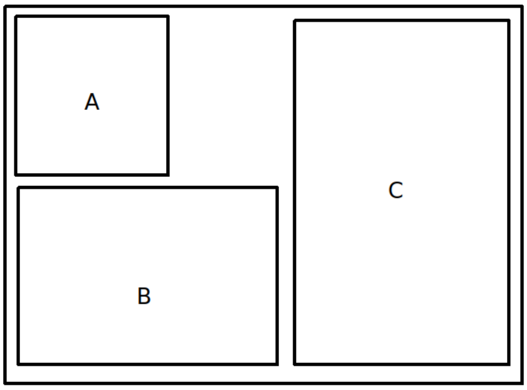

Figure 5.2: Visualisation of a window tiling arrangement that allows many windows to be presented unobscured within a capture area. Windows A, B, & C can each be captured simultaneously with `BitBlt` calls without obscuring each other.

Despite the discussed limitations, `BitBlt` operations were found to have an acceptable performance when capturing individual windows, allowing for a reliable 5 – 10 frame per second (FPS) of capture depending on window size and number. For parallel capture the performance was less reliable with greatly varying capture rates found due to the requirement of manipulating window z-order and position. This capture rate was found to be sufficient for viewing simple tasks where frame rate and responsiveness is not a concern such as document authoring and consumption, as well as for many precision tasks such as CAD where the lower than usual frame rate had limited impact. Visual multimedia applications did suffer as a result of this low frame rate, leading to mixed experiences for consumption.

5.2.2 PrintWindow

The `winuser` function `PrintWindow` was considered next as a window capture candidate. Unlike with `BitBlt`, this function allows for individual window contents to be captured on demand rather than capturing the whole, or a subsection of, the content displayed on the screen. This allows for obscured window contents to be captured, with any overlapping content excluded from the result. Captures taken through this method were found to be reasonable facsimiles of their target windows, but some notable differences could be observed in addition to the multimedia content issues discussed above.

Desktop compositing effects were noticeably absent in resulting window captures, with window transparency effects not shown, and rounded window corners showing black pixels in the negative space. While disappointing, these limitations could be avoided by disabling compositing effects within the operating system. This did not make these features appear correctly in captures, instead bringing parity to the captured content and the application as displayed on the conventional display.

It was found that for some applications which drew their own window decorations, or painted over the system decorations, that the system decorations were overlaid over the application content in resulting captures. This had the effect of obscuring controls for these applications and was most notably problematic with tabbed browsers that position tabs and page title text over where the system title-bar would normally lie, obscuring the tabs and associated controls in the capture. No suitable capture workaround was found that resolved this, however many of the applications that demonstrated this issue had options to disable custom window decoration, and to restore the native decorations which eliminated the issue, at the expense of displaying the window decorations that would normally be obscured.

Performance for `PrintWindow` calls was poor, but consistent across individual and simultaneous captures – averaging 1 – 3 FPS of capture for each window. This capture rate was considered acceptable for simple windows that did not update regularly such as text editors but had significant impact on applications that rely on regularly redrawing the screen.

5.2.3 DWM Thumbnail API

The DWM Thumbnail API was considered as a high-performance individual application window capture candidate, however this was not found to be suitable for the intended purpose following experimentation. This API backs the “window peeking” functionality of the Windows 7 taskbar, which allows for small live previews of windows when hovering over their taskbar entries. These thumbnails can be seen updating in real-time, with minimal observable performance drop when compared with the application window they represent.

The DWM Thumbnail API allows for requesting the content of a window be rendered scaled into an alternative device context, allowing for window contents to be rendered inside another window, but does not allow for the contents to be kept in memory for later use. This resulted in a prototype that could rapidly capture the contents of another application window and display it on another surface but would then require this thumbnail window to be re-captured by an alternative method rather than being able to directly utilise it. This is noted as a deliberate design choice by Microsoft, with the documentation overview stating that “DWM thumbnails do not enable developers to create applications like the Windows Vista Flip3D (WINKEY-TAB) feature” (Microsoft Corporation, 2018b). Examples of using this API to render into a memory context were found, but all relied on replacing or hooking versions of system libraries that could not be replicated in the test environment.

5.2.4 DXGI Desktop Duplication API

The DXGI Desktop Duplication API was noted as a promising candidate for capture of the screen contents, however this was not investigated as a part of this system design. This API was not available on the Windows 7 system that was available to the researcher and as such was not a viable path to pursue. It has been noted that for applications mirroring an entire display as presented on a traditional display, this API offers one of the higher performance capture options available on the Microsoft Windows platform, and future environment investigation using newer operating system versions should pursue this as an alternative to `BitBlt`. It is further noted that modern virtual reality (VR) desktop environments such as Desktop+, developed more recently than the environment described in this chapter, use this method to achieve high performance display capture.

5.2.5 General Optimisations

During experimentation it became apparent that `BitBlt` and `PrintWindow` calls were computationally expensive, with large amounts of data being copied in memory for each call. Small gains were made by reusing allocated memory for the destination of these calls which reduced some overhead, but the underlying expense of each call was still present.

Attempts were made to detect changes in window content so that captures could be updated on change rather than at regular intervals — freeing up resources for capturing more regularly updating windows — but no simple or reliable solution could be found. Intercepting window messages destined for the target window and looking for `WM\_PAINT` messages initially showed promise, but many windows were found to be updating without these, or any other messages being generated, so updates were missed.

Ultimately it was determined that further optimisations were not going to be easily made, and that brute force capturing windows as quickly as possible would be the simplest initial approach.

5.3 VR Environment Exploration

With window capturing approaches investigated, transferring the captured window data into a virtual environment became the next concern. A test-bed platform was developed on the Microsoft Windows 7 platform, making use of the DK1 IVR HMD. This platform was developed in C#, using the SharpOVR Oculus library, and the SharpDX DirectX9 library. The platform was developed using a right-hand coordinate system, with the y axis controlling elevation. The virtual camera was positioned at the origin, facing $\left[001\right]$. The basic platform consisted of a simple cornflower blue environment, with a single textured mesh having its centre point positioned at $\left[001\right]$, facing toward the origin.

Captured application windows were converted into Direct3D compatible formats, then used as textures which were applied to the mesh. For initial exploration, window data captured using `PrintWindow` from a single application window was utilised, with the target window being identified using WinSpy++. Updates to the captured texture were run once per second, displaying an updating texture that allowed for proving the concept, without introducing possible overhead related to the capture process to the virtual environment development.

The geometry of the tested mesh was one of the first considerations when testing this platform, with several options considered, and two evaluated in depth. These are discussed in this section.

5.3.1 Planar Mesh

The first mesh utilised was a simple planar mesh, consisting of 4 vertices positioned equidistant from each other. Captured textures were applied to this mesh, and necessary transforms calculated to have the orientation and colour profiles of the displayed texture match that of the underlying real window from which the texture was sourced.

After establishing that captured textures could be applied to this mesh, the first modification made was to have the aspect ratio of the planar mesh match that of the underlying real window. To accomplish this, the capture process for real windows was modified to first capture the window dimensions before capturing the contents. After conversion of the captured texture into an appropriate form, the mesh onto which the texture was to be applied was recreated with dimensions calculated from the real window dimensions (See Section 5.3.4).

Legibility of text was immediately noted as a concern within the environment when using this mesh. When positioned at $\begin{bmatrix}001\end{bmatrix}$, text was easily readable on narrow windows, but as the mesh width increased to accommodate wider underlying window content, text towards the edges of the mesh became significantly harder to read. The reasoning for this is obvious then considering this configuration as a right-angle triangle formed from the perceiver position at the origin, to the centre of the mesh and an outer edge (See Figure 5.3). If the distance between the perceiver and the mesh centre remains constant as the mesh width increases, then the length of the triangle hypotenuse increases while the incidence angle θ decreases. This makes the edge of the mesh appear further away from the perceiver and feels like reading from an angled surface. This observation is made by both Larson et al. (2000) and Buttner et al. (2020), who identified that for text rotated at more than 60° with respect to the perceiver, that text became difficult to read, requiring an increased font size to become readable. To reduce this effect, the distance at which the mesh is positioned from the perceiver can be increased, however this results in a decrease in clarity of text at a distance.

While readability concerns of this mesh geometry at larger sizes are noted, for meshes with smaller dimensions the planar mesh supports reading tasks well. Informal observations made by several users experiencing the system was that even at larger sizes they felt that this mesh geometry was what

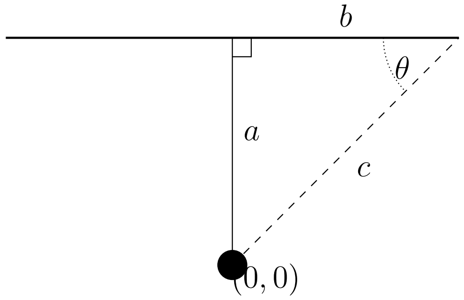

Figure 5.3: A visualisation of the issues observed with planar meshes. Length a is the mesh distance from the perceiver, length b is half the mesh width, and length c is the apparent distance of the mesh edge. θ shows the incidence angle of perceiver gaze on the mesh edge, which can be contrasted with the 90° angle of incidence for the mesh centre.

they desired for windows within a virtual environment. This opinion was not shared by all users however, with others suggesting that this mesh geometry became "unusable" after the width exceeded a critical point. When displaying other content such as visual multimedia, sentiment towards this mesh geometry was more universally positive with users experiencing them, with comments suggesting the experience was like that of a cinema.

5.3.2 Curved Distortion Mesh

To address observed issues with the planar mesh, an alternative curved mesh geometry — hereafter referred to as "curved distortion" mesh — was developed. This new mesh took the form of a four-sided section of a "display sphere" — with a radius equal to the z-distance measure used with planar meshes — with the top and bottom edges falling along latitudinal parallels equidistant from the equator, and the left and right edges falling along longitudinal meridians. A visualisation of this can be seen in Figure 5.4.

By curving around the surface of a display sphere, the distance between any given point on mesh and the perceiver will be $\approx r$, and the incidence angle of the user gaze will be $\approx \frac{\pi}{2}$ — assuming sufficient intermediate subdivisions of the mesh following meridian and parallel lines. This addresses the primary issues observed with planar meshes — that for wide or tall meshes, the distance between the perceiver and the mesh surface increases and the angle of incidence with the mesh becoming increasingly acute as focus moves towards the mesh edge.

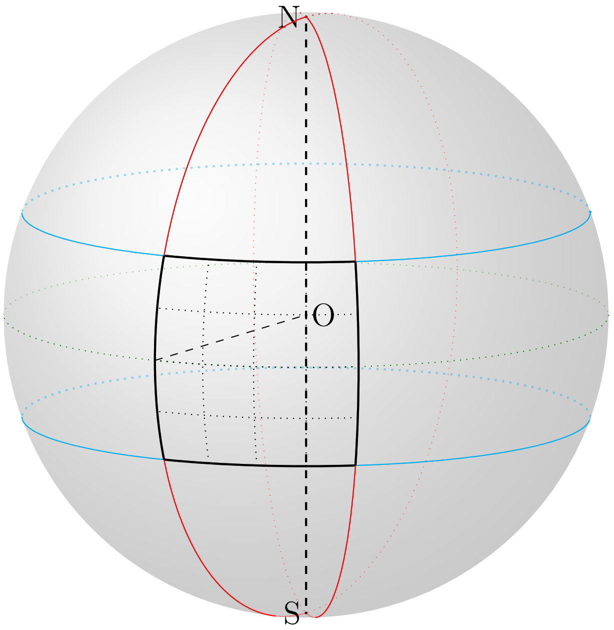

Figure 5.4: A visualisation of the curved distortion mesh geometry on the surface of a sphere. longitudinal meridians are marked in red, latitudinal parallels in cyan, and the equator parallel in green. Mesh edges are indicated with solid black edges, with intermediate meridians and parallels shown in dotted black. Note that in the worst-case vertical dimension mesh with a subtended angle of π, the left and right edges of the mesh will continue along the meridian lines to the poles giving the top and bottom edges a length of 0.

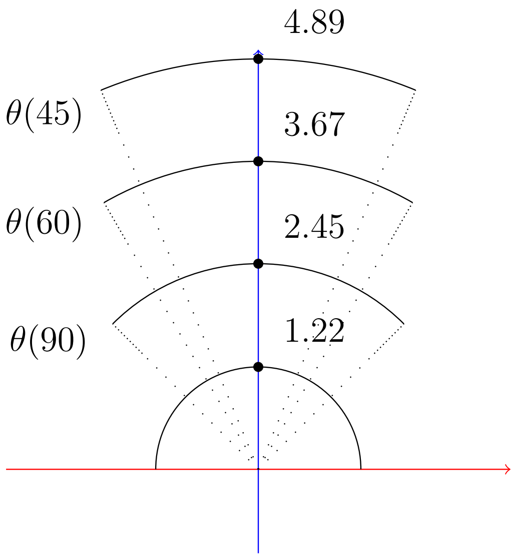

Figure 5.5: 2-dimensional visualisation of curved distortion mesh arcs calculated from display spheres of varying radii. Arcs displayed have a constant arc length of 3.84 in-engine units (IEUs), equating to a texture with dimension of 1920 pixels across the measured axis (See Section 5.3.4).

Width and height of meshes displayed in this fashion are presented as circle arc lengths, measured across the centre meridian and parallel lines the mesh is based around. These arc lengths can be expressed as angular arc measures for ease of manipulation, which is calculated through $L = r\theta$, where L is the arc length, r is the radius of the circle, and θ is the angular arc measure in radians. The radius of this calculation is equivalent to the z-distance used for positioning planar meshes.

Increasing the radius of the display sphere decreases the angular arc measure required to display a given mesh, resulting in reduced apparent curvature of meshes of equivalent dimensions, at the cost of displaying the mesh at a greater distance from the perceiver. This can be seen in Figure 5.5.

While this mesh geometry addresses noted issues with the planar mesh, it is not without issues. Due to the nature of the left and right edges following longitudinal meridians, the arc length along any parallels across the mesh decreases as the azimuth angle increases. This leads to a pinching effect on the top and bottom edges of the mesh as the mesh height increases, with the worst case of $\theta = \pm\frac{\pi}{2}$ resulting in the top and bottom points converging to the poles of the sphere, forming a mesh in the form of a spherical lune. The arc length of the top or bottom mesh edges compared to the arc length of the

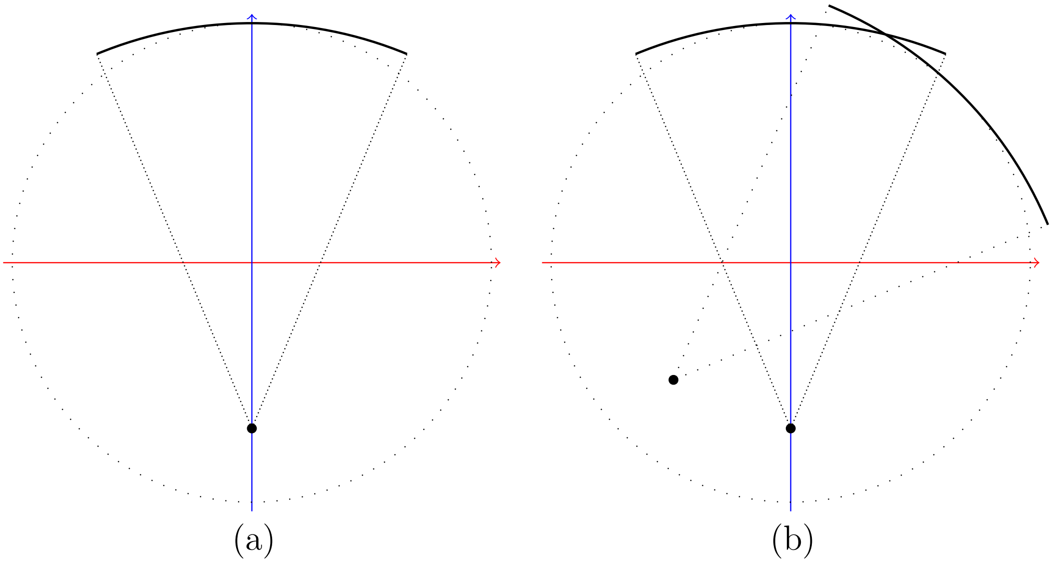
(a) (b)

Figure 5.6: 2-dimensional visualisations of virtual windows curved around a point behind the perceiver; *a* single virtual window, *b* two virtual windows demonstrating intersection when rotating about environment origin.

equatorial parallel for the mesh can be expressed by $\cos\left(\frac{\theta}{2}\right)$, where θ is the azimuth angle in radians above or below the y plane, in range $\left[0, \frac{pi}{2}\right]$. This pinching effect can be mitigated by increasing the radius of the display sphere as described above.

An additional concern noted with this geometry is the extent to which a user needs to move their head when observing points on larger meshes. Informal feedback by users of this system indicated that while content was clearer on meshes with this form, that the head movements were tiring and gave a feeling of motion sickness for some users. Kim and Shin (2021) assessed physical discomfort and simulator sickness symptoms when performing common office tasks involving reading within a virtual environment, noting that study participants developed feelings of muscular discomfort and simulator sickness twice as quickly within the virtual environment when compared to a conventional display. This is in part attributed to the increased head and neck movements required of the participants when using the tested IVR HMD, but the weight of the HMD producing flexion movements within the cervical spine was also noted as a contributing factor.

Outside of physical discomfort, it was also noted by some that these meshes with larger dimensions felt “unnatural” and “oppressive” as they began to enclose the user. For some this feeling was quite pronounced, with a desire to push the mesh away to escape the feeling expressed by one user. Others simply noted that they would prefer the mesh to be less enclosing but did not express strong desires to leave the environment. Guidance given by McKenzie et al.

(2017) suggests curving content around a point behind the perceiver, giving a greater curvature radius (See Figure 5.6), may reduce this effect. When trialed, users indicated that the effect did reduce the oppressive feeling but reported that this approach to curvature harmed the readability of text on the outer edges of virtual windows. This approach to correcting this effect also presented concerns around virtual window layout, as neighbouring virtual window meshes would intersect each-other if rotated around the environment origin rather than the curvature origin. Due to these concerns, adjusting the curvature origin was not pursued.

5.3.3 Other Meshes

Several additional mesh geometries were evaluated during this process, however when evaluated initially by users of the system these were discarded for further testing due to similarities with the above described meshes, or due to unsuitability. Of note, a cylindrical mesh was created, that wrapped around the perceiver horizontally but not vertically. This form of curvature is seen in many modern general purpose virtual environments such as DesktopPlus (2021) and Virtual Desktop, Inc. (2021), but was discarded for this environment following user feedback suggesting that text at the top and bottom of virtual windows was difficult to read. Some users expressed an interest in this geometry, as it exhibited some of the properties of the planar mesh that they preferred, however when asked to identify the best geometries for different tasks, the planar and curved distortion meshes were always chosen in preference to the cylindrical mesh.

5.3.4 Mesh Dimensions

Dimensions of virtual windows were calculated from the size of the texture to be presented upon them, using the equation:

$$d_v = \frac{d_r}{500}$$

where d_r is a texture dimension in pixels, and d_v is a dimension in IEUs. The scaling factor of 500 is arbitrary but was chosen as this was a conversion factor used in example projects shipped as Oculus development examples.

As a relation between pixels and IEU is established, and pixels have a physical dimension when displayed on a conventional display of known pixel density, virtual windows can be described in size and distance from the perceiver using real world dimensions. Take for example a virtual window from the visualisations in Figure 5.5, a virtual window 3.84 IEU wide positioned at

4.89 IEU from the perceiver. Converting these virtual dimensions into pixels yields a window 1920 pixels wide, positioned 2445 pixels from the perceiver. If relating these measurements to an image displayed on a 38 pixel per centimetre (ppcm) density display, these can be equated to viewing a 50.5 cm wide image on a 38 ppcm display from a distance of 64.3 cm.

5.4 Generation 1

To create a more complete system, the above-described prototype was extended and modified to support multiple meshes, referred to henceforth as virtual windows, each displaying the contents of a different conventional application window. Virtual windows could be created or destroyed in response to new applications opening or closing on the host Microsoft Windows 7 environment and were made responsive to window management interactions such as resizing, maximising, and minimising.

Mouse support was added to allow for a virtual mouse cursor that could interact with the window content, and window decorations to support resizing and moving. Keyboard input support was added, supporting hotkey interactions with the environment, and interaction with the application windows.

Interaction with underlying application windows was accomplished through intercepting device input and manipulating the ordering of real windows within the Microsoft Windows environment in response to observed input. By modifying the window ordering, the environment could control which real window would receive input from the mouse or keyboard, with the topmost window receiving the bulk of the input. Intercepting the input allowed for recreation of features like the mouse cursor within the virtual environment and allowed the opportunity to modify the positioning of the real cursor to reflect the virtual positioning. This input proxying approach allowed for a simple, but effective, solution which allowed for familiar interactions and behaviours with applications.

During initial development this system targeted only the DK1 device, but support for the DK2 was added at the time of its release. Due to the DK2 featuring 6 degrees of freedom (6DoF) tracking with the DK1 only offering 3 degrees of freedom (3DoF), controls were added to support moving the camera position using keyboard shortcuts when the DK1 was in use.

```
while running
  enumerated_windows := EnumWindows()

  for hwnd in enumerated_windows
    if hwnd not in tracked_windows
      if allowable_window(hwnd)
        new_v_window := VirtualWindow(hwnd)
        tracked_windows.insert_first(new_v_window)

  for v_window in tracked_windows
    if v_window.hwnd not in enumerated_windows
      tracked_windows.remove(v_window)
```

Listing 5.1: Pseudocode showing the process for updating the tracked window list within the virtual environment.

5.4.1 Multi-window support

To support detection of application windows being created, deleted, and modified, system calls that allowed access to windows using their corresponding window handles (HWND) were used.

Windows displayed in the host environment were enumerated using the win32 function `EnumWindows`, which returned a list of HWNDs for the active system. Many of these handles did not represent visible application windows that the user would typically interact with, and that therefore should not be displayed in the virtual environment. Some of these could be ignored by inspecting window properties, and ignoring applications that had, or lacked, certain properties and value. Other HWNDs were identified as non-interactable through matching of predefined property values, such as the application window title, or process name to which the HWND was parented. HWNDs that were determined to be ignorable were maintained as a list, which candidate HWNDs were first checked against before being marked as a visible and interactable window. Valid HWNDs determined to be of interest were added to a tracked window collection and were verified on subsequent updates to ensure they were still displayed in the host system. If HWNDs in this collection were not found to be present in the host system, then they were removed from the tracked window collection. On each update, the virtual environment updated the tracked window collection, collecting information about each HWND which was maintained for future usage. The position and size of the application window represented by each HWND was recorded, and when changed triggered a reactive behaviour within the virtual environment.

Within the virtual environment, a virtual window was created for each tracked window and associated with the corresponding HWND. The virtual window structure maintained information collected and translated from the underlying application window that it represented within the environment, with dimensions translated into IEUs understood by the virtual environment. Positional information for the centre of each virtual window was maintained simply as azimuth and inclination rotations about the environment origin, that were stored within a rotation matrix. No rotation about the virtual window centre was supported, so as to ensure that virtual windows always appeared upright within the environment, with their top and bottom centre-points aligning to the same meridian line. This was decided during initial testing based on user feedback, which indicated that rotation about a window centre-point was confusing to manage and offered little benefit.

After being associated with a virtual window, tracked application windows were re-positioned within the host environment such that they were displayed on the primary conventional display, positioned at the display origin. They were additionally restored to a windowed state if they were previously minimised or maximised. This action allowed for updates of application window properties to detect if changes had occurred to the window that should then be reflected in the virtual environment.

5.4.2 Mouse interaction

To support mouse interactions, a fixed size, named curved distortion virtual window, referred to henceforth as the virtual cursor, was created within the virtual environment to represent the mouse position and cursor. This was placed closer to the camera than the topmost virtual window, ensuring that it would never be obscured. The host system was queried on every update to identify the currently displayed mouse cursor, and this was displayed on the virtual cursor. Cursors were only captured once, converted to a texture with correct offsets on the first usage, with subsequent occurrences of a cursor reusing an existing texture. During creation, transparent cursor textures were populated using the cursor bitmap obtained from the host environment, with the bitmap data offset such that the specified “HotSpot” of each cursor would be centred in the texture.

Shared access to the mouse input was acquired from the DirectInput subsystem, and X-Y movements of the mouse mapped to scaled azimuth and inclination rotations of the virtual cursor. Inclination rotations were clamped within limits in the virtual environment, to prevent the virtual cursor from flipping over the top of the environment and becoming inverted. Azimuth rotations were unbounded, allowing for rotations to “wrap around” to the starting point as they passed through the environment. As with other virtual windows, no rotation about the virtual cursor centre was permitted, ensuring that the virtual cursor always remained ‘upright’ from the user perspective.

```
function intersection(v_window, v_mouse)
  % traverse triangles in left-to-right, top-to-bottom order
  for (triangle, index) in v_window.mesh
    intersect_pt := intersect(triangle, v_mouse.rotation)
    if intersect_pt
      tess_w := get_width(v_window.hwnd) / v_window.tess
      tess_h := get_height(v_window.hwnd) / v_window.tess

      max_x := max(triangle.va.x, triangle.vb.x, triangle.vc.x)
      max_y := max(triangle.va.y, triangle.vb.y, triangle.vc.y)
      min_x := min(triangle.va.x, triangle.vb.x, triangle.vc.x)
      min_y := min(triangle.va.y, triangle.vb.y, triangle.vc.y)

      tess_x := (index mod v_window.tess) / 2
      tess_y := index / (v_window.tess * 2)
      pct_x := (intersect_pt.x - min_x) / (max_x - min_x)
      pct_y := (intersect_pt.y - min_y) / (max_y - min_y)

      real_x := (tess_x + pct_x) * tess_w
      real_y := (tess_y + pct_y) * tess_h

      return Point(to_int(real_x), to_int(real_y))
  return nil
```

Listing 5.2: Pseudocode showing the translation of a intersection between a virtual window and mouse from virtual environment rotational relationships, to Cartesian coordinates relative to the top-left of the real window

On each frame update in which the virtual cursor, or any other virtual window within the environment had their position updated, a Möller–Trumbore ray-triangle intersection calculation (See Listing 5.2) was used to determine if the virtual cursor was over the topmost virtual window within the virtual environment. If it was, then the underlying application window represented by the virtual window was made topmost in the host environment z-order, then the host environment mouse cursor was moved to a point over the application window that corresponded to the intersection point of the virtual cursor and the virtual window. This movement of the host cursor triggered appropriate cursor glyph updates, and behaviours in the underlying application window which were then captured by the virtual environment. Mouse click events were observed, but otherwise allowed to pass directly to the topmost window in the host environment. If the virtual cursor was not over the topmost virtual window, then the underlying host system cursor location was set to a known location away from any underlying application windows, where click events would not cause damage.

```
% Proxy behaviour
while running
for (v_window, index) in tracked_windows
  if index is 0
    v_window.texture := BitBlt(v_window.hwnd)
    window_intersect := intersection(v_window, v_mouse)
    if window_intersect
      set_window_order(v_window.hwnd, 0)
      set_mouse_position(window_intersect)

      if get_mouse_glyph() not v_mouse.glyph
        if get_mouse_glyph() not in v_mouse.glyphs
          v_mouse.glyphs.add(load_mouse_glyph())
        v_mouse.glyph := v_mouse.glyphs[get_mouse_glyph()]
    else
      set_window_order(environment_window, 0)
      set_mouse_position(0, 0)
  else-if rr_counter is tracked_windows.index(v_window)
    v_window.texture := PrintWindow(v_window.hwnd)

  rr_counter := (rr_counter + 1) mod length(tracked_windows)
```

Listing 5.3: Pseudocode showing virtual windows being updated, and proxy behaviour for input.

If a mouse click event occurred when the virtual cursor was not over the topmost virtual window, then each virtual window was sequentially tested for intersections. If an intersection was detected, this triggered a change in virtual window ordering, pulling the intersected virtual window forward to be the topmost virtual window, and triggering a re-ordering of other virtual windows.

5.4.3 Keyboard interaction

Keyboard interaction was accomplished by requesting shared access to the device though DirectInput and listening for environment specific hotkeys. Input was otherwise allowed to pass directly to the topmost real window without modification. As the receiving window was controlled through manipulation of window ordering in the host system, keystrokes that should not be received by an underlying window would be captured by the virtual environment window that was brought topmost in the host system ordering.

5.4.4 Virtual Window Management

Window management features were added to the virtual environment to support common behaviours seen in the host environment. These features were implemented by allowing input to pass through to underlying application window controls and observing the underlying application window properties.

```
while running
  for (v_window, index) in tracked_windows
    w_pos := window_position(v_window.hwnd)
    if w_pos not v_window.real_pos
      % Calculate virtual rotation from relative pixel
      % difference
      q_x := (w_pos.x - v_window.real_pos.x) * dx_quat
      q_y := (w_pos.y - v_window.real_pos.y) * dy_quat
      rotate(v_window, q_x, q_y)
      set_position(v_window.hwnd, v_window.real_pos)

    if window_size(v_window.hwnd) not v_window.real_size
      % Update and recreate mesh to reflect altered size and
      % aspect ratio
      v_window.real_size := v_window.hwnd.size
      v_window.mesh := create_panel(v_window.real_size, index)

    if window_title(v_window.hwnd) not v_window.real_title
      v_window.real_title := window_title(v_window.hwnd)
```

Listing 5.4: Pseudocode showing the process for updating window attributes, and responses to changes.

Movement of virtual windows was accomplished by identifying when the underlying application window moved from a set location. The observed relative changes in position would be caused by the user clicking and dragging on the window title bar as they would do in the host environment. These relative movements were recorded, then translated into scaled azimuth and inclination rotations about the origin of the virtual environment, which were applied to the virtual window. After recording the change in position, the underlying application window was relocated back to the set location, allowing for accurate capture and future movement detection. Application window captures were observed to contain artefacts when the application window being captured was moving. These were due to the underlying application window changing position while the application window capture operation was in progress, resulting in a `BitBlt` operation at incorrect coordinates. To avoid this, application window capture updates were suspended while the virtual window was in moving mode, with a red border applied to the texture to highlight the freeze.

Resizing of virtual windows was accomplished by again observing the underlying application window size property, and when a change was noted the virtual window mesh would be updated as appropriate. Users could interact with the window decorations and resizing handle visible upon the virtual window to change the dimensions of the window, and the virtual environment would reflect these changes. Some difficulty was observed in interacting with the window resizing controls due to their small size and pointing inaccuracies within the environment. To address this, virtual window textures were made

```
% Update window z-ordering
while running
  if left_button in get_mouse_keys()
    for (v_window, index) in tracked_windows
      if v_mouse.rotation intersects v_window
        swap_index(tracked_windows, 0, index)
```

Listing 5.5: Pseudocode showing the process for updating the topmost (active) virtual window.

a few pixels larger than necessary prior to capturing application window content, leaving transparent space on the right and bottom of the virtual window. This effectively made the virtual window slightly larger than the underlying application window, and when translating the virtual cursor location into host environment coordinates, the resulting point was clamped to the application window dimensions. This made targeting the resizing controls easier by reducing the pointing error within the environment.

Minimising of virtual windows was also supported by allowing the user to interact with the underlying application window controls and observing the properties of the window. If an application window was determined to have been iconified, it would be immediately restored within the host environment making it visible once more. Within the virtual environment, then the virtual window mesh and texture were replaced with a plain coloured rectangular window containing only the application window title text, and the virtual window was marked as iconified. Iconified virtual windows could be moved within the virtual environment by clicking and dragging them to desired locations. Double clicking an iconified virtual window would remove the iconified state and trigger an update for the virtual window — reacquiring the application window dimensions before recreating a mesh, then capturing the window contents for display.

5.4.5 Hotkeys

To support usage of this virtual environment, support for hotkey combinations were added. Several hotkeys were defined and bound to functions that addressed concerns identified about the virtual environment.

The first hotkey triggered a re-centring action for the HMD, causing the current rotational orientation of the user and HMD to be recorded as “forward” withing the environment. This was found to be necessary when using the DK1, as over time the environment would appear to have rotated with respect to the users seated position due to the gyroscopic drift exhibit by the rotational sensors. This was not necessary when using the DK2, as the positional tracking camera was able to correct this drift automatically, but it did enable a form of complete environmental rotation that a minority of users found useful.

A second hotkey allowed for modification of the mouse scroll-wheel behaviour. When pressed, scroll events generated in response to scroll-wheel movements were intercepted, and converted into positive and negative values depending on scroll direction. The value captured was then used to modify the virtual environment radius, moving all virtual windows closer to, or further away from the perceiver. Initial implementation of this feature allowed for either a single virtual window, or all virtual windows to be moved in response to the scrolling, determined by which virtual window (if any) was determined to be active. This behaviour was removed following feedback from users, who identified several issues with this behaviour. The first issue reported was that it was too easy to accidentally move a single virtual window when moving all windows was what was expected, and the second that once a single virtual window had been moved, it was difficult to return it to its original location. General feedback suggested that in almost every circumstance, moving all virtual windows was desired, so the hotkey behaviour was restricted to this one operation.

A third hotkey was defined to assist with locating the virtual cursor within the virtual environment. Several users observed that when working in the virtual environment, that the virtual cursor could end up behind them or out into their peripheral, and that locating it introduced task delays and was frustrating. When tapped, this hotkey would cause the virtual cursor to rotate to align with the user gaze over 0.25 seconds, with an accompanying movement animation to help identify the movement. This feature allowed for the virtual cursor to be rapidly re-centred, bringing it to bear on the user's task. Alternative approaches were considered to manage virtual cursor location, including having the virtual cursor slowly move into the user field of view (FOV) if it was left in a location that could not be seen, and having a visual indicator in the user's peripheral vision pointing towards the virtual cursor. Feedback given on these proposals indicated that both were less desirable than a direct solution.

5.4.6 Test Environment

In order to conduct user studies evaluating suitability of this virtual environment, a copy of the environment was made, and modified to disable most of the window management functionality, leaving it capable of simply displaying static textures upon specified virtual windows. This test environment allowed for textures to be displayed according to a provided order, upon virtual windows of specified geometries. Support was added to allow for changing virtual window positions using the mouse scroll wheel with no additional hotkey keybindings, and a hotkey was added to display the next texture and virtual window combination in the supplied order. Initial position of virtual windows could be specified along with the provided sequence, and timing, virtual window geometry, displayed texture, and chosen position information could be logged for later analysis.

This modified environment formed the testing base for the study described in Chapter 6.

5.5 Generation 2

Following the user study conducted using the generation 1 virtual environment, a second generation of virtual environment was created. While the core functionality of the system remained much the same, several pain points identified during development and testing were addressed, to make future development and user studies easier.

The virtual environment was rewritten in C++ using the native Win32, DirectX11, and Oculus LibOVR frameworks. This rewrite was made necessary by the abandoning of the SharpOVR C# library providing support for the Oculus Rift devices, limiting access to new devices and features — most notably the newly released CV1 HMD, and the asynchronous timewarp framework feature that dramatically improved the perceived experience within the Oculus Devices. Potential performance benefits were identified in moving from DirectX9 to DirectX11, specifically around texture loading performance which allowed for faster texture updating.

Beyond simple rewriting and performance improvements, the generation 2 virtual environment also added several features to those found in the generation 1 implementation. This section will discuss several of the major additions.

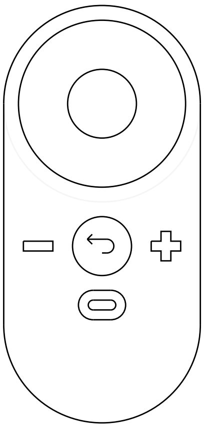

Figure 5.7: Oculus Rift remote device.

5.5.1 Positional Tracking

Positional tracking support was added for the CV1 and adjusted for the DK2 to allow for the user to move within the virtual environment. This support allowed the perceiver to move their position closer to virtual windows, bringing smaller or unclear features into view, or to move further away to gain an overview of several windows. This improved the sense of user presence within the virtual environment, as this movement reportedly felt natural to users, supporting a behaviour that they might use when working with conventional displays.

To support further testing, this option was made toggleable both at the launch of the environment, and at run-time. When disabled at launch, the environment origin was used as the pose position, and when disabled at run-time, the current position of the pose was frozen with all rotations occurring around this point. A further hotkey was bound to reset the pose position to the environment origin.

5.5.2 Input Methods

A gaze pointer was enabled within the environment, allowing the user's heat orientation to support pointing operations. This method of pointing was identified as both deliberate and accurate Kytö et al. (2018) for these operations, and as faster than use of a mouse or other pointer Jalaliniya et al. (2014). As mouse was still the primary pointing input mechanism, the existing infrastructure was utilised without major modification. When gaze pointer was in use, head movements were bound to the virtual cursor position, changing the position of this as the head was moved. When over a virtual window, this cursor glyph was shown, and when not over a window a configurable glyph was displayed instead.

While the implemented gaze pointer was able to move a mouse cursor, it was not able to replace the mouse completely as it lacked button and scroll-wheel support. To provide those features, support for the Oculus Remote device (See Figure 5.7) was introduced, with the upper directional ring controlling vertical and horizontal scrolling when pressed, and the centre button contained within acting as a left-click when clicked, and right-click when held.

5.5.3 Run-time Configuration

To increase testing flexibility, all features within the system were made toggle-able during both startup and run-time. Hotkeys were created to allow these to be toggled while using the environment, which simplified the debugging and sample scenario creation greatly.

A remote-control protocol was created, allowing an external application to launch the environment and control its behaviour. Using this protocol, controlling applications could launch the environment with features enabled, and run-time options could be toggled as desired. Virtual applications could be displayed at will within the environment, capturing window content for specified HWNDs once or repeatedly, or displaying fixed textures specified as a file URI or byte stream. Virtual windows can be specified as interactable, and keypress or mouse interactions would be reported, identifying the interacted window, along with where interaction occurred and what sort of interaction was observed.

5.5.4 Testing Harness

To support further user studies, an external control application, referred to as the "testing harness", was developed that allowed experimental plans to be executed, performing automated configuration and actions within a running generation 2 virtual environment instance. Test plans were defined by implementing an interface within the testing harness, containing step-by-step actions to be performed within the virtual environment. All features of the virtual environment were mapped to actions, allowing features of the environment to be toggled at run-time, and virtual windows to be created, resized, moved, and destroyed at will. Actions could request a response from the virtual environment, identifying state of the environment, or details of virtual windows that were interacted with.

Pre- and post- test actions that could be performed by the testing harness, included generation of participant IDs and collection of specified demographic data, along with corresponding logging of this information in an optionally anonymised way. Data collected for a participant was automatically collected into a folder structure, containing log files produced by the testing harness and the virtual environment, along with run-time variable information that could be used to repeat the experiment as executed.

Step-by-step actions were executed by the testing harness, with pauses for responses from the connected virtual environment.

5.6 Discussion

While developing the virtual environment discussed in this chapter, observations around the virtual environment behaviour and interactions, as well as general communication concerns were made by the researcher and environment testers. In this section, a selection of these observations are discussed.

5.6.1 Virtual Window Geometries

While initial versions of the environment utilised planar meshes as the primary virtual window geometry, it was identified quickly by the researcher and initial users that these were not suitable for displaying larger windows. An investigation into alternative geometries was conducted, with several candidate geometries implemented. Following feedback from users, the curved distortion that is found in the final implementations was identified as the most popular and usable option. This option allowed content to be seen at any point on the virtual window, at the expense of increased head movement amplitudes.

During evaluation however, a subset of users voiced a strong preference for the planar geometry virtual windows despite their size limitations that the curved geometry sought to address. The reason for the preference was mixed from these users, with several unable to identify the precise reason for their preference. Those that were able to identify a reason stated they felt that distorting virtual windows gave them an “unnatural” feeling, transforming the familiar planar windows into something unusual that they found difficult to engage with. They suggested that tasks such as reading, and comprehension of read content were negatively affected when consuming content from curved virtual windows, and that they believed their performance was better on planar windows even at larger sizes.

5.6.2 Input

Keyboard and mouse were chosen as the primary input methods due to the reasons described in Section 5.1, and the experience of using these in an IVR environment yielded mixed results. As would be expected when visually isolated, users had trouble locating these devices and positioning their hands correctly on them. As the environment allowed the user to orient themselves in any desired direction, placing the devices in consistent positions was insufficient to overcome the issue. Initial trials of placing visual references on the environment skybox to help locate the desk surface were promising, with users able to consistently orient themselves towards the desk surface. This approach did not assist with identifying devices on the desk however, meaning it was of questionable use, and it further suffered due to virtual windows obscuring the skybox when the environment was in use. An alternative approach that did help was placing a surface on the user's legs on which a wireless mouse and keyboard could be placed. This allowed users to consistently locate the devices regardless of physical orientation, however positioning of hands on the devices was still noted as an issue.

Once a user had located the keyboard, they were able to interact with it to varying degrees based upon their experience and practice with typing. User with touch-typing experience did not struggle with typing, or with toggling hotkeys, but did indicate that their inability to glance down at the keyboard was disorienting. Users that relied on hunt-and-peck typing, or who only had limited touch-typing experience struggled more, and reported that the experience was unpleasant. Gaps in the HMDs faceplates around the nose allowed some users to see out of the HMD to the keyboard, which allowed them to function to a limited extent within the environment. This approach was noted as being quite uncomfortable due to neck and eye strain, as the orientation of these gaps required them to look down while raising their head to view their surface. Use of cross reality (CR) techniques to visually identify hand and keyboard positions, as seen with commercial offerings such as Immersed Inc. (2021), would help to address this issue.

Mouse interaction proved to be remarkably well received, with the general feedback indicating that if additional features were added to help physically locate the mouse, no other modifications would be needed to use this inside virtual environments of this configuration. With the design as implemented, linear movements of the mouse translated directly to rotational movements about the user, limited to two dimensions as is expected. When used with curved virtual windows, the movements of the mouse cursor across the windows with respect to physical mouse movements appeared the same as on conventional displays and windows. Interaction with planar windows was identified as problematic, as parallax error was introduced between the cursor glyph and window content, leading to inaccuracies with pointing tasks on larger virtual windows.

5.6.3 Mouse Clamping

Early in development, it was identified that allowing the virtual cursor to rotate to the poles of the display sphere introduced issues with control inversion and general inaccuracies around those regions. Several different approaches were trialed to address the issues, however all had shortcomings of their own, and most were abandoned. The solution that was settled on was to limit the inclination of the virtual cursor to a fixed range, introducing "dead zones" to the poles where the cursor could not enter, to prevent the problematic conditions from occurring. 22.5° from each pole was chosen as the inclination cut-off by the researcher, as it was observed that this marked the approximate point at which issues began to manifest.

Reception of this limitation was mixed by users, with several indicating that they wanted to be able to position virtual windows at the poles, which this limitation largely prevented as the virtual cursor was unable to drag windows into these areas completely. Positioning windows below the user was found to be possible, as the top title bar which controlled window position could still be positioned by the virtual cursor to this point, leaving content below the title bar inside the dead zone. Positioning content above the user was not possible with this approach. Providing a method to more easily position virtual windows to these extreme positions whilst maintaining the mouse clamping behaviour may be desirable for future iterations of this virtual environment.

5.6.4 Window Capture Performance

Application windows were captured using `PrintWindow` and `BitBlt` methods as described above. While we were able to capture content at a fair rate when in the foreground, feedback from users indicated that the experience was poor for content that was shown in the background. Users expressed that one of their reasons for having multiple windows visible at once would be to view background content including video. If the performance was not sufficient for this use-case, then they would have limited use for the system.

This highlights the importance of platform support for such an environment. If we were to replace conventional computing environments with this implementation, we would need mechanisms to support real-time capture of application windows, or support for redirection to off-screen surfaces. Using newer APIs from Microsoft such as those provided by `Windows.Graphics.Capture` may be an approach that could be used for newer implementations of this environment.

5.7 Summary

This chapter has discussed the design and implementation of a virtual environment that provides a platform for investigation into Research Questions 3 and 4. The virtual environment described can display and interacting with conventional applications running on a Microsoft Windows 7 host system using existing peripherals without modification to the applications being displayed. It can additionally be configured to present static content in a variety of forms to facilitate user studies and testing in a controlled fashion.

Following informal evaluations of the developed environment, the design choice of using conventional keyboard and mouse for environmental inputs was noted as functional and allowing for natural interactions with conventional applications presented. The immersive nature of the utilised HMDs restricting view of these peripherals is noted as a issue that should be considered and addressed in future environments if these inputs are to be used. Observations of alternative mesh presentations of virtual windows identified the planar and curved distortions as most preferred by users, though opinions on the suitability of each to different tasks presented mixed views that require evaluation.

Chapter 6

Reading Study

In this chapter, a user study is presented that measures reading performance within the virtual environment described in Chapter 5, and contrasts this against baseline performance of reading from a conventional display. Different sizes and representations of virtual window, and different immersive virtual reality (IVR) head-mounted displays (HMDs) are tested to examine their impact on reading performance within the virtual environment. This seeks to address Research Question 3, which stated:

> *With respect to [reading], to what extent does changing from a conventional environment to a virtual one affect performance?*

It is demonstrated that it is possible to achieve a baseline equivalent reading performance within the virtual environment when using small windows, with performance falling below baseline on larger samples. Curved window distortions are shown to be more suitable for reading tasks than the planar alternative at larger window sizes, with reading accuracy matching baseline while incurring a speed penalty. Distances at which text is considered readable is identified as variable among study participants, and feedback given by participants is discussed, identifying text attributes and environmental changes that may affect the readable distance and overall usability of the virtual environment.

Supporting material for this chapter can be found in Appendix C. Parts of this chapter contents were included in a conference paper, the full text of which can be found in Appendix B. Evaluation of reading performance with respect to error rate is presented in this chapter additional to the material presented in this paper.

6.1 Questions

In Section 4.3, reading performance was identified as the primary metric for evaluating the suitability of a virtual environment for replacing a conventional computing system and workflow such as that observed in Chapter 3. Through measuring the reading performance of a user within the virtual environment presented in Chapter 5 and contrasting that against their performance when reading from a conventional display, it can be determined if the introduction of the virtual environment and associated IVR HMD had a positive or negative effect on reading tasks.

During development of the virtual environment, usage of the environment for extended periods was undertaken by the researcher. It was noted that several factors affected the readability of text, namely the form the text content took, the size and geometry of the virtual window on which text was presented, the distance the virtual window was positioned from the perceiver, and the IVR HMD used to display the virtual environment. To evaluate the suitability of the environment fairly, each of these needed to be controlled and evaluated.

6.1.1 Text Content

Content presented within the virtual environment is not designed with this method of presentation in mind, instead presenting conventional application windows captured directly from the host system. The applications drawing these windows are not aware that their destination is within a virtual environment, and do not make any attempt to follow best practice for display of text within such an environment. While some applications can be modified such that their window contents can be read with greater ease in this unexpected environment, many others cannot, and reading content from these windows presented more of a challenge to users of the virtual environment during development.

When considering the suitability of the virtual environment as a possible replacement for existing workflows, there will likely be many applications that fall into the latter category of not easily modifiable windows. As such, when evaluating the suitability of the virtual environment, reading performance should be measured from windows that have not been altered to improve their appearance to improve performance, instead testing the worst-case presentation of window content.

6.1.2 Virtual Window Size

When displaying text of any form within the virtual environment, it was observed that smaller virtual windows felt more comfortable to read from, with wider windows becoming increasingly uncomfortable to read. While no measurements of reading performance were undertaken, feedback from users indicated that there was a belief that reading performance was negatively impacted as the window width increased. It is well established that the line width of text has an impact of reading performance, with overly narrow or wide lines negatively impacting reading performance. This performance impact is seen both in print (Paterson & Tinker, 1942) and digital (Dyson & Haselgrove, 2001) reading tasks, with 52–55 characters per line identified as optimal for reading speed, however the degree to which this impacted performance within a virtual environment is not clear.

To determine the impact of differing virtual window sizes, several different sizes representing window sizes that might be encountered while using a conventional computing system should be trialed. This should similarly be trialed using a conventional display to allow for fair comparison between display devices.

6.1.3 Virtual Window Geometries

When presenting a virtual window within the virtual environment, it was immediately apparent that use of a planar mesh disadvantaged content positioned to the extremes of the mesh, as this content was necessarily further from the perceiver than that found at the centre. Use of a curved virtual window geometry sought to address this issue by ensuring that all points on the mesh were approximately equidistant from the perceiver, treating the content more fairly. While users of the virtual environment during development broadly agreed the curved virtual window presentations were superior for reading tasks, several strongly disagreed and indicated that they believed the curved presentation negatively affected their reading performance.

Additional feedback suggested that larger curved windows made users feel uncomfortable, with the sensation of being enveloped by the content not well received. Several users described the experience as “oppressive” and “unnatural feeling”, which maps to the claustrophobic effect noted by McKenzie et al. (2017) when content is curved about the perceiver. Not all users reported this sensation even when questioned about it specifically, and even those that noted this effect did not indicate that they believed the sensation impacted on their reading performance.

As feedback from users is mixed around preferred geometries for reading, both should be utilised to measure the degree of reading performance impact that these have when compared to the baseline. In addition to measuring performance, opinion data should be collected for preferred geometry to determine if users preference aligns with the most performant presentation, or if less ideal options are preferred due to feelings within the environment.

6.1.4 Presentation Distance

Position of content within the virtual environment can be controlled by the user, with virtual windows able to be freely rotated around the perceiver, and moved nearer to, or further from the perceiver. When initially presenting the environment to users, common feedback suggested that the chosen default distance at which virtual windows were presented was non-ideal, however opinion differed as to whether the content was too close and not making ideal use of space, or too far to the point that readability was impacted.

While exact positioning of content is likely to be driven by individual user preference, determining a distance at which text is comfortably readable whilst making efficient use of available space is necessary to both present a good first impression of the virtual environment, and to calculate the space available for window arrangement. This position should be determined through bringing content forward until it is readable by users, and the position recorded to see if this is consistent between users.

6.1.5 Device Preference

The generation 1 virtual environment was designed to be displayed on both the Oculus Rift DK1 (DK1) and Oculus Rift DK2 (DK2) IVR HMDs and offers feature parity across both devices to ensure a fair comparison. It was observed during development that the DK2 was preferred by most users when experiencing the virtual environment, however several users commented that text appeared to be “crisper” with easier to distinguish characters on the DK1, with text in the DK2 occasionally described as “like reading through smeary glasses”. While this opinion was not shared by all users and was called out as an intermittent issue by those that mentioned the effect of the DK2 on reading, this suggests that there may be a measurable performance difference between these devices that may point to HMD upgrades having an undesired impact on reading.

To determine the impact on reading performance the choice of HMD may have, both devices should be tested with the same virtual window sizes and geometries and compared.

6.2 Hypotheses

From the discussion above, three groups of hypotheses can be made for evaluation, covering environmental configuration, reading performance, and user opinion of the virtual environment. For the category of environmental configuration, determining the user's preferred virtual window geometry and positioning yields the hypotheses:

Hypothesis 6.1. *Curved distortion for virtual windows will be preferred by participants for reading over planar virtual windows.*

Hypothesis 6.2. *Participants will identify an optimal reading distance for virtual windows displayed within the virtual environment.*

Hypothesis 6.3. *Optimal reading distances within the virtual environment identified by participants will be consistent amongst participants.*

For the reading performance, identifying the degree of impact that virtual window geometry has on reading performance, the impact of different HMDs on reading performance, and the overall impact of the virtual environment on reading performance when compared with conventional systems yields:

Hypothesis 6.4. *Reading performance recorded for participants reading from a curved virtual window will be better than those recorded for the same participant reading from the same sized planar virtual window.*

Hypothesis 6.5. *Reading performance will be highest for experiments using a conventional display, followed by experiments using the DK2, with experiments using the DK1 having the lowest performance.*

And finally, determining the user opinion on the virtual environment experience when compared with a conventional system:

Hypothesis 6.6. *Participants will prefer the conventional display for reading, followed by the DK2, and the DK1 being least preferred.*

6.3 Experimental Design

To address the questions and hypotheses outlined above, a reading experiment was designed where participants would be required to read aloud, with their responses timed and verified for correctness. This experiment required participants to read text from virtual windows presented within the DK1 and DK2 HMDs, and windows displayed on a conventional display to provide baseline reading performance values for comparison. Virtual windows and conventional windows were displayed in three different sizes (See Table 6.1) and using curved and planar virtual window distortions to capture reading performance for a range of different reading scenarios. Static captured versions of conventional applications were chosen as the display target to ensure consistency of content and appearance between experiments and participants. Audio recording was used to collect participant spoken words when reading, and verbalised feedback, for later analysis. Written questionnaires were utilised to capture participant opinions and feedback about reading in each of the tested scenarios, and overall feedback regarding the test.

As reading of text was the subject of this test, the complexity of the text was required to be manageable by all possibly selected participants. Additionally, the chosen conventional application and simulated workflow needed to represent a familiar task and experience for these same participants.

To ensure that chosen text was understandable by the greatest number of participants, text samples were required to have a Flesch reading ease (FRE) score of 60–70, which represents a "Plain English" readability score understandable by 13–15 year old[1] students. The FRE score was calculated using the following equation.

$$FRE = 206.835 - 1.015\left(\frac{total\_words}{total\_sentences}\right) - 84.6\left(\frac{total\_syllables}{total\_words}\right)$$

The conventional application and workflow chosen for reading samples was the Google Chrome web browser displaying a wikipedia-style article, where participants would be required to first seek out a section of text to read, then read it aloud. Outside of the slightly unusual reading aloud of the text, this represented a familiar task and presentation of text within a known conventional application for the chosen participant group. Articles from the English Wikipedia scored low on the FRE scale, indicating more complex text than was desired, so articles about major world cities found in the Simple English Wikipedia were used instead. Articles found within the Simple English Wikipedia are intended for learners of the English language and attempt to use only the 1,000 most frequently used English words to compose articles, and many articles exhibit high FRE scores reflecting this simpler English language.

---

[1] Ages derived from converting the given American school grade range of "8^{th} and 9^{th} grade" given by Flesch, 1979, Chap. 2, into New Zealand school years of Year 9 & 10, then taking the upper and lower ages from this range.

Table 6.1: Sizes of captured text samples rendered with the Google Chrome browser. Pixel resolution and aspect ratio is given, along with the aliases used within this chapter to describe each sample size.

Resolution	**Aspect Ratio**	**Name**
800 × 600	4:3	**S**mall
1366 × 768	16:10	**M**edium
1920 × 1080	16:9	**L**arge

Table 6.2: Collections that sample articles were grouped into, detailing the number of samples collected at each sample size.

Collection	**S**	**M**	**L**
Seq 0	2	2	2
Seq 1	1	1	1
Seq 2	2	2	2

Fifteen articles were selected, and local copies of the pages were created. From each article, a body of text containing 110–130 words with the required FRE score of 60–70 was identified as a reading target. Chosen bodies of text were modified to remove Wikipedia superscript citation markers and other formatting not relevant to the text being read. Articles were grouped into three collections, with two of these containing 6 articles, and the other containing 3. The text of each of these chosen bodies of text, and the collection to which they were assigned can be seen in Appendix C.8.

Each local article was then opened with Google Chrome, and the contents of the main browser frame — excluding the browser controls and window decorations — was captured at one of the three different sizes seen in Table 6.1. Within each collection of articles, an equal number of Small, Medium, and Large sizes were collected, as shown in Table 6.2.

Using a modified version of the generation 1 virtual environment described in Section 5.4, an experimental setup was designed to allow for captured text samples to be displayed on either the curved or planar virtual window distortions within the virtual environment. This modified environment allowed for changing the displayed distance of a virtual window through use of a mouse scroll-wheel, and for the distortion and contents of virtual windows to be changed according to predefined orders. These modifications allowed for reproduction of repeatable scenarios that could be presented to participants within the virtual environment. This modified environment was set up to support both the DK1 and DK2 headsets, with only rotational tracking enabled. The virtual cameras were positioned centred about the environment origin, oriented facing towards $\begin{bmatrix} 0 & 0 & 1 \end{bmatrix}$.

A test plan was created for testing within each of the chosen HMDs, that would allow for either of the two large sample collections (6 sample images) to be displayed sequentially on planar and curved virtual windows at three different sizes within the described environment. The order of display was configured such that the virtual window distortion would alternate between the two options, and sizes would be presented in pairs of the same size. Initially the virtual presented a simple green environment, with a multicoloured rotating cube floating ahead of the user, the position of which could be changed by the participant through use of the mouse scroll-wheel to demonstrate the function of this input. The first sample, and successive samples were displayed by pressing the space-bar key on the connected keyboard. This would first clear the virtual environment of all existing elements and set the world colour to cornflower blue, before placing a new virtual window of appropriate size and distortion and displaying the sample upon it. Samples displayed within the environment had an initial position that was a large distance away from the perceiver, far beyond where any content could be read, requiring the virtual window to be moved closer in order to be read. When the virtual window position was changed, a log file entry was created, specifying the timestamp and current position of the virtual window within the environment. Once all samples had been displayed, the world colour changed to red to indicate completion of the experiment, and no further samples were displayed.

A second test plan was created for testing on a conventional display, displaying the remaining smaller collection (3 samples) of captured article images in a randomised order on the centre of a conventional display, with any surrounding margin coloured black. The initial view of the display was completely black until the first sample was shown and returned to a full black screen when all samples had been displayed. As with the virtual environment experiments, the first and subsequent samples were displayed in sequence, with the spacebar on the keyboard advancing the samples.

To collect participant general participant feedback following experiments, two questionnaires were designed. The first of these, referred to as the post-test questionnaire, was to be given following completion of experimentation with each of the devices under test. This collected participant feedback on discomfort or disorientation during the experiment, the experience with reading during the test, and distortion preferences for experiments within the virtual environment. A copy of this questionnaire can be found in Appendix C.6.

The second questionnaire, referred to as the exit interview questionnaire, was given to participants following completion of all three device experiments. This collected demographic data about participants, and overall device preferences having experienced the available choices. Additional room was provided for additional feedback about the virtual environment design, and the experience of reading within it. A copy of this questionnaire can be found in Appendix C.7.

6.3.1 Equipment

The experimental software was run on a computer installed with the Microsoft Windows 7 operating system. This computer used an Intel i7-4770 processor, with 16GB of system memory installed, and an NVidia GeForce GTX750Ti graphics card, with 2GB of dedicated video memory. Acting as the conventional display, a Dell U2312HMi LCD monitor utilising a 23-inch panel running at a 1920×1080 resolution at 60Hz was used. Brightness and contrast settings for this display were reset to the factory setting of 75%. A 104-key US-layout keyboard was provided along with a 2-button mouse with mouse-wheel for interaction with the experiment.

The IVR HMDs used were the DK1 and DK2. Both headsets were configured using the “A” lens sets, which are suitable for long sighted users, or users with no requirement for vision correction. Eye relief dials were set to their furthest setting, with the distance between the lens and display at the maximal setting. The DK1 external control box was used to configure to brightness and contrast of the device to 50% of it’s maximal value, providing a similar appearance to both the DK2 and conventional display. The DK2 positional tracker was positioned 1 metre away from the participant’s seat, facing directly towards the participant at approximately eye level.

For experiments utilising HMDs, a secondary 22-inch monitor was connected to the system, which displayed a monocular mirror of the participant’s view for the researcher to observe. This secondary display was disconnected for experiments using the conventional display.

6.3.2 Method

On arrival, participants were introduced to the HMDs they would be using for the study, an overview of the experimental process was given, and a randomised participant ID was assigned. Participants were seated at a desk with their chair adjusted to a comfortable height, and a conventional display was adjusted ahead of them such that the participant was looking at the horizontal and vertical centre of the display when positioning their head at a neutral forward position. Audio recordings were started at this point to capture participant responses and feedback for later analysis.

Following this initial setup, the first HMD was fitted and calibrated for the user, with the Oculus inter-pupillary distance (IPD) utility used to calculate the participant IPD. The first virtual environment experiment was then started, and the participant's hands were positioned on the mouse positioned on the desk in front of them. The participant was prompted to manipulate the mouse scroll wheel to demonstrate the functionality of this using the displayed cube, and to ensure the participant was able to interact with the peripherals. The experiment was then advanced by the researcher, with the first of a sequence of virtual windows appearing at an unreadable distance ahead of the participant's viewpoint. The participant was prompted to use the mouse wheel to position the virtual window nearer and further away from their viewpoint and was instructed to position the virtual window such that the text on the windows was at the greatest distance away from them whilst still remaining comfortably readable. The participant was then prompted to read an indicated paragraph of text from the presented sample aloud. Once complete, the researcher advanced the process to display the next sample, at which point the process was repeated, continuing until all samples had been read by the participant.

On completion of all reading tasks, the HMD was removed from the participants head, and a post-test questionnaire was provided for the participant to complete. While completing the questionnaire, additional detail for provided responses and general feedback about the experience was requested by the researcher and captured in the continuing audio recording.

The participant's attention was then drawn to the conventional display ahead of them, and the participant was positioned in front of the display at a distance of 750mm as measured from the participants nose to the display surface. The participant was asked to remain at that distance, with the researcher prompting them to move back or forward if they adjusted during the experiment process. No peripherals were provided to the participant. The researcher displayed the first text sample, centred on the conventional display with black borders filling any additional space. The participant was then prompted to read an indicated paragraph of text aloud. Once complete, the researcher advanced the process to display the next sample, where the process was repeated. Once all samples had been read by the participant, a post-test questionnaire was provided and completed by the participant as described above.

The participant was then fitted with the remaining HMD, which was recalibrated in the same manner as the first HMD. The experimental process as described with the first HMD process was repeated in the same manner and concluded with a post-test questionnaire.

Finally, the participant was provided with an exit interview questionnaire which collected their demographic and device preference information. Additional feedback was requested by the researcher, and additional detail around participant observations was sought. Following this, audio recordings were stopped, and the participant was allowed to leave.

6.4 Participants

Participants for this study were sourced from the University of Waikato student population, with the majority being from within the Faculty of Computing and Mathematical Sciences undergraduate student body. Potential participants were approached by the researcher within the faculty buildings, and were provided with an overview of the study, with contact details for the researcher provided if the approached person expressed interest. Through this process, and word of mouth via interested persons, a pool of volunteers was established. 24 participants were ultimately selected from this pool of volunteers, with preference given to those with no vision issues, or those who had minor vision issues. Participants requiring the use of spectacles for correcting vision were excluded from selection due to unreliable compatibility with the HMDs. No preference was given to participant age, gender, or ethnicity during selection.

Of the selected participants, 16 identified as male, and 8 as female. 19 participants reported no vision issues, 2 participants had corrected vision through the use of contact lenses, and 3 were classified as having uncorrected vision issues affecting a single eye. This is summarised in Figure 6.1.

Participants were briefly surveyed prior to study commencement to determine their preference for reading environments, and for experience with IVR devices. 23 participants indicated a strong preference for consuming text on-screen, with 1 participant indicating a preference for print media. Only 1 participant indicated extensive experience with IVR devices, with other participants indicating slight or no experience with these.

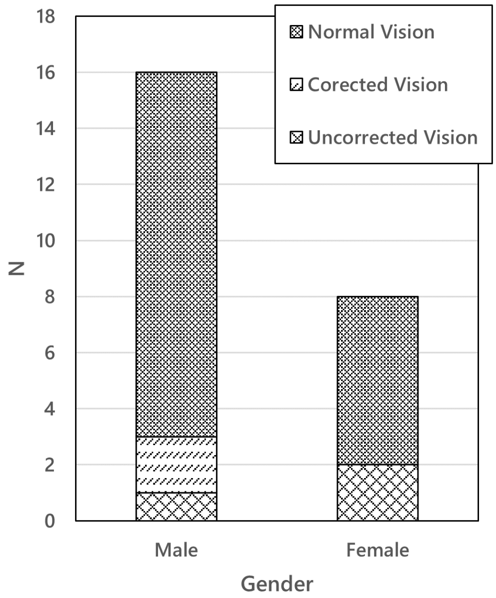

Figure 6.1: Study participant demographic information showing numbers of participants with vision issues.

6.5 Data Processing

Audio recordings collected during each participant and experiment were processed to split each experimental step into separate files for easier processing. Recordings for each participant were organised into folders, grouped by participant ID, device under test, and individual sample. Files for each participant were split into individual clips for each sample, with the clips beginning at the point each new sample was displayed and ending when the experiment moved to the next sample. This ensured that participant observations and feedback for each sample were captured, along with the sample itself.

Individual participant sample clips were further processed to isolate only the read-aloud text. These clips began at the point the first word of the sample was vocalised and ended following the last word of the sample. These clips were edited to remove comments from the participant that were clearly directed to the researcher, including but not limited to observations such as “There are weird colours around the text here” or “It feels weird having such a large panel in front of me”. Other comments which were clearly vocalised commentary on the read material, such as “Huh, I didn’t know that” and “That’s a weird word” were similarly removed. Removal began at the leading edge of the audio waveform making up the comment and ended at the trailing edge of the same waveform. Silence proceeding and following the removed sections were retained. Vocalisations that were not clearly directed to the researcher, or that were not deliberate asides such as “I think” following an unfamiliar word, or filler words like “Um” and “Ahh” were not removed from these clips.

Individual sample clips were then processed using the running record techniques described by Clay (2000), comparing each spoken word against the written text to establish the correctness of the read-aloud text. This process recorded an error for each omitted, inserted, or substituted word from the text. Initialism and contraction of words was not recorded as errors, allowing “we are” and “we’re” substitutions, and “U.S.” and “United States” substitutions without penalty. Substitution errors with proper nouns were not recorded in cases where the substitution was due to a mispronunciation of an unfamiliar word, with “Chesapeake” pronounced as “Cheese peak” serving as a notable example observed in multiple participants. In cases where an error described above was made, but the participant corrected themselves, the error was removed and recorded as a self-correction. Vocalisations made as part of the reading process, such as “I think”, “Umm”, and “Ahh”, were deliberately excluded from processing and treated as if they were not present in the recording.

Following completion of the running record process, several different scores were calculated using the number of errors made for a sample, total words in the written sample, and the time taken to read the sample — which was taken as the length of the processed audio clip the running record was performed on. The first of these was an accuracy of reading value, calculated for each sample and participant, using the following equation.

$$AR_{sample} = 100 - \frac{errors}{wordcount} \cdot 100$$

The second score calculated was an accurate words per minute (AWPM) value, which was calculated using the following equation.

$$AWPM = \frac{wordcount - errors}{\mathcal{T}} \times 60$$

Each of the accuracy of reading and AWPM scores collected from HMD trials was then compared against the baseline result obtained by reading from the conventional display. Values were compared against the baseline for the same participant, from the same sample size between these devices. The percentage difference for each of these was calculated with the following equation.

$$V_{diff} = \frac{V_{baseline} - V_{sample}}{V_{baseline}}$$

Table 6.3: Virtual window distortion preferences collected during the post-test questionnaires for each participant. Results are overall opinions across all virtual window sizes.

Distortion	DK1	DK2
Flat	2	2
Curved	16	19
No Pref.	5	3

Log files produced by the testing application were processed to extract the final distance at which samples were placed prior to the participant reading aloud the sample. These values were recorded in a tabular file for later use.

6.6 Results

This user study sought to answer several different hypotheses relating to configuration, performance, and preference within an IVR environment. In this section, each of the hypotheses stated in Section 6.2 will be addressed in turn.

6.6.1 Preferred Distortion

To address Hypothesis 6.1 which stated:

> *Curved distortion for virtual windows will be preferred by participants for reading over planar virtual windows.*

the data collected in the post-test questionnaires given following experiments with each of the tested devices was tabulated, which can be seen in Table 6.3. From this data, the curved distortion was preferred by participants on both devices, with 67% of participants indicating curved as their preferred distortion when using the DK1, and 79% when using the DK2. A minority of participants indicated that they preferred the appearance of planar panels on both HMDs regardless of the displayed virtual window size, with the remaining responses indicating no preference for either distortion.

In feedback collected along with these ratings, several common feedback items were present. The first, given by all participants, was that the participants believed that for larger virtual window sizes the curved distortion was able to present the text content of the application windows in a manner that allowed for easier reading. As a caveat to this point, several participants expressed that this easier reading came at the cost of comfort, both within the environment where the curved virtual windows felt oppressive, and when reading due to the larger angles that their necks were required to bend to track the lines of text. Some participants made comments during the study process, that reading from the curved distortion virtual windows felt “unnatural”, and “like reading off the side of a side of a building” This feedback and result mirrors observations made about the virtual environment during development (See Section 5.6) and suggests that while the curved distortion can solve the problem of reading from large virtual windows in the developed environment, this does come with trade-offs.

Participants also indicated that for small virtual windows, that they felt there was no difference in reading ease when using the curved or planar distortions. This was also the case with medium virtual windows displayed on the DK2, however for medium virtual windows displayed on the DK1 several participants noted that the curved distortion did offer a slight advantage in reading ease when compared to the planar representations.

From these results, it can be concluded that for larger samples that the curved distortion was preferred by a large proportion of participants, but for smaller sizes there is no strong preference to either distortion. From a general perspective, it can be concluded that curved distortions would be preferred over a wider range of virtual window sizes for reading, but this may not hold for other tasks where perceived downsides may have a stronger effect.

6.6.2 Optimal Reading Distance

To address Hypothesis 6.2, which stated:

> *Participants will identify an optimal reading distance for virtual windows displayed within the virtual environment.*

final virtual window positions selected by participants prior to reading the displayed text samples aloud were extracted from experimental logs. Values for each tested HMD, distortion, and sample size were tabulated, and descriptive statistics calculated for each combination. The results of this can be seen in Table 6.4.

From these results, several observations can be made. The most immediate of these is that virtual windows displayed using the DK2 were positioned at greater distances from the perceiver when compared to virtual windows with the same distortion and size on the DK1. Distances chosen within the DK1 and DK2 vary, with no consistent position identified for each HMD independent of any other display variable.

Table 6.4: Descriptive statistics for final selected virtual window distances within the virtual environment by study participants.

Device	Distortion	Size	N	$\bar{x} \pm s$ ($\mathbb{D}_n$)	c_v	Q_0	Q_1	Q_2	Q_3	Q_4
DK1	Planar	S	23	1.02 ± 0.19	19.05	0.56	0.90	1.00	1.08	1.46
DK1	Planar	M	23	0.97 ± 0.23	23.77	0.62	0.79	0.90	1.19	1.42
DK1	Planar	L	22	0.88 ± 0.22	25.41	0.50	0.70	0.88	1.01	1.53
DK2	Planar	S	24	1.49 ± 0.31	20.71	0.86	1.26	1.55	1.69	2.20
DK2	Planar	M	23	1.48 ± 0.31	20.67	0.99	1.20	1.50	1.70	2.13
DK2	Planar	L	24	1.33 ± 0.41	30.59	0.60	1.05	1.32	1.50	2.09
DK1	Curved	S	23	1.11 ± 0.19	17.13	0.78	1.01	1.08	1.20	1.51
DK1	Curved	M	23	1.21 ± 0.16	13.02	1.00	1.05	1.20	1.30	1.50
DK1	Curved	L	23	1.25 ± 0.18	14.67	0.92	1.10	1.29	1.40	1.56
DK2	Curved	S	24	1.48 ± 0.30	20.26	1.01	1.22	1.50	1.69	2.10
DK2	Curved	M	24	1.60 ± 0.28	17.30	1.14	1.37	1.64	1.72	2.27
DK2	Curved	L	24	1.69 ± 0.39	22.83	0.90	1.41	1.69	1.95	2.56

Secondarily it can be seen that two trends in chosen distance are observed, dependent on the distortion used. Planar virtual windows show a trend downwards with chosen distance as the size of samples increases, while curved virtual windows show an upward trend with the distance increasing as the positioned further away as the sample size increases.

From these results, it can be concluded that the optimal reading position for virtual windows displayed in this virtual environment is not constant between participant, device, distortion, or sample size, with each playing a part in the chosen position.

6.6.3 Inter-participant Reading Distance Consistency

To address Hypothesis 6.3, which stated:

> *Optimal reading distances within the virtual environment identified by participants will be consistent amongst participants.*

the coefficient of variation (c_v) values presented in Table 6.4 were examined, and an interval plot generated showing 95% confidence intervals for mean selected distance, seen in Figure 6.2.

From examining the c_v values for each sample, significant variation is seen in selected positions across devices, distortions, and sample sizes. c_v values trend upwards along with sample size for planar virtual windows viewed on either tested HMD, suggesting a range of acceptable positions for larger planar panels. c_v values recorded for curved virtual windows do not show a clear trend across sample size but are noted as being lower than the values found for the

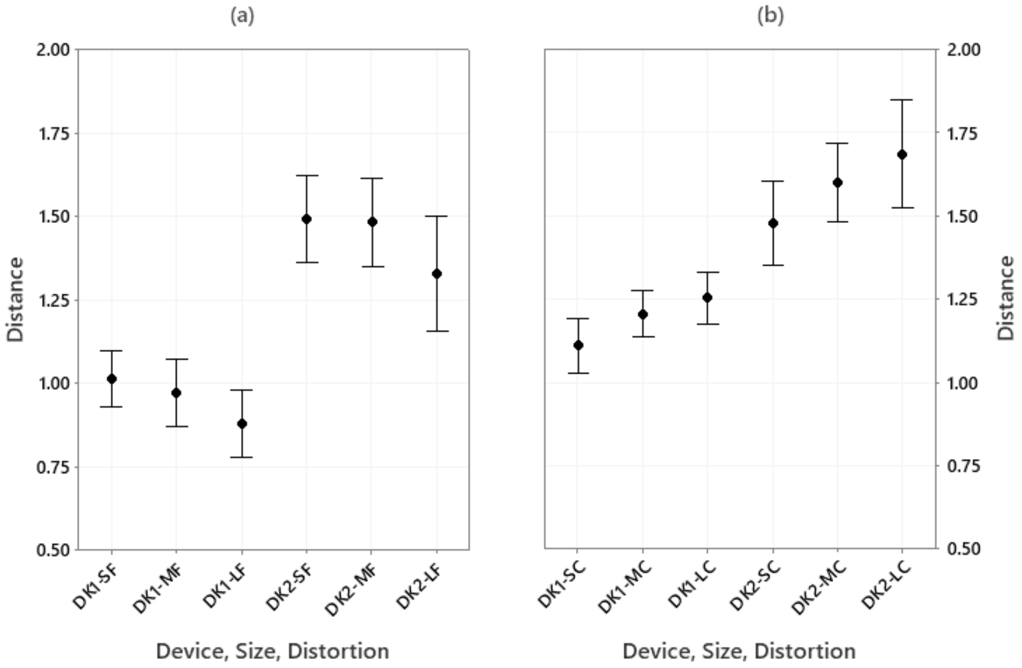

Figure 6.2: Interval plots for final selected virtual window distances within the virtual environment; *(a)* for planar distortion meshes, *(b)* for curved distortion meshes. A 95% confidence interval is shown along with the mean distance for each device and sample size.

same device and sample size displayed on a planar virtual window. This indicates that while variation was observed in positioning of samples in all cases, that the curved virtual windows presented a narrower range of distances that were considered optimal by participants when compared to samples displayed on planar virtual windows.

When examining Figure 6.2, displayed 95% confidence intervals presented for samples displayed on the DK1 are narrower than those shown for samples displayed for DK2 distances. While this could be presented as evidence that some attribute of the DK1 allowed for a more consistent selection of optimal reading distance, it is more likely attributed to a smaller range of acceptable reading distances available within the DK1 compared to the DK2, relating to the effective device resolution.

From these results, it can be concluded that the optimal reading distances are not consistent for participants, with participants instead identifying a broad range of acceptable reading distances within the virtual environment. This suggests that personal preference plays a significant part when selecting readable distances when minimal constraints on selection are given.

Table 6.5: Paired T-tests for distortions pairs, over the percentage difference from baseline of same sample size.

Device	**Size**	$\bar{x} \pm s$ (Planar)	$\bar{x} \pm s$ (Curved)	$\bar{x} \pm s$ (Difference)	**95% CI**	**t**	**p**
DK1	S	$-6.12\% \pm 15.20$	$-3.90\% \pm 20.36$	$-2.22\% \pm 16.83$	$(-9.50, 5.05)$	$t(22) = -0.63$	0.533
DK1	M	$-24.87\% \pm 13.66$	$-12.00\% \pm 15.43$	$-12.87\% \pm 12.90$	$(-18.45, -7.30)$	$t(22) = -4.79$	0.000
DK1	L	$-57.59\% \pm 14.32$	$-22.25\% \pm 17.36$	$-35.33\% \pm 17.62$	$(-43.15, -27.52)$	$t(21) = -9.41$	0.000
DK2	S	$3.51\% \pm 11.85$	$3.57\% \pm 16.89$	$-0.06\% \pm 16.72$	$(-7.12, 7.00)$	$t(23) = -0.02$	0.987
DK2	M	$-7.39\% \pm 10.36$	$-5.70\% \pm 7.76$	$-1.69\% \pm 10.54$	$(-6.25, 2.86)$	$t(22) = -0.77$	0.449
DK2	L	$-21.88\% \pm 14.58$	$-6.51\% \pm 9.63$	$-15.37\% \pm 15.53$	$(-21.93, -8.82)$	$t(23) = -4.85$	0.000

6.6.4 Reading Performance Superior with Curved Distortion

To address Hypothesis 6.4, which stated:

> *Reading performance recorded for participants reading from a curved virtual window will be better than those recorded for the same participant reading from the same sized planar virtual window.*

the AWPM values calculated for each participant were biased against the baseline reading performance collected using the conventional display sample of the same size, then evaluated using a paired t-test, comparing the difference in AWPM against baseline performance for curved versus planar distortions for each device. This test was run with the hypotheses:

$$H_0 : \mu\left(AWPM_{Planar}\right) - \mu\left(AWPM_{Curved}\right) = 0$$

$$H_1 : \mu\left(AWPM_{Planar}\right) - \mu\left(AWPM_{Curved}\right) \neq 0$$

The results of this evaluation can be seen in Table 6.5, and a visualisation of the data can be seen in Figure 6.3.

From these results, for large virtual window samples displayed on both tested HMDs, that there were statistically significant differences in AWPM values ($p = 0.000$), with the curved distortion showing a greater mean AWPM value versus baseline than the planar distortion virtual windows. This difference is most dramatic for the DK1, where the planar distortion offered a mean AWPM drop of -57.59% compared to baseline, with the curved distortion only incurring a -22.25% mean drop in performance. The impact was not as severe for the DK2 but was still significant, with a -21.88% drop in mean AWPM against the baseline, with a -6.51% drop in performance with the curved distortion virtual windows.

For medium virtual window samples displayed on the DK1, a statistically significant difference in AWPM values was also observer ($p = 0.000$), with similar results to those seen with large panels. The curved distortion recorded a better mean AWPM performance drop versus the baseline than with the

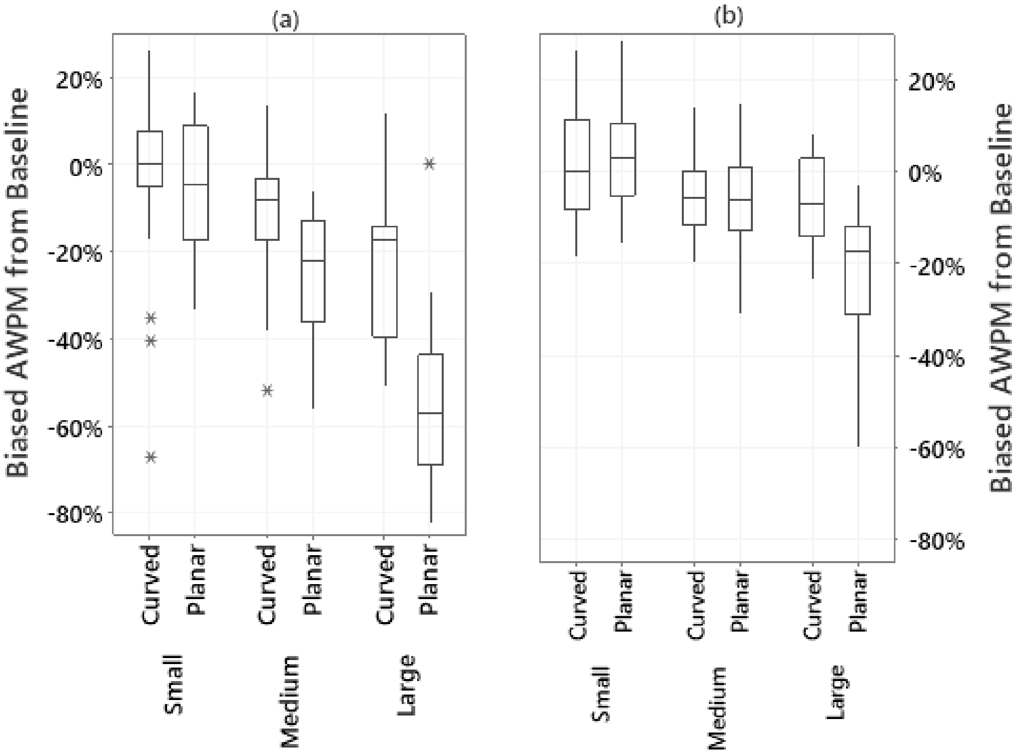

Figure 6.3: Mean AWPM percentage differences for sample sizes against baseline AWPM for the same sample size, grouped by device, virtual window size and distortion; *(a)* results for DK1, *(b)* for DK2.

planar distortion, with a -24.87% mean AWPM drop when using planar virtual windows, and a -12.00% drop when using a curved distortion. For small virtual window samples, and medium samples displayed on the DK2, no statistically significant difference in performance was noted between the curved and planar distortion mean AWPM performance against baseline values.

From these results, it can be concluded that for larger samples that the curved virtual window distortion offers measurable reading performance benefits over the baseline when compared with planar virtual windows. While this benefit was not seen for smaller virtual windows, no decrease in performance was noted when using a curved virtual window distortion when compared to planar virtual windows.

6.6.5 Conventional Display Outperforms Virtual Environment

To address Hypothesis 6.5, which stated:

> *Reading performance will be highest for experiments using a conventional display, followed by experiments using the DK2, with experiments using the DK1 having the lowest performance.*

the accuracy rating and AWPM scores for individual virtual window sizes and distortions were compared against the baseline figures for samples of the same size. These results were evaluated using paired t-tests with the hypotheses:

$H_0 : \mu\left(Metric_{IVRDevice}\right) - \mu\left(Metric_{Conventional}\right) = 0$

$H_1 : \mu\left(Metric_{IVRDevice}\right) - \mu\left(Metric_{Conventional}\right) \neq 0$

The results of these tests are presented below, in Tables 6.6 to 6.9.

First considering the AWPM metrics for planar virtual window distortions seen in Table 6.6, it can be seen that apart from small virtual window samples read within the DK2, that a statistically significant difference in AWPM scores was noted ($p \leq 0.05$). The baseline AWPM values were consistently higher than those collected from within the virtual environment, with significant differences noted for the larger virtual windows. The samples read from the small virtual window within the DK2 were the only ones to not record a statistically significant result ($p > 0.05$), with the mean recorded AWPM value being close to that recorded as the baseline metric.

When examining the AWPM metrics for curved virtual window distortions seen in Table 6.7, no statistically significant difference ($p > 0.05$) was noted in AWPM performance between the baseline value, and the values collected for the small virtual windows on either device. For medium and large panels on both tested HMDs however, a statistically significant difference ($p \leq 0.05$)

Table 6.6: Paired t-test comparing AWPM performance for planar virtual windows to baseline AWPM results for collected from the conventional display test using the same sized sample.

Device	**Size**	$\bar{x} \pm s$ (Planar)	$\bar{x} \pm s$ (Baseline)	$\bar{x} \pm s$ (Difference)	**95% CI**	**t**	**p**
DK1	S	134.84 ± 22.82	144.93 ± 19.53	-10.09 ± 23.34	$(-20.18, 0.00)$	$t(23) = -2.07$	0.050
DK1	M	115.68 ± 25.44	154.40 ± 20.41	-38.73 ± 23.32	$(-48.81, -28.64)$	$t(23) = -7.96$	0.000
DK1	L	67.75 ± 21.54	162.88 ± 25.35	-95.13 ± 32.60	$(-109.58, -80.67)$	$t(22) = -13.69$	0.000
DK2	S	151.00 ± 25.70	146.15 ± 20.02	4.85 ± 17.33	$(-2.46, 12.17)$	$t(24) = 1.37$	0.183
DK2	M	142.29 ± 21.21	154.40 ± 20.41	-12.11 ± 17.00	$(-19.46, -4.76)$	$t(23) = -3.42$	0.002
DK2	L	127.67 ± 31.46	162.90 ± 24.52	-35.22 ± 23.33	$(-45.08, -25.37)$	$t(24) = -7.40$	0.000

Table 6.7: Paired t-test comparing AWPM performance for curved virtual windows to baseline AWPM results for collected from the conventional display test using the same sized sample.

Device	**Size**	$\bar{x} \pm s$ (Curved)	$\bar{x} \pm s$ (Baseline)	$\bar{x} \pm s$ (Difference)	**95% CI**	**t**	**p**
DK1	S	138.26 ± 30.58	144.93 ± 19.53	-6.66 ± 29.85	$(-19.57, 6.24)$	$t(23) = -1.07$	0.296
DK1	M	135.06 ± 27.01	154.40 ± 20.41	-19.34 ± 25.03	$(-30.16, -8.52)$	$t(23) = -3.71$	0.001
DK1	L	125.91 ± 30.31	162.33 ± 24.91	-36.41 ± 29.27	$(-49.07, -23.75)$	$t(23) = -5.97$	0.000
DK2	S	150.92 ± 28.27	146.15 ± 20.02	4.78 ± 22.48	$(-4.72, 14.27)$	$t(24) = 1.04$	0.309
DK2	M	146.61 ± 21.34	155.27 ± 20.41	-8.66 ± 11.77	$(-13.63, -3.69)$	$t(24) = -3.61$	0.001
DK2	L	151.07 ± 18.98	162.90 ± 24.52	-11.83 ± 17.58	$(-19.25, -4.40)$	$t(24) = -3.30$	0.003

was noted in mean AWPM performance in the virtual environment versus the baseline. It can be seen that the baseline values, collected from the conventional display, exhibited a higher mean AWPM value than the compared virtual environment value.

Considering accuracy rating as a metric for planar virtual window distortions seen in Table 6.8, it can be seen that for small and medium virtual window samples presented in either tested HMD, that no statistically significant difference in performance is noted ($p > 0.05$). Indeed, it can be observed that for small virtual windows, that the mean accuracy rating was higher with a smal-

Table 6.8: Paired t-test comparing accuracy rating of read text for planar virtual windows to baseline accuracy rating results for collected from the conventional display test using the same sized sample.

Device	**Size**	$\bar{x} \pm s$ (Planar)	$\bar{x} \pm s$ (Baseline)	$\bar{x} \pm s$ (Difference)	**95% CI**	**t**	**p**
DK1	S	$98.763\% \pm 0.746$	$98.132\% \pm 2.058$	0.631 ± 2.025	$(-0.244, 1.507)$	$t(23) = 1.49$	0.149
DK1	M	$97.875\% \pm 2.165$	$98.860\% \pm 1.348$	-0.984 ± 2.330	$(-1.992, 0.023)$	$t(23) = -2.03$	0.055
DK1	L	$61.08\% \pm 19.98$	$98.39\% \pm 1.38$	-37.31 ± 20.08	$(-46.21, -28.41)$	$t(22) = -8.72$	0.000
DK2	S	$98.538\% \pm 1.481$	$98.209\% \pm 2.049$	0.329 ± 1.556	$(-0.328, 0.986)$	$t(24) = 1.04$	0.311
DK2	M	$98.422\% \pm 1.791$	$98.860\% \pm 1.348$	-0.437 ± 1.505	$(-1.088, 0.213)$	$t(23) = -1.39$	0.177
DK2	L	$93.04\% \pm 9.69$	$98.43\% \pm 1.37$	-5.39 ± 9.42	$(-9.37, -1.41)$	$t(24) = -2.80$	0.010

Table 6.9: Paired t-test comparing accuracy rating of read text for curved virtual windows to baseline accuracy rating results for collected from the conventional display test using the same sized sample.

Device	**Size**	$\bar{x} \pm s$ (Curved)	$\bar{x} \pm s$ (Baseline)	$\bar{x} \pm s$ (Difference)	**95% CI**	**t**	**p**
DK1	S	$98.324\% \pm 1.435$	$98.132\% \pm 2.058$	0.192 ± 1.769	$(-0.573, 0.958)$	$t(23) = 0.52$	0.607
DK1	M	$98.869\% \pm 1.297$	$98.860\% \pm 1.348$	0.009 ± 1.651	$(-0.704, 0.723)$	$t(23) = 0.03$	0.979
DK1	L	$96.96\% \pm 5.73$	$98.36\% \pm 1.36$	-1.40 ± 6.08	$(-4.03, 1.23)$	$t(23) = -1.10$	0.281
DK2	S	$98.031\% \pm 1.630$	$98.209\% \pm 2.049$	-0.178 ± 1.929	$(-0.993, 0.636)$	$t(24) = -0.45$	0.655
DK2	M	$98.608\% \pm 1.243$	$98.873\% \pm 1.320$	-0.265 ± 1.475	$(-0.888, 0.358)$	$t(24) = -0.88$	0.388
DK2	L	$98.897\% \pm 1.265$	$98.425\% \pm 1.374$	0.472 ± 1.316	$(-0.084, 1.028)$	$t(24) = 1.76$	0.092

ler standard deviation than the recorded baseline. For large virtual windows, there is a statistically significant difference in accuracy rating ($p \leq 0.05$), with the baseline recording higher accuracy rating scores, most noticeably with the large sample read from the DK1.

Evaluating the same accuracy rating metric for curved virtual window distortions, seen in Table 6.9 yields further results concluding that no statistically significant difference in accuracy rating can be seen for any sample size or device when compared against the baseline accuracy rating ($p > 0.05$). As with the small samples noted above, in some cases samples read within the virtual environment exhibited greater mean accuracy ratings with smaller standard deviations than those recorded on the baseline, though this is not consistent across devices or sample sizes.

From these results, it can be concluded that when using AWPM as a metric, reading text from a conventional display will result in an AWPM score equal to, or greater than that which could be recorded on either HMD, with any sample size or distortion. From observing the mean AWPM scores between the DK1 and DK2, the DK2 consistently provides better AWPM scores than the DK1, for any sample size or virtual window distortion. This result aligns with the hypothesised device ordering presented within this chapter.

When considering accuracy rating as the metric for reading performance however, this order does not hold as definitively. When examining the accuracy rating of text read from a planar virtual window, the ordering is maintained, with larger windows showing the conventional display having the greatest accuracy rating, followed by the DK2, then the DK1. For curved virtual window distortions though, this is not seen at any sample size, with no statistically significant difference in accuracy ratings observed between tested devices.

Table 6.10: Overall device preferences results collected from the exit interviews with each participant.

Device	1^{st}	2^{nd}	3^{rd}
Monitor	22	2	0
DK2	1	18	5
DK1	1	4	19

6.6.6 Device Preference Order

To address Hypothesis 6.6, which stated:

> *Participants will prefer the conventional display for reading, followed by the DK2, and the DK1 being least preferred.*

data collected in the exit interview conducted at the conclusion of the study was tabulated, with the results presented in Table 6.10.

From this data, a clear device preference order can be seen, with 92% of participants indicating the conventional display as their preferred device, followed by 75% of participants preferring the DK2, and finally 79% indicating the DK1 as the device at the bottom of their preference order. Feedback collected along with the device preference order provides some insight into the rationale behind this ordering.

Most of the participants indicated that the conventional display was their most preferred device due to its familiarity, and the clarity of the text presented on it. The appearance of text on the conventional monitor was described as "crisp" and easy to read, leading to a perceived increase in reading performance on this device. Several participants identified that due to the amount of experience they had enjoyed over years of reading from such displays, that there was nothing they had to become accustomed to in order to read effectively from the display, which they identified as a strong factor in their perceived reading performance on this device.

A minority of participants ranked the DK1 over the DK2 in their device rankings, which was unexpected given the specification differences between these devices. When questioned on their choices, participants who indicated this preference order explained that while the DK1 required larger head and neck movements to read larger samples than with the DK2, that this disadvantage was outweighed by an increased clarity of text within the DK1.

From these results, it can be concluded that device preference order aligns with the hypothesised order, with the familiar conventional display ranked above either of the tested HMDs. While the ordering of the HMDs matched the predicted order for most participants, there are factors which can change this for a subset of participants, who may prioritise certain attributes of content displayed on the devices than others.

6.7 Discussion

In addition to the results presented above, several common observations were made by participants during and after the study. Some of these observations provide context for the above results, while others identify possible considerations for future experiments. Additional context and explanations for presented results have been identified during results analysis and deserve additional discussion. Each relevant observation and explanation for results is presented in this section.

6.7.1 Display Distance

In evaluating Hypothesis 6.2, it was identified that selected distances for planar virtual windows exhibited a downward trend as sample size increased, with larger samples being positioned nearer the perceiver than with smaller samples. This trend was reversed when considering curved virtual windows, with larger samples displayed at greater distances from the perceiver than smaller samples. As these trends align with the utilised distortion and sample sizes, it is valuable to identify the reasoning behind such placements to inform distortion choices for differing virtual window sizes in future.

When considering the curved distortion trend, identifying the range of head and neck movements required to centre the perceiver's view on alternate edges of the virtual window informs a possible reason behind the choices made. Using the equations described in Chapter 5, the angular measure for virtual window arcs sized for each sample size at the participant identified mean and max distances were calculated. The results of these calculations can be seen in Table 6.11. From this data, as the virtual window size increases, so to does the angular measure. This corresponds to the degree of head and neck movement required by participants to scan across the virtual window. It can be seen when comparing the angular measures presented that by increasing the distance at which the virtual window is positioned from the perceiver, that the angular measure decreases, requiring a reduced amount of head and neck movements to

Table 6.11: Angular distance between left and right edges of displayed curved virtual windows when positioned at the mean and max distances from the observer for given device and sample size combinations.

Device	Size	θ_{mean}	θ_{max}
DK1	S	82.43°	60.7°
DK1	M	129.6°	104.4°
DK1	L	175.4°	141.0°
DK2	S	61.97°	43.7°
DK2	M	97.7°	69.0°
DK2	L	130.4°	85.9°

read text upon it. The desire to reduce required head movements as the virtual window size increased may then have been the driver behind increasing mean distances correlating with increasing sample size, with participants willing to accept a decrease in text legibility to reduce physical discomfort.

When considering the decreasing distance trend observed with planar virtual windows, the reasoning behind this is more apparent than with curved distortions. Due to the geometry of the planar windows, the distance between the perceiver and the virtual window varies depending on the focus point within the virtual window, with the centre-point of the virtual window representing the minimum distance and the distance value recorded. Angle of incidence for the perceiver gaze and the virtual window also varies, starting perpendicular at the centre of the virtual window, and becoming increasingly acute towards the edges. When positioning planar windows, participants required distances nearer to the perception point, to reduce the distance between themselves and text content found on the outer edges of the virtual window, while ensuring that it wasn't positioned so close that characters at those locations could not be made out due to the acute incidence angle.

6.7.2 Distortions

An item of feedback given by participants during this study was that curved virtual windows felt “oppressive” at larger sizes, mirroring observations made by users experiencing the virtual environment during development. Some participants giving this feedback indicated that this feeling affected their reading performance, while others indicated it had no impact on their performance but would make them less inclined to use such an environment when alternatives exist. This feedback can be seen reflected within the distortion preference results, where a minority of participants indicated the planar distortion as preferred over the curved distortion, with others indicating no preference. This no preference response was given in greater numbers when considering the experience within the DK1, which suggests that the given responses were related to this point of feedback, as the enveloping nature of the curved virtual windows is exaggerated for windows positioned at distances nearer the perceiver which was observed for the DK1. Despite the undesirable feelings brought about by this effect, the larger planar panels were still considered harder to read, resulting in responses not preferring either distortion over the other.

When considering all reading performance results and feedback given regarding the different distortions, it is clear that that curved distortion does allow for improved performance when evaluating reading performance, however aspects that would drive potential users away from a virtual environment should be investigated to determine if they can be addressed. An obvious solution to this wrap-around feeling of curved virtual windows is to simply position them further away, however the distance at which virtual windows can be placed and still be used effectively is limited by the HMD specifications, which may take time to develop to the point where this can be accomplished. Alternatives distortions that minimise the wrap around effect whilst still improving on planar representations should be investigated, as well as alternative techniques that seek to address the issues of legibility of conventional applications through means other than mesh geometry.

6.7.3 AWPM Performance

When using AWPM as a performance metric for reading, a drop in performance was noted for large, curved distortion windows when compared to the recorded baseline. When using accuracy rating as the metric, no such performance drop was observed. This indicates that accurate reading was possible, but not at the same speed as on the baseline. This was not noted for small and medium virtual windows using the same distortion while positioned at comparable distances from the perceiver.

A likely reason for this decreased reading speed relates to physiological factors in reading when combined with increased head and neck movements required to view all parts of a large curved virtual window discussed above. When reading, eye movements called saccades take place where the eye rapidly jumps through points in text, becoming stationary (fixating) at points throughout the text. This movement is primarily driven by the eye itself, independent of the head and neck, though some small movements are seen during the normal course of reading. This saccadic behaviour is disturbed when viewing text through the lenses of a HMD, where clarity and detail is maximised in a narrow cone ahead of the user, and total field of view (FOV) within the HMD is limited. This impact necessitates an increased amount of head and neck movements to track text across a line, which slows the process of reading.

The reading process is further complicated when reaching the end of a line and requiring relocation to the start of the following line — a process known as a return-sweep in the context of saccadic eye movements. Ordinarily when reading, a lawnmower pattern will be followed, where as the eye approaches the end of a line of text, a return sweep will take place to allow fixations at the beginning of the following line. When the start of the line is not within the available FOV however, head and neck movements are required to reorient the user's view to allow this new-line fixation. This adds a delay to the return sweep process, and can delay visual acquisition of the new line.

When reading on these larger curved virtual windows, several participants remarked on the experience of reading long lines of text. Feelings of discomfort and a recurring "unnatural" feeling were common observations when reading along the lines, and comments on the difficulty of reacquiring the start of the new line were also common. One participant described the process as disorienting and explained that they were unable to easily identify the new line as they would when reading a smaller window, noting that they needed to scan the first words of each line in an area to identify the previously read line in order to locate the next line, adding significant mental overhead and delay.

To resolve this effect on larger virtual windows, improvements in HMD FOV and optics are needed. While limited with current HMD specifications however, a conclusion that can be drawn is that smaller virtual windows should be preferred when AWPM is a critical performance aspect. Larger virtual windows should be avoided for longer reading tasks due to the decreased reading speed and increased discomfort experienced, though can be used for shorter reference-based reading tasks if required. Investigation into supporting virtual environments could be considered to reduce the amount of head and neck movement required, such as positional changes for the virtual window being read based on head and neck position, or eye movements if an eye tracking solution is available.

6.7.4 Reading Accuracy

When evaluating reading performance within the virtual environment, it was noted that on some occasions text read within the virtual environment exhibited a greater mean accuracy rating, with smaller standard deviation values when compared to the baseline reading performance. While claiming this is a benefit of such an environment would be ideal, this is more likely attributed to increased focus of participants due to using an unfamiliar device and environment to consume text. As participants were focusing more on the process of reading within the virtual environment, they made fewer simple mistakes than they did when reading from a familiar device.

While no immediate conclusions should be drawn about this then, given the small difference in performance, sample size and novelty of the virtual environment, attention should be paid to this metric in future studies. As users become more accustomed to reading within such environments, a more conclusive measurable difference may be presented that could suggest benefits or disadvantages related to reading in these environments.

6.7.5 Legibility of Text

During the study, some participant indicated that when their head was stationary and they were focusing on a word that it was difficult to read, but when moving their head, the word became readable. This observation was often accompanied by a description of the text as being “wriggly” or “moving like it is alive” while reading. A participant with a typographical background observed that the kerning of the font appeared poorly tuned, with some characters merging forming different characters — such as neighbouring ‘r’ and ‘n’ characters appearing to form a single ‘m’ character — or losing distinctive font details that would help distinguish similar characters, such as ‘i’ and ‘l’.

A visualisation of these observations on a single word was captured using a camera directed through the lenses of the DK2, which can be seen in in Figure 6.4. It can be seen in this figure that the characteristics of characters are indeed distorted or lost as the HMD is moved, with consecutive frames showing visually different glyphs. This is most noticeable with the ‘i’ character losing distinction between the diacritic dot (tittle) and the stem of the glyph, making it appear like the adjacent ‘l’. The ascender of the ‘h’ glyph is also diminished in several frames making it appear as an ‘n’, and the counter within the ‘e’ is not clear in several frames allowing it to be confused for a ‘c’.

The culprit for this character distortion is the texture filtering mechanism used within the virtual environment, nearest-neighbour interpolation, which determines the pixel value displayed on the physical device when there is a miss-match between texels and available pixels. This filtering can select non-ideal texel values for display, leading to inconsistency of appearance between frames, and potential distortion.

Figure 6.4: The word ‘while’ as seen over 9 frames through the lenses of the DK2.

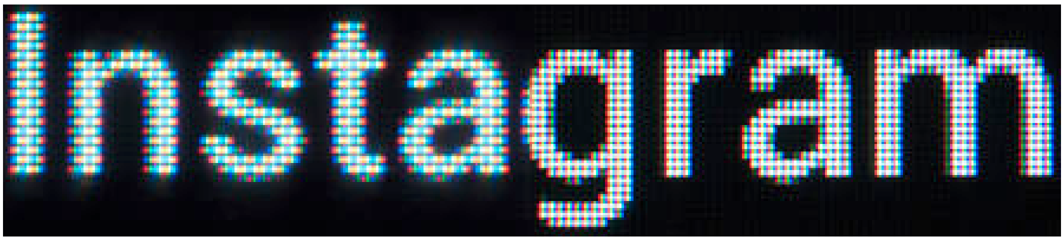

Figure 6.5: Magnified view of the word ‘Instagram’ shown on RGB and RGBG subpixel arrangement displays (Rakver, 2020). Left side demonstrates RGBG, right side demonstrates RGB.

An approach to resolve this that was suggested by several participants was the use of an alternative font for body text that would be more tolerant to being displayed in the virtual environment. Many participants suggested a serifed font, such as “Georgia” used in section headings within the read texts, would provide better results than the “Arial” body text. This suggestion can possibly be attributed to the presence of different fonts within the document allowing comparison, with the identified font being identified as a possible improvement based on the appearance within the environment. This comparison is not necessarily valid due to a difference in font-size with elements displayed with either font, however it does warrant further investigation to identify if a difference in readable distances or reading performance is seen when using a different font.

A final observation made by several participants related to selection of the DK1 and the preferred HMD over the DK2, with clarity of text noted as an advantage. A possible candidate for the cause for this observation is the differing displays used within the HMDs, with the DK1 using an RGB LCD display and the DK2 using an RGBG organic light-emitting diode (OLED) display. A visualisation of the difference in text appearance between these display types can be seen in Figure 6.5. As this was only noted by a minority of participants, the impact of this has not been judged to be severe, however given that newer headsets such as the Oculus Rift CV1 (CV1) continue to make use of RGBG OLED displays the impact of this should be considered going forward.

6.7.6 Text Haloing

An observation made by nearly all participants was that text displayed within the virtual environment appeared to "glow", with some describing the effect as "a colourful halo" around displayed text. It was observed that these halos changed appearance as participants read, with a "shimmering" effect described as the colour of the halo changed. The colour of these halos was not consistent, however primary colours were noted as the most prominent, with blended colours being less apparent.

There are several candidates for the cause of this effect, each of which may be solely or partly responsible. The first of these is chromatic aberration caused by the biconvex lenses used with the tested HMDs. This aberration is caused by flaws in the lenses which causes light colours to not converge correctly to a point, resulting in a prismatic effect with red, green, and blue light channels separated slightly. While this effect aligns with the observed halo colours noted, it would be expected that if this were the cause that the aberration would be seen throughout the displayed image and not isolated to the text. As the study focused on text primarily, it is possible that participants did not inspect other page elements within the virtual windows to see if this effect was present, which may account for this.

A potential contributing factor to observed chromatic aberration is a mismatch between participant and HMD IPDs, which is noted by McCleary (2009) as a potential cause of prismatic effects and double vision. Both the DK1 and DK2 have fixed lenses with a constant IPD of 63.5mm, relying on software correction within the Oculus runtime to account for the difference between this and the actual user IPD. This correction is accomplished by providing the user IPD at runtime, with a utility provided to calculate user IPD if it is not known. This utility is imprecise however, which may lead to continuing mis-

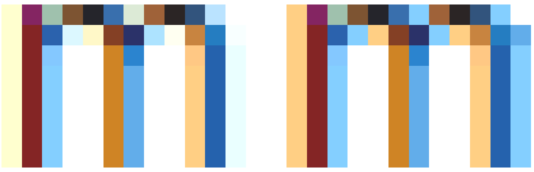

Figure 6.6: Letter ‘m’ extracted from a sample image (left), zoomed to display the ClearType colour fringing. The same ‘m’ with faint colours amplified is displayed on the right.

match in IPD values, possibly amplifying negative effects. Future experiments should consider an external tool such as a digital pupilometer to calculate participant IPDs with precision rather than relying on a software utility if fixed lens positions continue to be standard.

The second and more promising candidate is the effect of the sub-pixel anti-aliasing technique employed by Microsoft ClearType. This font-hinting technique is employed to improve the visual qualities of text displayed on RGB LCD displays (Microsoft Corporation, 2001), with studies (Sheedy et al., 2008) identifying that the effect is preferred by readers though no improvement in reading comfort or speed was observed as a result of it. This technique functions by adding additional coloured pixels alongside character glyphs and replacing solid pixels making up these glyphs with coloured shades to take advantage of the subpixel layout within an RGB LCD display, with this effect applied directly to the text and thus captured when collecting samples for display. When magnified, the effect of this font-hinting can be seen clearly as seen in Figure 6.6.

In this image, the colours utilised by the ClearType effect do appear to match with the observed halo colours, suggesting this is a likely contributor to the effect. While participants did make observations about this halo effect, it was not clear if the effect negatively impacted their reading performance. Future investigations should consider comparing readability of text with ClearType, or other font-hinting techniques, enabled or disabled to determine if the effect has an impact on readability.

6.8 Summary

In this chapter, reading performance within a developed immersive virtual environment was compared against baseline performance on conventional displays when reading from unmodified conventional application windows. It was shown that reading performance equal to baseline is achievable for narrower windows displayed on either curved or planar surfaces, with equal reading accuracy achievable on wider windows displayed on curved surfaces. A range of acceptable reading distances were identified, and issues with reading from wide planar windows were demonstrated. It was demonstrated that the use of an IVR HMD with greater resolution and pixel density improved reading performance on wider windows and allowed for more distant positioning of windows while retaining reading performance. Curved windows were shown to be preferred by most study participants for reading, and a minority preference for planar panels was noted.

Chapter 7

Tuning Text

In this chapter, a user study is presented that measures the effect of modifying text parameters within conventional applications on the readable distances of text samples positioned within the virtual environment described in Chapter 5. Font, ClearType, and interface zoom are tested to examine their impact on readable distance, identifying attributes that can be modified within many conventional applications that can improve readability. This seeks to address Research Question 4, which stated:

> *With respect to [reading], what software configuration changes can be easily made to conventional applications that may improve performance within a virtual environment?*

It is demonstrated that font choice did not affect readable distances of text within the virtual environment, while both ClearType and interface zoom are shown to have statistically significant impacts. Device improvements between this and the previous study are shown to reflect greater readable distances within the virtual environment, though the degree of improvement is shown to be greater than what the observed improvement in pixel density would suggest. Factors that may have contributed to these observed results are discussed and identified for control in future evaluations.

Supporting material for this chapter can be found in Appendix D.

7.1 Questions

In Chapter 6, it was demonstrated that reading performance within the created virtual environment was capable of matching that of reading from a conventional display. It was also identified that the distance at which text was considered readable varied between participants, with a spread in results that was greater than what was expected. To determine if improvements could be made to the readability of text displayed on conventional interfaces within the developed virtual environment, study participant feedback and post-study analysis was evaluated to identify possible modifications that could be made within conventional applications or the host operating system that would address concerns raised.

When evaluating study participant feedback, two aspects were identified as affecting the chosen distance at which samples were positioned prior to reading. The first of these was “text haloing” which was observed to affect the appearance of text in a dynamic manner, and the second was the choice of fonts used within the presented samples which some considered to be non-ideal for the task. During post-test analysis, two further items were identified that may have affected the sample distances chosen. The first of these was the choice of allowing participants to freely choose the distance prior to reading, as natural preferences for reading distance are likely to vary and not identify actual limits on readability. The second was that text appeared to distort when viewed through the tested head-mounted displays (HMDs) when the view position was changed, which was identified as a possible texture filtering issue when mapping from texels to pixels.

7.1.1 Text Haloing

When observing text samples from within the tested HMDs, participants noted a shimmering effect around text that changed as content was read. While some participants indicated that they found the colour “haloing” effects unpleasant, it was not clear if this impacted the legibility of the text or simply added an element of distraction. To determine if this does impact on performance, factors contributing to this effect should be identified and addressed where possible.

The previous chapter identified three possible contributors to this effect — chromatic aberration caused by imperfections in the HMD lenses, a miss-match between the perceiver’s inter-pupillary distance (IPD) and that configured on the HMD, and the sub-pixel anti-aliasing feature ClearType which was enabled on the host operating system when preparing samples for the study. During informal evaluation, disabling ClearType on the host operating system was observed to almost eliminate the effect, suggesting that this is the greatest contributor to the effect and should be evaluated primarily.

On conventional displays, the use of ClearType has been observed to have benefits for lexical decision and sentence comprehension tasks (Gugerty et al., 2004), however no reading performance benefits were observed when it was in use. Sheedy et al. (2008) supported this conclusions but did note that users

Table 7.1: Fonts used in the readability study conducted by Ali et al. (2013), grouped by the media they were designed for, and font family.

Designed for	Serif	Sans-serif
Print	Times New Roman	Arial
Screen	Georgia	Verdana

preferred moderate levels of anti-aliasing (ClearType levels 1 & 2) to no anti-aliasing. While this does suggest that disabling ClearType should not have a negative impact on reading performance on a conventional display, this does not necessarily translate into the performance within a virtual environment.

7.1.2 Fonts

Several participants of the previous study questioned the choice of the fonts used in the text samples. Some users stated that they disliked the font, with a subset of these believing this dislike affected their decisions in the experiment. It was suggested by these participants that by changing the font from the sample sans-serif Arial font to a different sans-serif font, or to a serifed font, they would find the text more readable. It was not clear whether this suggestion of increased readability would relate to speed and accuracy of reading at their chosen sample distances, or if such a change may allow for the samples to be read at greater distances.

Using conventional displays, Ali et al. (2013) conducted a readability experiment that assessed the readability of 4 fonts shown in Table 7.1. These fonts were chosen as examples of serif and sans-serif fonts designed for print and screen media. This study concluded that there “was no significant difference in text readability on the computer screen between serif and sans-serif fonts” (Ali et al., 2013) within each media category, but did observe a difference in readability scores between fonts designed for screen and for print purposes, with the fonts designed for screen exhibiting a greater readability score than those designed for print.

As this study showed that the choice of screen or print fonts influenced readability on conventional displays, an investigation to determine if this is similarly the case with samples displayed in the virtual environment would be valuable. While the above study suggests that the choice of serif or sans-serif fonts will not affect readability of text samples, the preferences of study participants may influence their willingness to award greater readability scores if the scoring is taken as a user-provided value. Investigating a change in fonts considering these factors would help to determine if simply changing the font used in text samples improves the effective readable distance at which samples can be placed.

7.1.3 Engine Distortions

During development of the virtual environment, different texture filtering mechanisms were not investigated as the DirectX default of nearest-neighbour was found to be predictable and initially sufficient. During the previous experiment, the appearance of text was observed to change as the orientation of the HMD was modified, often producing frames where identifying characteristics of font glyphs were lost. It was suggested that a possible method of addressing this issue was to make use of higher density textures, displayed on a virtual window sized for a lower density texture. This was identified as possibly providing a greater number of information-bearing texels to the filter, allowing for a better-quality displayed result.

A simple approach proposed to test the effectiveness of this was to capture application windows that had been resized to twice the usual horizontal and vertical size, and to present the captured texture on a virtual window sized and positioned for the normal size. Content displayed within the displayed window would appear scaled down with this presentation if left unmodified, however resizing the font to twice the normal size or zooming the interface to 200% was found to counteract this effect, which is achievable in many different conventional applications. During initial investigation with this technique, a difference in appearance of text was noted, however the impact on this change on the readability of the text was unclear. Measuring any change in readability of text is then necessary to determine if simple zooming and scaling of content is a viable approach to improving the quality of displayed text.

7.1.4 Readable Distances

The previous study sought in part to identify and optimal distance at which text could be displayed, which could then be used as a basis for further development efforts. Rather than identifying a single distance agreed upon by all participants, the study showed there was a range of tolerable reading distances within the virtual environment, with no “break points” in distances where text moved from readable to unreadable able to be accurately determined. While this is not unexpected, as individual participant preferences and tolerances are likely to vary, it was not clear if this was solely down to preference and tolerance, or due to a poor instruction and design around the process for positioning windows.

At the beginning of the study process, participants were instructed to place the text samples at the furthest distance from themselves where they could still comfortably read the text. It was demonstrated to participants that the distance could be adjusted back and forth with some granularity, and they were encouraged to spend some time identifying the greatest readable distance for a sample before committing to the position. While participants were observed to move the sample position forward and back prior to reading the content — often at speed — it is not clear if each position was fairly considered for reading comfort. It is therefore possible that greater distances at which text could be comfortably read were not identified by some participants which may have resulted in the wider than expected range of readable distances identified.

To address this, presenting samples at a range of fixed distances from participants and asking them directly if the text was comfortably readable could be considered. This approach would require participants to consider each position and may result in observations of the expected "break points" where text moves from readable to unreadable cleanly for each participant. If this is observed, these breakpoints could be compared amongst participants to determine if these are objective values seen consistently with all participants, or if a subjective tolerance to text appearance plays a role in selection.

7.1.5 Device Improvement

When conducting the previous experiment, two different immersive virtual reality (IVR) HMDs were utilised, and the results compared. It was observed that the Oculus Rift DK2 (DK2) device allowed participants to position text samples at a greater distance while remaining comfortably readable than was seen on the Oculus Rift DK1 (DK1) device. While this was expected due to display improvements seen between the two devices, it could not be extrapolated to predict readable distances on future devices due to a lack of data points. During the process of conducting the previous study, a new device in the same Oculus family, the Oculus Rift CV1 (CV1), was released which further improved the display, tracking, and optics used. If collected, the readable distances of text using this new device could be compared with the previously collected data to establish if improvement continues to be seen with a newer device, and if any distance improvement is proportional to display improvements seen between devices.

7.2 Hypotheses

From the above discussion, five hypotheses can be made, with three of these grouped into the category of potential modifications that could be made to conventional applications that would result in greater readable distances for text. For the categorised hypotheses, evaluating the impact of font, ClearType, and texture density on readable distances yields:

Hypothesis 7.1. *Using a font designed for screen-based media will result in text samples being considered readable at greater distances than for samples using a font designed for reading on paper.*

Hypothesis 7.2. *Enabling and disabling ClearType for samples will result in variation in acceptable distances at which text samples can be placed and considered readable.*

Hypothesis 7.3. *Increasing the information density of textures will improve the appearance of text samples, allowing them to be considered readable at greater distances than those using the current density of textures.*

For evaluating readable distances more precisely:

Hypothesis 7.4. *When presented with fixed locations for evaluation, the variation in distances where text is considered comfortably readable will be narrower than what was seen in the previous study where free selection of distance was given.*

And finally, determining the improvement on readable distance seen by moving to a more modern HMD:

Hypothesis 7.5. *The upper limit of where text is considered readable using the CV1 will show an improvement over the limit found with the DK2 proportional to the improvement in display pixel density between the two devices.*

7.3 Experimental Design

Using the experimental harness described in Section 5.5.4, an experimental setup for the generation 2 virtual environment was created that allowed for text samples to be displayed within the environment, with accompanying response controls allowing participants to give ratings for the readability of the displayed sample. Text samples could be presented within a fixed range of distances chosen by the researcher, that spanned from a point where the researcher believed that text could be comfortably read, through a point beyond where

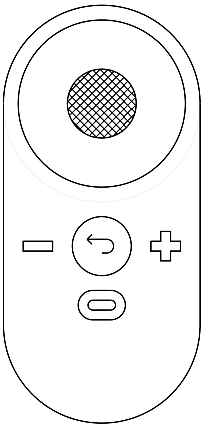

Figure 7.1: Oculus Rift remote with the centre button indicated.

the researcher considered the text to be unreadable. Three rating responses were provided alongside each text sample, allowing participants to indicate that the text was clearly readable, readable with effort, or unreadable at the displayed position.

Within the environment, captured text samples were rendered on a circumferential virtual window of a fixed size corresponding to a 610 × 458 pixel sample on a 38ppcm display. This virtual window was positioned directly ahead of the participant, positioned at specified distances from the environment origin, with the radius of the circumferential panel updated to match this distance. The gaze pointer for the environment was enabled and shown ahead of the user, following the direction of the headset.

Below the sample at a fixed distance, rotated $-25°$ below level gaze, three pictorial response button panels were shown in a row, visualised in Figure 7.2. Interaction with the response buttons was accomplished by positioning the gaze pointer over the desired button, then pressing the centre button of the Oculus remote peripheral (See Figure 7.1).

62 text samples created using text of the short stories found in 'GrimmFairy Tales', obtained from Project Gutenberg[1]. Each of these short stories were processed into simple HTML documents, consisting of a single `<div>` element containing multiple `<p>` elements. The text for each short story was placed inside this structure, with each sentence of the story placed inside separate `<p>` elements. A common CSS stylesheet was applied to each document to give consistent results. The stylesheet used can be seen in Listing 7.1.

---

[1] `https://www.gutenberg.org/ebooks/2591/`

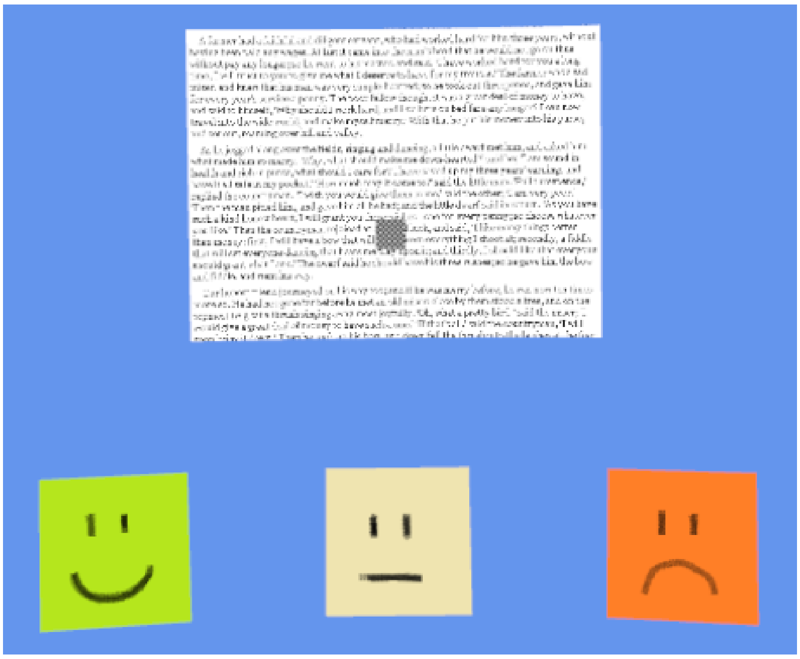

Figure 7.2: Experimental text sample shown with response buttons. Gaze pointer is represented by a textured square over the sample.

Text samples were rendered using Mozilla Firefox running on a Microsoft Windows 7 computer, then captured for later use by taking a screenshot of the `<div>` and contents using the screenshot tool built into Mozilla Firefox. The screenshot tool was instructed to capture only the `<div>` content, with the document background excluded. Screenshots were saved in folders based on the permutation of options used for the sample, with sequential numbering as their filenames.

Each sample was captured 8 times with different permutations of options. The list of options and their states can be seen in Table 7.2. When modifying the font for samples, the CSS attribute `font-family` — marked with underscores in Listing 7.1 — was set to “Arial” or “Georgia” as appropriate. When modifying the zoom factor for a sample, the Mozilla Firefox browser zoom setting was configured to 100%, or 200% prior to capturing the sample. When

Table 7.2: Attributes under test, showing their possible states.

Attribute	A	B
ClearType	Enabled	Disabled
Font	12pt Arial	12pt Georgia
Scale	100%	200%

```
p {
    text-indent: 1em;
    margin: 0.5em 0;
    color: #222222;
    font-family: _____; /* Swapped as necessary*/
    font-size: 12pt;
}

div {
    width: 600px;
    height: 450px;
    overflow: hidden;
    text-overflow: clip;
    background-color: white;
    padding: 5px 0.5em;
}

body {
    padding: 2em;
    background-color: black;
}
```

Listing 7.1: CSS stylesheet used when rendering text samples within Mozilla Firefox. `font-family` is marked and was substituted with “Arial” or “Georgia” as appropriate. Note that the `width` and `height` properties specify the content area dimensions of the `div`. When captured, the resulting file dimensions included the `padding` dimensions as specified.

modifying the ClearType setting, ClearType was either disabled, or set to the default state by running the “Adjust ClearType text” configuration wizard, and selecting options $\begin{pmatrix} 1 & 0 & 0 \end{pmatrix}$, $\begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix}$, $\begin{pmatrix} 1 & 0 & 0 \end{pmatrix}$, and $\begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix}$ from the sequence of visual options presented.

Examples of rendered text samples can be seen in Figures D.1 to D.2. An enlarged example of a single word rendered in all 8 permutations can be seen in Figure 7.3, with the effect on the appearance of text caused by each visible.

The presentation order of captured text samples and the distance at which they were displayed was randomised between experiments. The complete range of display distances under test were shown for a permutation of test attributes prior to moving to the next permutation set. The order of permutation sets was also randomised between participants. The logic for this expressed as pseudocode can be seen in Listing 7.2.

Results for the experiment were recorded in a log file for each participant, an example of which can be seen in Listing 7.3.

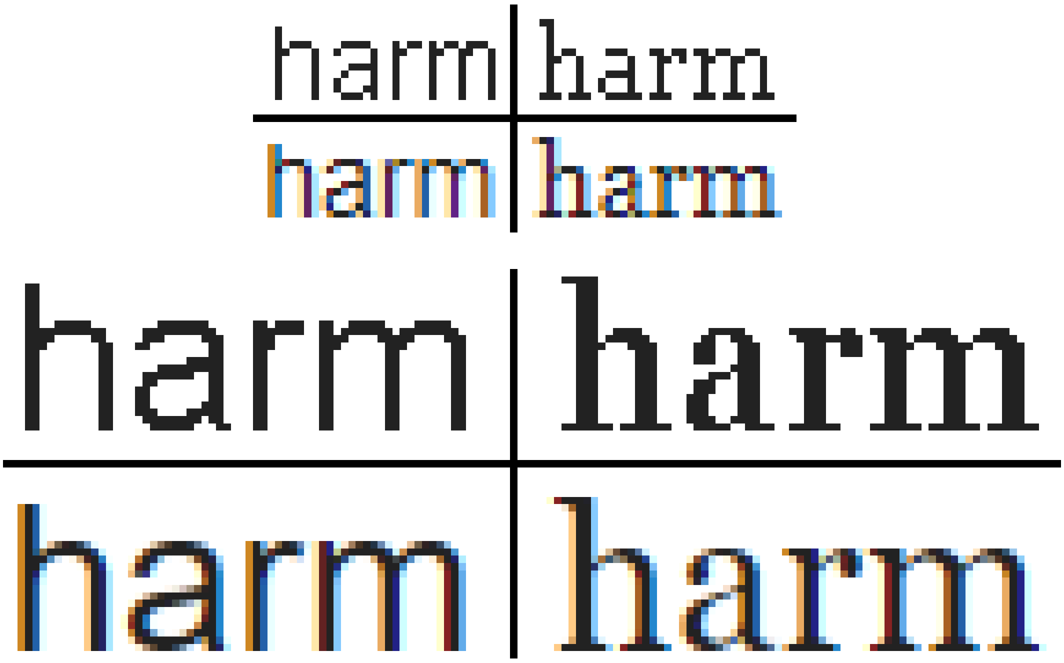

Figure 7.3: Magnified view of a single word within a text sample, rendered with each of the permutations of tested attributes as two grids. Within each grid, Arial font is displayed in the left column, and Georgia font in the right column. The top row of each grid is displayed with ClearType disabled, and the bottom with ClearType enabled with default settings. The top-most grid is shown at 100% zoom, and the bottom at 200%.

```
sets := {A, B, C, D, E, F, G, H}.Shuffle()

for tSet in sets
  distances := Range(MinDistance, MaxDistance, 0.1).Shuffle()
  images    := tSet.images.Shuffle()

  for tDistance in distances
    tImage   := images.Pop()
    response := Present(tImage, tDistance)
    RecordResponse(tSet, tImage, tDistance, response)
```

Listing 7.2: Pseudocode showing the experimental process of selecting sample sets, distances, and images for display. Sets A–H represent each permutation of independent variable combinations. `MaxDistance` and `MinDistance` represent the maximum and minimum render distances under test, configurable by the experimental supervisor. Distances are range inclusive, stepping by 0.1 in-engine units (IEUs) per step.

```
Experiment begins at UTC 2021-01-01 00:07:32
Experiment ends at UTC 2021-01-01 00:07:45
Participant ID: 0f1h7x
"Permutation","Order","Image","Distance","Rating"
"ARIAL_100_OFF",1,"part-002",2.0,2
"ARIAL_100_OFF",24,"part-049.png",3.1,1
```

Listing 7.3: Excerpt of an example experimental run results file, showing the start and end times of the experiment, the ID of the participant, and CSV data containing participant responses. Response rows indicate the permutation of attributes that is under test, the order in which a given sample was displayed, which captured text sample was used, the distance at which the sample was displayed in the environment, and the rating given by the participant.

7.3.1 Equipment

The experimental software and the accompanying harness application were run on a computer installed with the Microsoft Windows 7 operating system. This computer used an Intel Core i7-5930K processor, with 16GB of system memory installed, and an NVidia GTX1080 graphics card with 8GB of dedicated video memory. The IVR HMD used was the CV1 with 3 constellation sensors present. The constellation sensors were positioned at the 10 o'clock, 2 o'clock, and 6 o'clock positions with respect to the participant's seated position. An Oculus remote was used for selection input by the participant and was provided by the researcher at commencement of the study.

A secondary 24-inch monitor was connected to the system, which displayed a monocular mirror of the participant's view for the researcher to observe.

7.3.2 Method

On arrival, participants were introduced to the CV1 HMD and the Oculus remote control, and an overview of the experimental process was given. Participants were then seated at a desk with their chair adjusted to a comfortable height, and the CV1 was positioned on, and adjusted to comfortably fit, their head. The Oculus remote control was passed to the participant's dominant hand, and their thumb was placed over the centre button. Participants were prompted to select menu items within the Oculus home environment using the gaze pointer and the Oculus remote to select IPD configuration settings. They were prompted to adjust the headset vertical position, and the IPD slider of the CV1 until the test image presented was clear.

With use of the gaze pointer and Oculus remote for selection established, the environment via the configured testing harness was launched, and the first sample displayed. Participants were instructed to examine the displayed sample, taking time to read a portion of the text, then decide if they considered the sample easily readable, readable with effort, or unreadable. They were prompted to select the appropriate response button using the gaze pointer and Oculus remote, after which they would be given the next sample. Any unprompted observations made by participants were noted by the researcher, but no questions were directly asked of participants during this process.

Participants were instructed to repeat this process until no further samples were shown, at which point the CV1 was removed from their head. Participants were then asked to remain seated for a few minutes prior to leaving to ensure that they were not disoriented from the experience.

7.4 Participants

Participants for this study were sourced from the University of Waikato student population, with the majority being from the Faculty of Computing and Mathematical Sciences undergraduate student body. Potential participants were introduced to the study through presentations and written descriptions and were invited to contact the researcher if they were interested in taking part. A subset of participants engaged with this study due to a class requirement that they take part in a user study. The researcher was not involved with this paper, did not report attendance or results to the paper convener, and did not have influence on grades received for this paper.

22 participants were selected from the pool of volunteers, with preference given to those who did not require spectacles to correct vision. Participants that did require spectacles were accepted if they would not need to remove them to use the CV1. No preference was given to participant gender, age, or ethnicity during selection.

Of the selected participants, 9 identified as male, 12 as female, and 1 as neither of these two choices. 16 participants reported no vision issues, 3 reported corrected vision through the use of contact lenses or spectacles compatible with the headset, and 3 were classified as having uncorrected vision issues affecting only a single eye. Participants were assigned a randomised ID number during the experiment, with demographic and experiment data recorded against it. During data processing, participant IDs were simplified to sequential IDs $0 \leqslant ID \leqslant 21$ for ease of reference. This is visualised in Figure 7.4.

Table 7.3: Participant information summarising the number of samples observed and the total time taken between the start and end of the experiment. $\mathcal{T}_s$is calculated as the $\mathcal{T}$over the number of samples, providing a rough average time spent per sample in seconds.

ID	**Samples**	$\mathcal{T}$	$\mathcal{T}_s$
0	168	1209	7.20
1	168	630	3.75
2	248	1251	5.04
3	248	986	3.98
4	248	808	3.26
5	248	1475	5.95
6	248	1519	6.13
7	96	218	2.27
8	248	1123	4.53
9	248	895	3.61
10	248	895	3.61
11	248	632	2.55
12	248	819	3.30
13	248	1855	7.48
14	248	568	2.29
15	248	730	2.94
16	248	1585	6.39
17	248	1620	6.53
18	248	1133	4.57
19	248	1406	5.67
20	248	1922	7.75
21	96	854	8.90

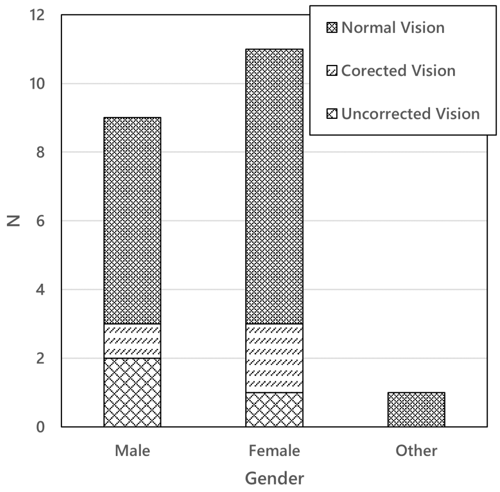

Figure 7.4: Study participant demographic information showing numbers of participants with vision issues.

Of the selected participants, 18 took part in the full-length experiment consisting of 248 data samples. 2 participants undertook a shorter 168 sample experiment, and the remaining 2 participants 96 samples. This variation is due to unexpected participant time constraints that required a shorter experiment time. Mappings between participant IDs and experiment sample length is shown in Table 7.3.

7.5 Data Processing

Raw data was collected from logs and processed into tables for each permutation of attributes (See Tables D.2 to D.9). Participant responses were recorded as 0, 1, 2 for the responses of "unreadable", "readable with effort", and "readable" respectively. Data was further processed to record the distance at which the last "readable" response was given ($\mathbb{D}_n$) before other responses were indicated, and the first "unreadable" result ($\mathbb{D}_f$) was produced for each participant and permutation of attributes. Participants that did not provide any "readable" rating for a given permutation had no $\mathbb{D}_n$ value recorded. A summary of $\mathbb{D}_n$ values for participants can be found in Table D.10, and $\mathbb{D}_f$ values in Table D.11.

Duration of experiment ($\mathcal{T}$) was recorded by taking the experiment start and end timestamps from the produced logs and calculating the difference in seconds. The average duration per sample seen ($\mathcal{T}_s$) was calculated as $\frac{\mathcal{T}}{\mathbb{N}_s}$. This data can be seen in Table 7.3.

7.6 Results

This experiment aimed to address hypotheses surrounding impact of changing independent attributes of text samples on the distance at which they were considered readable, as well as hypotheses around readable distance values given by participants. In this section, each hypothesis will be addressed in turn.

7.6.1 Font

To address Hypothesis 7.1, which stated:

> *Using a font designed for screen-based media will result in text samples being considered readable at greater distances than for samples using a font designed for reading on paper.*

the $\mathbb{D}_n$ and $\mathbb{D}_f$ values for samples using the Arial font were compared against those using the Georgia font. A paired t-test was run across these values, testing for:

$H_0 : \mu\left(\mathbb{D}_{Arial}\right) - \mu\left(\mathbb{D}_{Georgia}\right) = 0$

$H_1 : \mu\left(\mathbb{D}_{Arial}\right) - \mu\left(\mathbb{D}_{Georgia}\right) \neq 0$

The results of this test can be seen in Table 7.4.

Table 7.4: Two-way paired t-test for the $\mathbb{D}_n$ and $\mathbb{D}_f$ points, with font acting as the independent variable.

Cutoff	$\bar{x} \pm s$ (Arial)	$\bar{x} \pm s$ (Georgia)	$\bar{x} \pm s$ (Diff)	95% CI	t	p
$\mathbb{D}_n$	2.6963 ± 0.3656	2.6854 ± 0.3542	0.0110 ± 0.2320	$(-0.0400, 0.0620)$	$t(82) = 0.43$	0.670
$\mathbb{D}_f$	3.2651 ± 0.4104	3.2723 ± 0.4462	-0.0072 ± 0.2668	$(-0.0655, 0.0510)$	$t(83) = -0.25$	0.806

From the results of this test, the difference between both $\mathbb{D}_n$ and $\mathbb{D}_f$ values for these samples was not statistically significant ($p > 0.05$), suggesting that using 12pt Georgia in place of 12pt Arial fonts for text samples will not produce a difference in maximum readable distance within this virtual environment.

7.6.2 ClearType

To address Hypothesis 7.2, which stated:

> *Enabling and disabling ClearType for samples will result in variation in acceptable distances at which text samples can be placed and considered readable.*

the $\mathbb{D}_n$ and $\mathbb{D}_f$ values for samples with ClearType enabled were compared against those with ClearType disabled. A paired t-test was run across these values, testing for:

$H_0 : \mu\left(\mathbb{D}_{On}\right) - \mu\left(\mathbb{D}_{Off}\right) = 0$

$H_1 : \mu\left(\mathbb{D}_{On}\right) - \mu\left(\mathbb{D}_{Off}\right) \neq 0$

The results of which can be seen in Table 7.5.

Table 7.5: Two-way paired t-test for the $\mathbb{D}_n$ and $\mathbb{D}_f$ points, with ClearType status acting as the independent variable.

Cutoff	$\bar{x} \pm s$ (On)	$\bar{x} \pm s$ (Off)	$\bar{x} \pm s$ (Diff)	95% CI	t	p
$\mathbb{D}_n$	2.7205 ± 0.3987	2.6458 ± 0.3206	0.0747 ± 0.2459	$(0.0210, 0.1284)$	$t(83) = 2.77$	0.007
$\mathbb{D}_f$	3.2928 ± 0.4507	3.2410 ± 0.4076	0.0518 ± 0.2461	$(-0.0019, 0.1055)$	$t(83) = 1.92$	0.059

From the results of this test, $\mathbb{D}_n$ values show that toggling ClearType did have a statistically significant difference in readable distances ($p < 0.05$). Samples with ClearType enabled showed a greater mean readable distance ($2.72 > 2.65$) when compared to samples with this disabled.

For $\mathbb{D}_f$, the result is on the cusp of significance, with $p = 0.059$. While this on its own is not enough to state that the effect of ClearType has an effect of readable distances when using a confidence interval of 95% , when accompanying the statistically significant $\mathbb{D}_n$ result it does suggest that there may be an impact that could be significant if evaluated with a larger sample size.

These results suggest that the use of ClearType can have a positive impact of readable distances within this configuration of virtual environment. Further experimentation with ClearType, alternative font hinting methods, and larger numbers of participants would be required to identify if this effect is specific to ClearType, or if this can be generalised to other sub-pixel hinting approaches. Further study may also identify specific setting configurations for ClearType that could provide a stronger effect.

7.6.3 Zoom

To address Hypothesis 7.3, which stated:

> *Increasing the information density of textures will improve the appearance of text samples, allowing them to be considered readable at greater distances than those using the current density of textures.*

the $\mathbb{D}_n$ and $\mathbb{D}_f$ values for samples with zoom set to 100% enabled were compared against those with zoom set to 200%. A paired t-test was run across these values, testing for:

$H_0 : \mu\left(\mathbb{D}_{100\%}\right) - \mu\left(\mathbb{D}_{200\%}\right) = 0$

$H_1 : \mu\left(\mathbb{D}_{100\%}\right) - \mu\left(\mathbb{D}_{200\%}\right) \neq 0$

The results of which can be seen in Table 7.6.

Table 7.6: Two-way paired t-test for the $\mathbb{D}_n$ and $\mathbb{D}_f$ points, with zoom acting as the independent variable.

Cutoff	$\bar{x} \pm s$ **(100%)**	$\bar{x} \pm s$ **(200%)**	$\bar{x} \pm s$ **(Diff)**	**95% CI**	**t**	**p**
$\mathbb{D}_n$	2.7185 ± 0.3889	2.6407 ± 0.3240	0.0778 ± 0.2269	$(0.0276, 0.1280)$	$t(81) = 3.08$	0.003
$\mathbb{D}_f$	3.3024 ± 0.4353	3.2313 ± 0.4225	0.0711 ± 0.2564	$(0.0151, 0.1271)$	$t(81) = 2.53$	0.013

From the results of this test, the difference between both $\mathbb{D}_n$ and $\mathbb{D}_f$ values for these samples was statistically significant ($p < 0.05$). For $\mathbb{D}_n$ the zoom setting of 100% produced a greater mean readable distance than with the 200% setting ($2.7185 > 2.6407$), which is also seen for the $\mathbb{D}_f$ mean readable distance ($3.3024 > 3.2313$).

These results indicate for this virtual environment configuration, simply zooming text content to increase the resolution of displayed textures will not result in an increased readable range, instead harming it.

7.6.4 Fixed vs Free Distances

To address Hypothesis 7.4, which stated:

> *When presented with fixed locations for evaluation, the variation in distances where text is considered comfortably readable will be narrower than what was seen in the previous study where free selection of distance was given.*

a subset of data was extracted from the greater collected dataset for comparison. When comparing to locations identified in the study described in Chapter 6, it should be considered that the data collected in that study was gathered based on text samples captured at 12pt Arial font with ClearType enabled. Thus, when comparing those samples to the values collected in this study, only values resulting from evaluation of samples with the same constraints should be used. Additionally, the previous study measured distance for three different sizes of virtual windows, on two different distortions was collected. As each of these variables resulted in different placements, only the values collected from samples with the same distortion and approximate size should be used. In this scenario, samples from the previous study using the 'Small' panels distorted with the circumferential distortion will be used for comparison.

To evaluate the variation in readable distances, the coefficient of variation was calculated for each of the concerned samples. The results of this evaluation can be seen in Table 7.7.

From this evaluation for the samples compared, samples given at fixed positions rather than allowing free placement resulted in smaller coefficients of variation. This indicated that the variation in mean distance selected for the CV1 samples was smaller than that seen with the older devices. This suggests that requiring participants to consider samples positioned at fixed

Table 7.7: Coefficient of variation values for identified comfortable readable distance, displaying the size of the sample distance was evaluated for in pixels, and the zoom at which the sample texture was captured. Entries concern samples captured using a 12pt Arial font, with ClearType enabled.

Device	Sample Dimensions	Zoom	CV
Oculus Rift DK1	800 × 600	100%	17.13%
Oculus Rift DK2	800 × 600	100%	20.26%
Oculus Rift CV1	610 × 458	100%	15.12%
Oculus Rift CV1	610 × 458	200%	12.87%

locations within a region of interest will yield more precise results than when allowing participants to freely position samples within the same region. Free positioning could still be a useful tool for evaluating preferred locations within a range of acceptable locations for individuals.

7.6.5 Device Improvement

To address Hypothesis 7.5, which stated:

> *The upper limit of where text is considered readable using the CV1 will show an improvement over the limit found with the DK2 proportional to the improvement in display pixel density between the two devices.*

$\mathbb{D}_n$ values were collected from the collected results, and compared against those collected from the previous study and the pixel density improvements between tested devices. As noted above, when making comparison to data collected in the study described in Chapter 6, only samples with comparable attributes should be used. To support this, $\mathbb{D}_n$ values from this study having the attributes of Arial font, 100% zoom, and ClearType On, were compared against the samples in the previous study that utilised curved mesh distortion virtual windows. Data from the previous study was grouped by device and virtual windows size. This data can be seen in Table 7.8 and visualised in Figure 7.5.

From this data, the $\mathbb{D}_n$ values obtained from this study using the CV1 HMD were significantly higher than values from both the DK1 and DK2 used in the previous study. When comparing the small sample values from the previous study and the values collected in this study, we can see that the DK2 offered a median distance improvement of 38.4% over the DK1, and the CV1 a 90.6% improvement over the DK2.

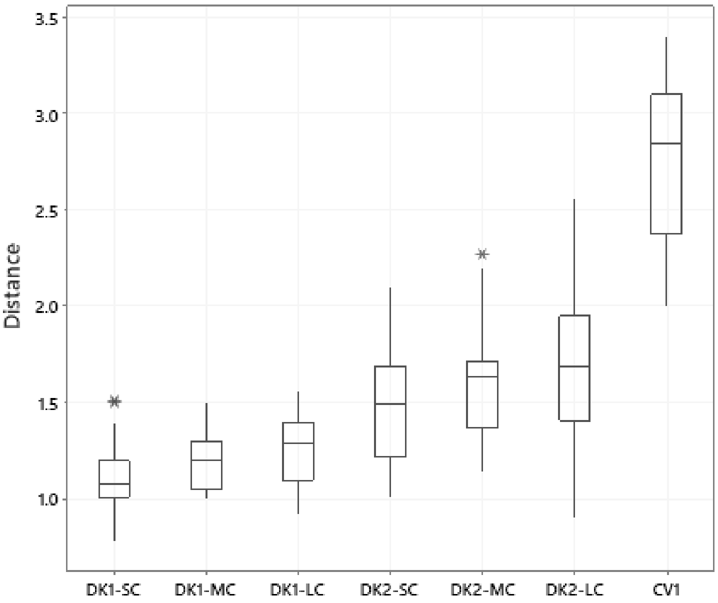

Figure 7.5: Visualisation of readable distance values on curved virtual windows, collected from the study described in Chapter 6, show alongside the $\mathbb{D}_n$ values collected from this study with the attribute permutation of Arial font, 100% zoom, and ClearType on.

Table 7.8: Descriptive statistics for the readable distance values on curved virtual windows, collected from the study described in Chapter 6, show alongside the $\mathbb{D}_n$ values collected from this study with the attribute permutation of Arial font, 100% zoom, and ClearType on.

Device	Size	N	$\bar{x} \pm s$ ($\mathbb{D}_n$)	Q0	Q1	Q2	Q3	Q4
DK1	Small	23	1.1122 ± 0.1905	0.7800	1.0100	1.0800	1.2000	1.5100
DK1	Medium	23	1.2078 ± 0.1572	1.0000	1.0500	1.2000	1.3000	1.5000
DK1	Large	23	1.2544 ± 0.1840	0.9200	1.1000	1.2900	1.4000	1.5600
DK2	Small	24	1.4792 ± 0.2998	1.0100	1.2200	1.4950	1.6875	2.1000
DK2	Medium	24	1.6029 ± 0.2774	1.1400	1.3700	1.6350	1.7175	2.2700
DK2	Large	24	1.6879 ± 0.3853	0.9000	1.4075	1.6900	1.9500	2.5600
CV1	-	22	2.7591 ± 0.4171	2.0000	2.3750	2.8500	3.1000	3.4000

Table 7.9: Summary of tested device per-eye display resolutions, given with the horizontal field of view (FOV) of the device and calculated horizontal PPD. Improvement is given with respect to the previous device PPD value.

Device	Resolution	FOV_h	PPD_h	Improvement	Layout
DK1	640 × 800	90°	7.11	-	RGB
DK2	960 × 1080	94°	10.21	43.6%	PenTile
CV1	1080 × 1200	88°	12.27	20.18%	PenTile

Referencing Table 7.9 where the specifications of these devices can be seen, the difference in display density expressed as pixels per degree (PPD) can be seen and used for comparison between devices. Comparing the improvements between the DK1 and DK2, a 43.6% improvement in display density is noted, is close to the 38.4% improvement in median readable distance observed. Between the DK2 and CV1 however, only a 20.18% improvement in pixel density is seen, which is significantly below the readable distance improvement of 90.6%.

When comparing the $\mathbb{D}_n$ improvements for the CV1 against the medium and large sample results for the DK2 — 74.31% and 68.64% respectively — the improvement is seen to be less dramatic, but still higher than the increase in pixel density between devices can account for. While this is not necessarily a fair comparison given the differences in the sample dimensions, it does serve to highlight the difference.

From these results, it can be seen that the hypothesis of readable display distance improvement being proportional to the increase in pixel density does not hold true between the DK2 and CV1. Interestingly, this does appear to hold true for the DK1 and DK2, which is surprising given the change in display technology and pixel layout between the devices, which is not a factor for the DK2 to CV1 step. This indicates that other factors are playing a part in the change in readable distance improvements for the DK2 to CV1 step that require further investigation.

7.7 Discussion

This section expands on the results discussed above, providing context and identifying areas of future work around the Font, ClearType, and Zoom results; limitations of the experimental design, in particular those impacting measurement of reading time per sample; and potential reasons for outlying results.

7.7.1 Fonts

While the results presented above do indicate that the change in font did not affect the $\mathbb{D}_n$ or $\mathbb{D}_f$ values for presented samples, observations made by participants did continue to suggest that there were strong preferences for font choices in samples. Even with feedback being unprompted, multiple participants expressed strong preference towards either of the two font options. While these participants were in the minority, it does suggest that font preference will continue to be a strong driver in selection of fonts within interfaces. It was noted that these observations were only made after the second font option had been shown, with no participants offering an opinion at the commencement of the study, regardless of the first font used in samples.

While this result presented above can conclude that for the given screen and print oriented fonts tested, that there is no difference in readable distance tolerances, this does not provide sufficient evidence to suggest that there is no difference for any two fonts in these categories. Many more combinations of fonts would need to be tested in order to establish this, which is beyond the scope of this research. It should also be noted that this research in no way attempted to identify an “ideal” or “best” font for use within such virtual environments, instead it only sought to determine if an obvious difference could be identified between two common fonts from different typographical backgrounds.

What can be taken away from this and the results presented, is that font choice for conventional applications intending to be presented within an immersive virtual environment such as the one presented in this thesis can make use of existing and established font choices without requiring significant consideration of the font purpose or classification. This is not a broad approval of all font choices; indeed it is obvious that many existing fonts will exhibit varying levels of acceptability within these environments as they do on conventional applications. However, this can be used as evidence that no new fonts need be developed in order to allow conventional applications to be displayed within a virtual environment effectively, nor that font choice be heavily restricted in order to be readable within these environments.

Figure 7.6: Text displayed with lost detail due to texel:pixel mapping mismatch. Text above is shown with ClearType disabled, and below with ClearType enabled. Left text shown at 100% zoom, right at 200%.

7.7.2 ClearType

The results outlined above show that when ClearType was enabled, both $\mathbb{D}_n$ and $\mathbb{D}_f$ values were greater than when it was disabled. It was also noted that while this suggests the use of ClearType is beneficial in this case, this result does not necessarily apply to other font-hinting methods or even for ClearType configurations that vary from the tested defaults. The results also do not identify the reason for the difference, instead only observing that the difference exists.

From comparing magnified text samples with ClearType enabled or disabled, there are two obvious differences. The first being that with ClearType enabled, most glyph pixel colors have been transformed from simple black to some coloured shade, and the second being that each glyph has been made larger horizontally through addition of coloured pixels. As these changes are the only notable differences between the samples, it is obvious that either or both effects are the factor that results in the different $\mathbb{D}$values observed. Of the two, the most likely candidate is the second, where and increase in available pixels for each glyph has resulted in a higher information density within the texture.

Because of the use of nearest-neighbour interpolation for texture filtering within the tested virtual environment, if there is not a 1:1 mapping of texels to pixels, a single texel value will be used from the one or more available texels, with the nearest value used. When dealing with simple black and white text with ClearType disabled, this results in the filtering producing either a black or white pixel on the device with no intermediate values, producing heavily distorted characters as the number of texels to pixels increases. When selecting from a sample with ClearType enabled however, the number of available texel values that can be selected increases, allowing for more variation in the output. An example of this can be seen in Figure 7.6. From this example, the samples with ClearType disabled give output that causes some character glyphs to be completely unidentifiable. This is contrasted with the ClearType enabled samples, that while exhibit a more complex output, retain more detail of glyphs enabling the words to be read more easily. This effect is complicated by the introduction of head tracking, as head movement will cause this filtered output to change as the fixation point moves, making characters “shimmer” in and out of recognisable forms, however the broad effect will remain the same, suggesting that that the additional pixel values given by ClearType will provide a better overall experience.

While the results and discussion suggest that enabling ClearType will improve readable distances of text samples, which as a metric suggests this is the better of the two options, it does not invalidate the observations made by previous study participants suggesting that the colour halos caused by ClearType were distracting. While opinions were not sought from participants regarding their opinion of the text appearance in this study, two participants did note that some of the text appeared “more colourful” than other samples, suggesting that the effect of ClearType did still result in visible colour haloing. As this was a less common observation than in the previous study, it could be concluded that this effect is less pronounced in the CV1 than it was in either the DK1 or DK2, possibly as a result of using different lenses or general improvements in sample positioning. As it is still an observation that was noted however, it does suggest that further investigation should be done on the effect of this on performance. Establishing if this colour haloing has negative effects on reading performance over time, and for aspects other than those already evaluated, such as comprehension and information retention would be a good route to pursue. Further investigation into alternate ClearType settings, or the use of alternate font-hinting methods to help minimise the effect if it is determined to be detrimental could follow.

7.7.3 Zoom

It can be seen from the results above that when displaying text samples captured with a simple zoom of 200%, that the distance at which the sample was considered readable suffered. This can likely be attributed to the zoom approach within Mozilla Firefox, of doubling all display units used within the page when configured with a 200% zoom. This approach would result in a doubling of the font size from 12pt to 24pt and redrawing the text with this

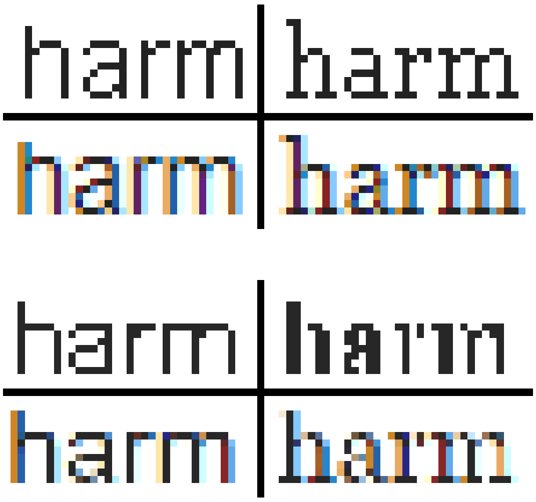

Figure 7.7: Text samples for all permutations of attributes shown as would be seen at a readable distance within the virtual environment, when applied to equally sized virtual windows. The first two rows show ClearType off and on for zoom 100% samples, the following two rows show ClearType off and on for zoom 200% samples. The left column shows Arial font, the right column shows Georgia font. While still readable, the appearance of the zoom 200% samples is noticeably worse than the zoom 100% samples.

new size rather than scaling up the 12pt font linearly. For a browser this approach makes sense, as a 12pt font scaled up would have visible aliasing around the glyphs with applied sub-pixel font hinting no longer applying to single pixels or sub-pixels. This would provide little benefit to users of the browser, for whom the implemented method of adjusting dimensions would likely give the result expected.

When looking at the zoomed samples shown in Figure 7.3, when changing from 12pt to 24pt font additional detail and features have been added to the character glyphs. This is most notable in the Georgia “a” glyph, which features a teardrop shaped counter and a distinctive ball terminal at the larger size. Glyphs displayed at smaller sizes lack these tend to lack these additional flairs, as clarity of characters is often of greater concern than apparent style. By capturing samples using more detailed glyphs and then scaling the texture down, additional visual noise was introduced through additional glyph details, making some characters harder to distinguish. This can be seen clearly in Figure 7.7 which shows the 200% zoom sample scaled down to the same dimensions as the zoom 100% sample, where the noted Georgia samples rendered at 200% are clearly distorted, with the ClearType off sample suffering the most. This

Figure 7.8: Text samples for all permutations of attributes shown as would be seen at a distance within the virtual environment, when applied to a equally sized virtual windows. The first two rows show ClearType off and on for zoom 100% samples, the following two rows show ClearType off and on for zoom 200% samples. The left column shows Arial font, the right column shows Georgia font. The degraded appearance of the zoom 200% samples is notable when compared to zoom 100% samples at the same physical size.

effect is amplified further when samples are scaled down further to demonstrate samples displayed at greater distances from the perceiver within the virtual environment in Figure 7.8, where more detail is lost when compared with the equivalently sized zoom 100% samples.

While the intention behind application of zoom was to increase the number of available texels for the nearest-neighbour interpolation filtering to select from, in the hope that it may achieve a similar result as with ClearType, this did not result in a benefit to readable distance ratings. While an alternate approaches to scaling samples, such as post-processing the 100% zoom samples to double the size of each captured pixel could be pursued, it seems that alternative modifications may yield better results. Outside of tuning ClearType or alternative font-hinting mechanisms, potential investigation routes could consider trialing the effect of other texture interpolation methods than nearest-neighbour, to determine changing this can give a positive effect to readable distance instead.

7.7.4 Acceptable Display Distances

Data collected from this study can be used in aggregate to identify a range of distances at which text is considered readable, and readable with effort for a given percentage of the population. When considering only participants who took part in the full-length experiment, a visualisation of acceptable distances can be generated and seen in Figure 7.9. It can be seen from this graph that there is no single $\mathbb{D}_n$ or $\mathbb{D}_f$ value that was agreed on amongst participants, instead variability in participant tolerances for text clarity form a regular trend for both values.

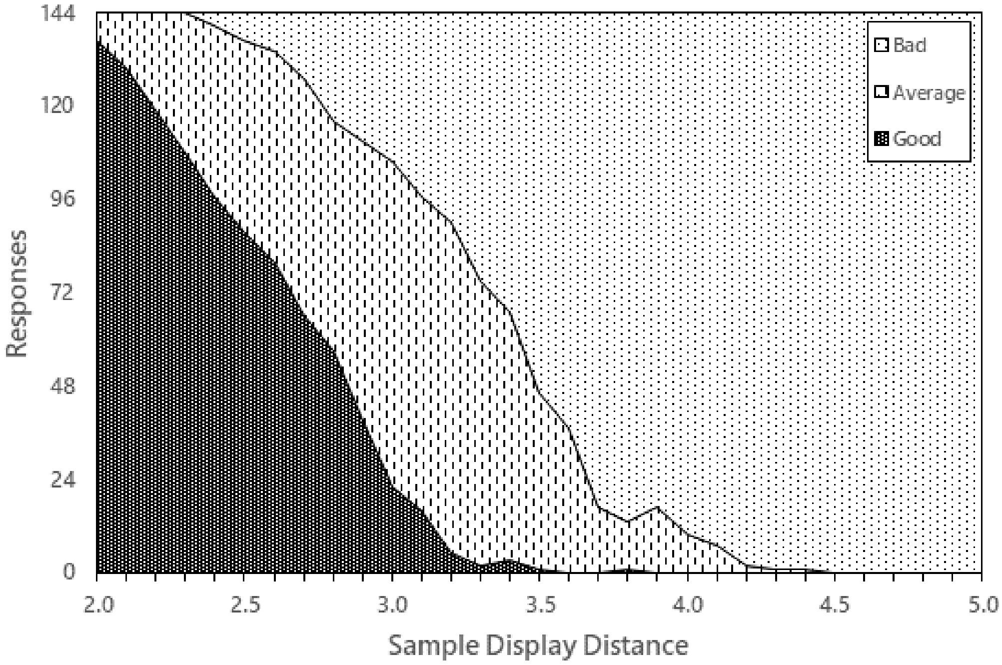

Figure 7.9: Area chart showing the readability ratings given by all participants taking part in the full-length study across all samples.

Using this data, rough guidelines for default positioning of text-bearing windows within the virtual environment can be determined. If aiming for a 50% acceptability rating, $\mathbb{D}_n \approx 2.7$ and $\mathbb{D}_f \approx 3.4$ can be identified as bounds for the maximum readable range of 12pt text within the environment. While these values will not be ideal for all users of such a system, they do identify points that should not require much adjustment for a significant proportion of users and would therefore be useful points to consider as defaults when considering future studies.

7.7.5 Device Improvement

In the results shown above, it was identified that factors other than simply display density played a part in the recorded $\mathbb{D}_n$ differences noted between the DK2 and CV1 devices. Several candidates have been identified as possible contributing factors that could be further investigated to determine impact.

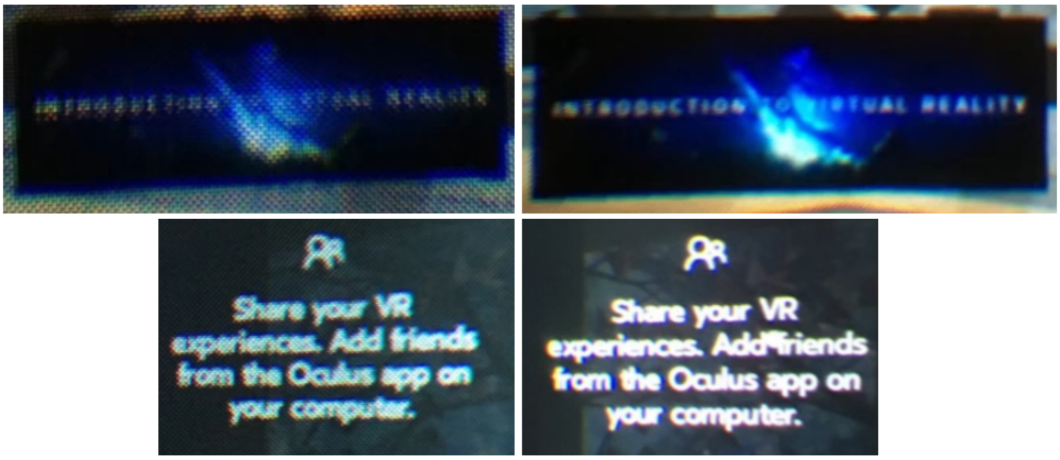

Figure 7.10: Screen door effect observed through the lenses of an Oculus Rift DK2 (DK2) and Oculus Rift CV1 (CV1), focusing on text samples. Images shown to the left are from the DK2, and the images on the right from the CV1. Images extracted from video comparing HMDs (VRdudeBRO, 2016).

As discussed with regards to fixed vs free distance selections, there is a difference in methodologies between the two studies for selection of $\mathbb{D}_n$. In the previous study, users were asked to position text-bearing virtual windows at the furthest distance where comfortable reading could be achieved, before being asked to read the positioned text. Variation was observed within participant positions selected for different sizes of virtual window, suggesting that the position they chose was dependent at least in part by the virtual window size. As participants knew they would need to read text aloud from these samples, they may have also been conservative when selecting positions to ensure they performed well. Given the free positioning allowed to them, participants may have selected the first comfortable reading position identified rather than hunting for the furthest point where it may have been achievable. This contrasts with the approach used in this study, where a series of fixed positions were given with no requirement to read the text aloud. This may have influenced participants to be more optimistic with ratings, giving samples that would have been challenging to read “clearly readable” ratings when “readable with effort” may have been more appropriate.

Outside of methodology differences, an observed difference in the DK2 and CV1 devices is the severity of the screen door effect seen on the display. This effect, where gaps between pixels on the display can be seen as a grid across the displayed image, is more pronounced on both the DK1 and DK2 devices than on the CV1. As can be seen in Figure 7.10, this difference has a noticeable impact on the clarity of text presented on the devices, with the DK2 presented text being less clear than on the CV1. Given both the DK1 and DK2 both

Table 7.10: Inter-pupillary distance (IPD) values for male and female U.S. Army personnel. Table summarises data presented by Gordon et al., 2014, p. 150. Values are expressed in millimetres.

Gender	$\bar{x} \pm s$ (IPD)	N	Q0	Q4
Female	61.7 ± 3.6	1986	51.0	74.5
Male	64.0 ± 3.4	4082	53.0	77.0

exhibited noticeable screen door effect on the images, the impact of it on the $\mathbb{D}_n$ values collected in the previous study is minimal, which is reflected in the $\mathbb{D}_n$ improvement value tracking closely with the display density improvements on the devices. With a noticeable difference in the screen door effect between the DK2 and CV1, it is a strong candidate for being a factor in the observed difference in $\mathbb{D}_n$ for these devices.

Another point of difference between devices is the lens shape used in the HMD optics. Both the DK1 and DK2 make use of aspheric biconvex lenses between the eyes of the user and the display, whilst the CV1 uses hybrid fresnel lenses. Biconvex lenses are known to cause a variety of different aberrations, including chromatic aberration caused by a failure of the lens to focus all colours correctly, leading to blur or visible colour separation in images. This effect was another possible candidate for the cause of observed colour haloing in text samples from the previous study, which was noted as not as influential as the effects of ClearType. This effect could however contribute to overall lower acceptable $\mathbb{D}_n$ values observed from the previous study, with the effects degrading text enough to require a nearer placement to the perceiver.

A final point of difference between the devices is the IPD management of each device. For the DK1 and DK2, the physical IPD was fixed to 63.5mm, whereas the CV1 offered a hardware adjustable range of 58–72mm IPD adjustment. A study by Gordon et al. (2014) identified a range of observed IPD values in soldiers, which shows that the fixed IPD values used for the DK1 and DK2 lie between the mean values observed for males and females, while the adjustable range of the CV1 comfortably covers one standard deviation from the mean for these values. While both the DK1 and DK2 offer software correction for IPD, the interaction with the physical lens positions with relation to the participant's actual IPD could result in a non-ideal viewing configuration. It is noted that a miss-match in IPD values for a user can cause increased eyestrain (Murray, 2017, p. 62), which may again have affected the $\mathbb{D}_n$ results from the previous study.

7.7.6 Time per Sample

While the overall time taken to complete the study was collected for each participant, no timing information was recorded for individual sample display and response times. It was observed in all cases that during the experimental runs, participants were quick to judge samples that fit into the “clearly readable” and “unreadable” categories but spent more time considering samples around and between $\mathbb{D}_n$ or $\mathbb{D}_f$. Future studies should consider recording individual consideration and response time per sample, as time spent per sample may provide insight into the exact level or readability that a sample presents.

7.7.7 Outliers

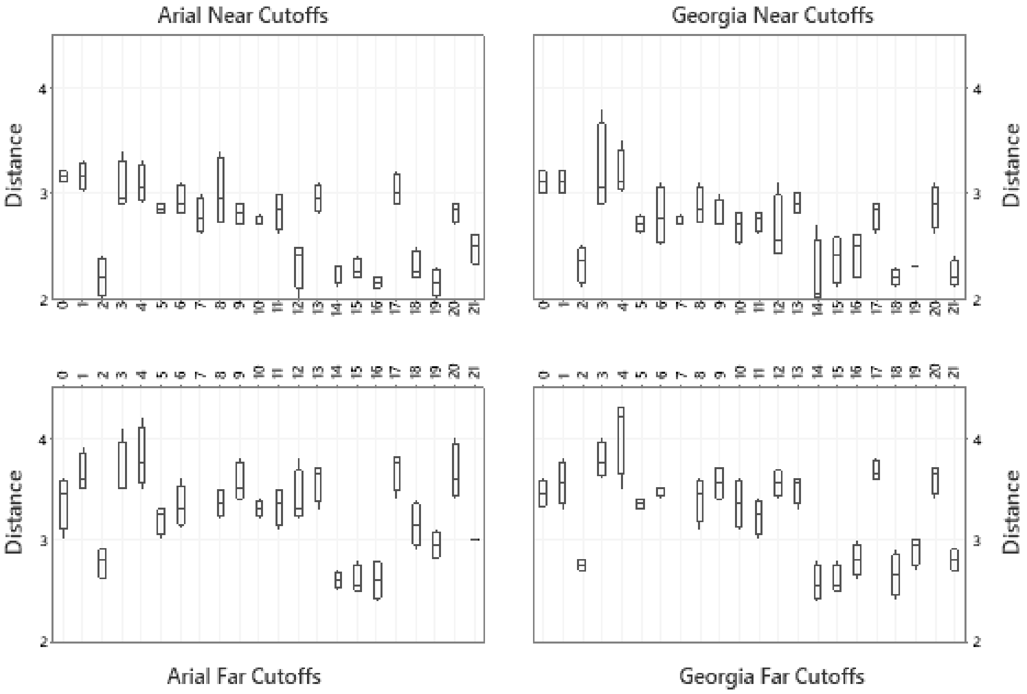

Figure 7.11: Boxplots showing the impact of independent variable font on near and far cutoff points for individual participants.

When inspecting the individual participant $\mathbb{D}_n$ and $\mathbb{D}_f$ values visualised in Figures 7.11 to 7.13, it can be seen that there is are two sets of participants, with one set having a lower distance tolerance than the other. When considering this low tolerance group with the demographic information collected, there is no obvious reason for this. As this set is consistent across all permutations of samples however, it suggests that these participants have a common connecting factor that causes this difference.

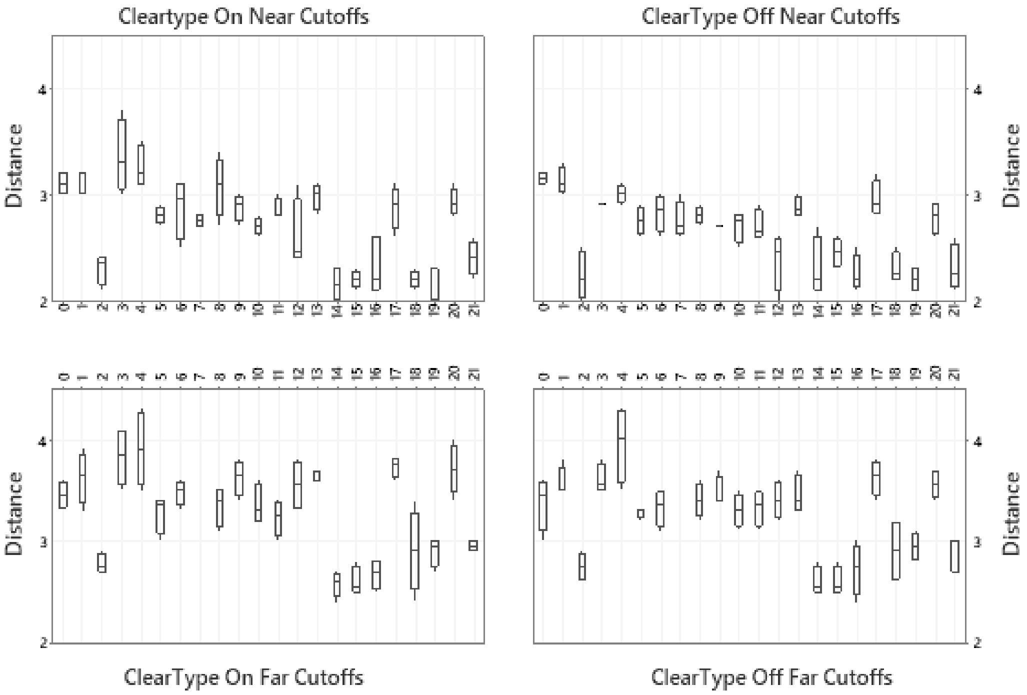

Figure 7.12: Boxplots showing the impact of independent variable ClearType on $\mathbb{D}_n$ and $\mathbb{D}_f$ points for individual participants.

Participants were selected from the University of Waikato student body, with no restriction on participation outside of possible interactions between spectacles and the CV1. As the University has a significant body of ESL international students, this represents a candidate for the common factor connecting members of this set. Guidelines given for the design of websites consumed by ESL students indicates that both font choice and text density are important factors for students navigating and consuming content (Liu et al., 2008), which could indicate a reduced tolerance for unclear text displayed in such an environment.

Future experiments should consider collecting additional demographic information from participants, including first language and writing system, to possibly provide insight into this difference.

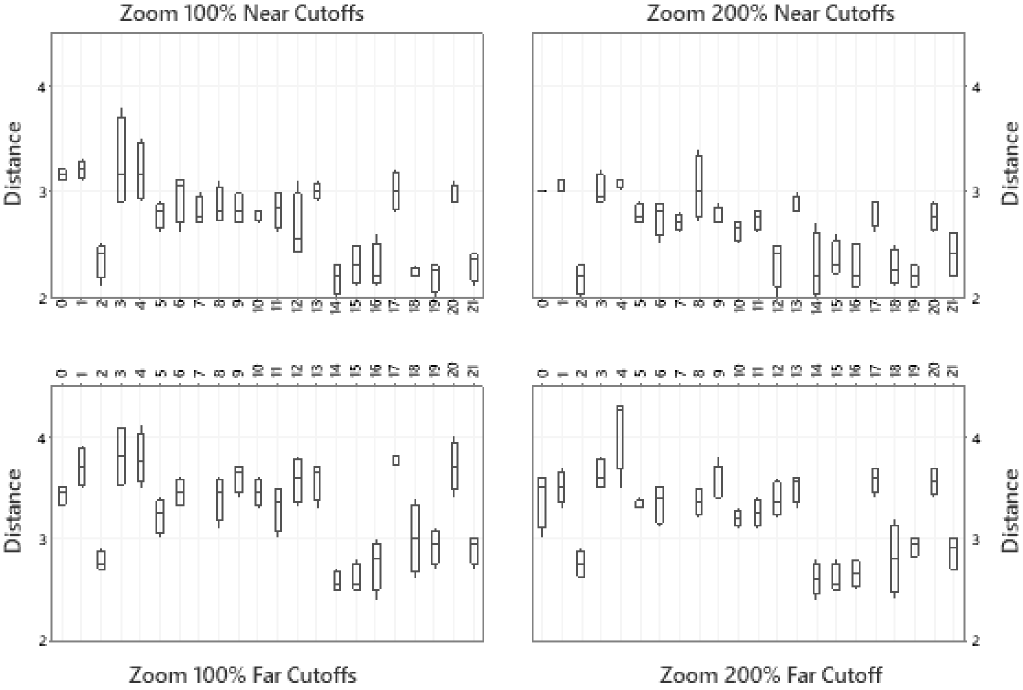

Figure 7.13: Boxplots showing the impact of independent variable zoom on $\mathbb{D}_n$ and $\mathbb{D}_f$ points for individual participants.

Chapter 8

Summary and Conclusions

This thesis investigated the extent to which conventional applications could be used, unmodified, within an immersive virtual environment designed for a seated office experience. This research was driven by the observation that despite improvements in both conventional display technology and window management techniques, users still had difficulty effectively managing increasing numbers of application windows for their increasingly digital task load.

Given that usage of window management techniques was minimal among an observed cohort of diverse professionals, a way of improving their ability to deal with their window load is to increase their available space through adding additional displays or replacing the display technology with a possible alternative.

We investigated constructing an immersive virtual reality (IVR) based general-purpose computing environment that displayed conventional windows surrounding the user. This was done on the basis that with any change in technology or capabilities there will be a lag time before applications taking advantage on the new technology will be developed. As such, supporting now legacy applications within such an environment would be helpful in allowing adoption of this technology.

8.1 Research Questions

The investigations covered in this thesis seek to address an overarching research question:

> ***To what extent can applications designed for conventional interfaces and displays be presented and used inside a general purpose immersive virtual reality environment with only user preference modifications?***

To address this question, four further research questions were identified:

- RQ1: What does a conventional computing workspace look like for modern office workers in New Zealand?
- RQ2: What performance metric should be used when comparing applications within differing environments?
- RQ3: With respect to the chosen metric, to what extent does changing from a conventional environment to a virtual one affect performance?
- RQ4: With respect to the chosen metric, what software configuration changes can be easily made to conventional applications that may improve performance within a virtual environment?

In this section, each of these questions is considered in turn.

8.1.1 RQ1: Conventional Computing

To address Research Question 1, which stated:

> ***What does a conventional computing workspace look like for modern office workers in New Zealand?***

A definition of conventional computing was first established in Chapter 2, identifying laptop and desktop personal computers as the devices of interest for this thesis. It was observed that conventional computing devices typically utilise pointing and typing devices for input, and graphical displays for output. Multiple forms of each were identified, but it was seen that the underlying functionality and broad appearance of these has changed little over years of device evolution. The most notable difference observed between historical and modern personal computers were the displays, which were shown to be trending larger both in diagonal size and pixel resolution and were seen to be used increasingly in multi-monitor configurations to increase working space for computing tasks.

Window management techniques and tooling were outlined, with compositing window managers using a stacking approach identified as the primary window management approaches seen with modern operating systems. Supporting and assistive tooling for window management was examined, identifying common features of operating system window managers that could be leveraged for window and space management. Benefits and drawbacks of the use of large or multiple display configurations were identified, highlighting that current window managers struggle with the increasing numbers of windows displayed by users.

Following this analysis, Chapter 3 presented a workplace survey conducted at a large New Zealand business. Office workers at this company were observed at intervals throughout a working day to identify the number of tasks being worked on simultaneously on their conventional computing devices, the percentage of their workday spent on digital tasks, and the number of windows utilised to support each task. Each office worker was additionally surveyed to obtain feedback about their digital work behaviours and difficulties they encountered while performing tasks. Information was collected about their computing environment, workspace, and their individual opinions on the equipment provided.

From this study, it was determined that office workers at the surveyed company largely made use of desktop computers with two or three displays laid out in a horizontal fashion across a fixed desk surface. All made use of physical mice and keyboards for input, and LCD displays diagonally larger than 22" on average for output. They were observed to spend most of their day working digitally, with most indicating they would prefer a fully digital workspace, working on multiple tasks each of which involved multiple windows. While most expressed satisfaction with their current computing system and connected displays, they also indicated they believed they would be more productive with larger and more numerous displays,

This desire for increased display space appeared related to both the number of windows open on their displays at any one time, and poor usage of supporting window management techniques provided by their operating system. A majority indicated they regularly had difficulty locating windows on their systems and believed that an increase in digital space would address this.

8.1.2 RQ2: Evaluation Metric

To address Research Question 2, which stated:

> ***What performance metric should be used when comparing applications within differing environments?***

Tasks performed by office workers in Chapter 3 were examined. Lifecycle management of application windows, layout of windows, identification of windows, and consumption of application content were identified as tasks performed by the observed office workers that might present issues when moving to an immersive virtual reality environment. Lifecycle management of windows was not observed to take a significant amount of time, with actions performed both quickly and infrequently during observation. Window layout actions were expected to be important tasks when commencing the survey, however the observed office workers were rarely seen performing this category of task, with the typical behaviour involving positioning and sizing each window once when opening the window, then leaving them in place until they were closed. Locating specific windows was a frequently performed task, however, was not observed to occupy a significant proportion of most office worker's days. Office workers with larger numbers of tasks and windows spent significantly longer than their peers on identification tasks but were seemingly more skilled at rapidly identifying windows of interest.

The final common task performed, that of content consumption, was by far the most time consuming and frequent task performed by observed workers, with the bulk of their day spent consuming content of varying kinds. Within the banner of content consumption, reading text was identified as the most time-consuming task performed in a working day for the surveyed office workers, which was supported by the observation that on average, 312 minutes per working day (65% of an 8 hour working day) is spent on reading tasks (White et al., 2010, p. 291). Due to the large proportion of the working day consumed by this task, reading was chosen as the primary metric for evaluation of a new interface, identifying that should a new interface negatively impact the performance of reading, then this would have a large impact on user's ability to complete tasks efficiently.

8.1.3 RQ3: Environment Performance

To address Research Question 3, which when updated to include mention of the RQ2 chosen metric of reading, stated:

> ***With respect to [reading], to what extent does changing from a conventional environment to a virtual one affect performance?***

A test-bed immersive virtual environment was created, running on the Oculus Rift series of IVR head-mounted displays (HMDs). This environment, described in Chapter 5, was utilised in the user study described in Chapter 6 to present static captures of text-bearing conventional application windows at varying sizes and distortions ahead of users. Study participants viewed these text samples through two different IVR HMDs, the Oculus Rift DK1 (DK1) and Oculus Rift DK2 (DK2) and were asked to position the presented windows at a comfortable reading distance before reading indicated passages from the presented text aloud. Conventional application windows of the same size were additionally presented on a conventional computer display, and users were asked to again read indicated passages from the presented text aloud. Recordings were taken of read text, and following processing described in Section 6.5, were against samples of the same size for each participant.

It was found that when displaying on narrow windows, reading content within the virtual environment could be performed with the same accuracy and speed as it could when read from a conventional display. This equivalent performance was noted on both tested HMDs across both tested window distortions. When displaying wider content applied to curved windows, reading speed dropped when compared to conventional display performance, though accuracy of reading remained equivalent. When wider content was displayed on a planar window, both accuracy and speed of reading was seen to decrease relative to the increasing width of the content being read.

8.1.4 RQ4: Software Configuration Changes

To address Research Question 4, which when updated to include mention of the RQ2 chosen metric of reading, stated:

> ***With respect to [reading], what software configuration changes can be easily made to conventional applications that may improve performance within a virtual environment?***

Feedback from participants of the previous study was examined, and three configuration changes were identified. Font choice was raised by several participants as a potential influencing factor in their reading performance and choice of window position, and a serifed font was selected for comparison with the previously used sans-serif font. A “haloing” effect around text was identified as distracting when performing reading tasks and following investigation the cause was identified as ClearType sub-pixel anti-aliasing which could be enabled or disabled. Loss of detail in displayed font glyphs was attributed to poor texture filtering, and a proposed solution was to capture and display higher density textures through increasing application window size and zooming or scaling content prior to capture, the efficacy of which required investigation.

A further user study was conducted using the developed virtual environment, presenting text bearing windows at fixed locations along with response buttons allowing for rating the readability of the displayed text. Text was presented with different permutations of each of the three identified configuration options at each location, and readability ratings were collected for each.

Captured ratings were processed for each permutation of configurations to identify the distance at which each readability response changed, identifying a point where readability transitioned from easily readable to readable with effort, and a second point where readability moved to unreadable.

From comparing the readability transition point distances, each configuration was examined to identify the effect of the change. It was identified that the choice of font had no statistically significant impact on readable distances, suggesting that the choice of either a common serif or sans-serif will not impact on the readability of text presented. ClearType status was shown to have a statistically significant impact on readable distance, with text displayed with ClearType enabled being rated as readable at greater distances than when disabled. Texture zoom level was also shown to have a statistically significant impact on readable distance, with text shown at 100% scale being readable at greater distances than text displayed at 200%.

8.1.5 Primary Research Question

With the supporting research questions addressed, the primary research question can be considered. This research question stated:

> ***To what extent can applications designed for conventional interfaces and displays be presented and used inside a general purpose immersive virtual reality environment with only user preference modifications?***

And was addressed in several stages. First, observations made during the workplace study described in Chapter 3 were analysed to identify common tasks performed by users. From the subset of common tasks, the task that occupied the greatest proportion of each working day was identified as content consumption through reading, which was used as a measure for performance going forward.

An immersive spherical virtual environment was then developed that allowed for presentation of conventional applications around the perceiver. Different geometries for virtual windows were investigated, and a curved distortion that ensured that every point on the surface was equidistant from the perceiver along with a simple planar presentation were selected as preferred representations for windows. Feedback identified that most users preferred the curved windows despite large windows inducing a sense of claustrophobia, while a minority of users identified planar presentations as preferable in spite of issues observed with viewing content on the window peripheries.

Interaction with the virtual environment was accomplished through orientation and position sensors accompanying the IVR HMDs on which the environment was displayed, the Oculus remote device, and conventional mouse and keyboard devices. A mouse cursor was bound to the movements of the mouse, converting translational movements into rotations within the environment, allowing for established application pointer interactions to be performed. Users of the environment identified the mouse interaction as pleasant and easy to use, with no learning curve required to interact with applications in established ways. Lack of ability to see the mouse and keyboard was identified as problematic, more-so for the keyboard than the mouse by users not experiences with touch-typing.

A reading experiment performed using the developed environment demonstrated that reading tasks could be performed on narrow virtual windows displaying conventional applications, with equivalent speed and accuracy as is seen when performing the same task on a conventional display. When reading from wider virtual windows, user preference of window distortion played a role in reading performance, with curved windows offering equivalent reading accuracy at lower speeds, and planar windows performing worse in both speed and accuracy. The newer DK2 IVR HMD exhibited lesser speed and performance penalties in all cases when compared to the older DK1 device, suggesting that with sufficient improvements in HMDs specifications, the observed performance penalties will be of lesser in all presentations.

A second experiment was performed using the developed environment, comparing the effects of font choice, ClearType sub-pixel anti-aliasing, and interface zoom on readable distances. It was found that both ClearType and interface zoom had impacts on the readable distance of text samples, with the presence of ClearType allowing for more distant position while remaining readable, and 200% zoomed content requiring placement nearer the perceiver to be readable. Font choice of common serif and sans-serif fonts were not seen to affect the readable distance.

From these results, it can be concluded that presenting and interacting with applications designed for display on conventional displays in an immersive virtual environment is possible and reading tasks can be performed with equal accuracy and speed when users perform appropriate preference modifications. The performance of reading tasks can be tailored to match conventional display performance by resizing application windows such that they present as narrow windows, and when wider windows are required reading accuracy can be maintained through use of a curved window. Readability of text can be improved through enabling of ClearType, allowing text to be read at greater distances.

8.2 Limitations

The work presented in this thesis took place over several years, and development of tools along with study processes are products of their time. Of note, the devices utilised for this research — the DK1, DK2, and Oculus Rift CV1 (CV1) — were released in 2013, 2014, and 2016 respectively, and represented the state of IVR HMD hardware readily available on the market at the time. Since conducting the studies presented in Chapters 6 and 7 using these devices, the state of the art has moved on, with modern devices offering superior resolutions, fields of view, optics, and software. With these newer devices available, should the presented experiments have been re-performed to make use of these devices, producing results directly applicable to the current state of the art?

The reading study presented in Chapter 6 was conducted in 2015, utilising the DK1 and DK2 which were current and relevant devices at the time. The results from this study were published when the results were new (See Appendix B), and the presented results have been used to justify design decisions by other researchers making use of more modern equipment (Dingler et al., 2018; Kojić et al., 2020; Dingler et al., 2020; Rzayev et al., 2021; Kim & Shin, 2021), suggesting that the results presented were still relevant to these researchers. Further, the method for measuring and reading performance utilising running records as presented has seen use by other researchers (Wei et al., 2020), evidencing contributions to the field applicable regardless of the utilised devices.

The CV1 device utilised in the study presented in Chapter 7, conducted between 2018 and 2020, continues to see significant use in research publications across a range of fields (Dingler et al., 2020; Labbe et al., 2021; Caserman et al., 2021; Zujovic et al., 2021; Luu et al., 2021; Yang & Kim, 2021). This continued use suggests that this device, and results collected using it, remain relevant even when newer devices with improved specifications are readily available.

While re-conducting these studies on newer devices may yield results that could have direct applications to the newer tested devices, the results collected utilising older devices still have application to the newer devices as the core findings remain relevant. It was identified in the Chapter 6 study that reading performance of conventional applications within the older devices was comparable to that seen on conventional displays at that time, and this remains true even as new devices are released. It can be expected that newer devices that improve on specifications and features would not degrade the experience, so the results presented represent a baseline from which results can be expected to improve in step with technological device developments.

8.3 Future Work

As is often the case with research tasks, additional questions and avenues for investigation were identified while addressing the above research questions. While many different directions for future research were identified over the course of this research, two were identified as of particular interest, and are presented in this section.

8.3.1 Curving & Dancing Windows

It was identified through user feedback in Chapters 5 and 6 that planar representations of virtual windows were preferred by a subset of users despite the limitations to reading performance when scaled to larger sizes. While it was demonstrated that utilising a curved representation of virtual window addressed display limitations, demonstrating superior reading accuracy and speeds, the planar presentation was still indicated as preferable to these users for a variety of expressed reasons, the most prominent of which being a perceived “unnatural” feeling when looking at curved windows. To make virtual window representations both visually acceptable to these users and readable with acceptable performance, different behaviours and geometric representations of virtual windows could be investigated.

A possible solution to the identified concern of curved windows feeling unnatural would be to reduce the curvature effect on smaller windows where the curve effect is not required, only applying the effect when necessary. Windows could be displayed as planar under a certain dimension and curved to the tested extent when an upper size is met, with intermediate sizes applying a less pronounced curve effect. This effect could be applied to vertical and horizontal sizes separately, utilising a single axis cylindrical distortion for wide-and-short, or tall-and-narrow windows in preference to a dual axis solution as presented. Investigation into the acceptability of such a solution, and determining upper and lower bounds for applying curvature, could contribute valuable guidance for interface design and presentation within immersive virtual environments, making these interfaces more acceptable to a wider range of users.

Addressing concerns of text readability could also be investigated using techniques other than distortion of windows. A possible avenue for this would be that of “dancing windows”, that move or rotate as the user interacts with them. As the user shifts their gaze across a planar window, translation or rotation of the window could occur to bring the observed content closer to the perceiver and move the angle of incidence of the gaze ray closer to 90°. If implemented correctly, this could provide a workable approach to improving reading performance on large planar windows, however this may not scale effectively when multiple windows are displayed and may be unpleasant to use if window movements are significant. Evaluation of different approaches to this idea could produce useful results, guiding future immersive virtual environment interface development.

8.3.2 Intelligent Text Extraction and Presentation

This research identified that text could be read effectively from unmodified conventional application windows presented within an immersive virtual environment, however clarity of character glyphs was identified as a factor that affected the visual quality of the text, and the distance at which the text was considered readable. A cause of the reduced glyph clarity was identified as being due to a miss-match between texels and device pixels, leading to lossy presentations of content when displayed on a device (See Figure 7.6). As underlying application windows are not aware of their presentation within an immersive virtual environment, they are unable to modify their appearance automatically to address this issue, requiring manual adjustment of application display settings or movement of windows within the environment to improve text quality.

A potential approach that could be considered for improving the quality of text would be for the environment to detect and extract text content from underlying windows, then re-displaying this content within the virtual environment, following virtual reality text presentation best practices to maximise legibility. Processing application window textures with optical character recognition (OCR) techniques would allow text content to be extracted as windows update, capturing the text content and regions of the window where the text was found. The environment could then overlay a paginated or RSVP presentation of the text with appropriate controls, displayed appropriately for the environment configuration. While this approach is unlikely to function perfectly with all conventional applications, a compatible list of applications could be maintained where this approach is automatically applied, with unsupported application's windows left alone. Investigation into the viability of this approach could be undertaken, and if shown to be functional, different presentations could be trialed with real-world workflows to examine the impact on both reading and task performance.

8.4 Closing Remarks

Affordable, high-quality, mass-market virtual reality devices have the potential to reinvent the way that users interact with computer systems. Providing users with access to novel interaction methods and a large, visually encompassing virtual display space unlocks the potential for equally novel interfaces and environments to be developed and utilised. The use of a IVR HMD to present an immersive virtual environment displaying a continuous virtual space upon which windows can be arranged as seen in this thesis, represents a potential intermediary environment that could be seen in a hypothetical transition from conventional computing systems to portable virtual reality environments.

Barriers for adoption of new or alternate technologies are numerous, with lack of applications — both practical and software — for new devices often seen as the death knell for novel systems. Ensuring that existing software, workflows, and tasks can be performed on new devices may be an important contributor to the success or failure of virtual reality device adoption for general purpose computing. Allowing users to replace existing computing display systems to leverage virtual reality specific behaviours and interactions, whilst retaining the ability to work in a familiar way with known software may reduce friction in adoption of virtual reality technologies in the workplace and household.

This thesis has demonstrated that this approach to transitional interfaces is viable, showing that existing applications can be displayed, common tasks performed, and existing input devices utilised within an immersive virtual environment presented on an IVR HMD. While additional work is still required to address other potential points of friction towards adoption, this initial step towards a potential next step in computer interaction provides a view of an exciting possible future.

References

Agarawala, Anand and Ravin Balakrishnan (2006). ‘Keepin’ it real: pushing the desktop metaphor with physics, piles and the pen’. In: *Proceedings of the SIGCHI conference on Human Factors in computing systems - CHI ’06*. CHI ’06. New York, NY, USA: ACM, pp. 1283–1292. ISBN: 1595933727. DOI: `10.1145/1124772.1124965`.

Ahmad, Norhafizan, Raja Ariffin Raja Ghazilla, Nazirah M. Khairi and Vijayabaskar Kasi (2013). ‘Reviews on Various Inertial Measurement Unit (IMU) Sensor Applications’. In: *International Journal of Signal Processing Systems*, pp. 256–262. DOI: `10.12720/ijsps.1.2.256-262`.

Ali, Ahmad Zamzuri Mohamad, Rahani Wahid, Khairulanuar Samsudin and Muhammad Zaffwan Idris (2013). ‘Reading on the Computer Screen: Does Font Type has Effects on Web Text Readability?’ In: *International Education Studies* 6.3. DOI: `10.5539/ies.v6n3p26`.

Amir, Guy, Ayala Prusak, Tal Reiss, Nir Zabari and Dror G. Feitelson (2021). ‘Use and Perceptions of Multi-Monitor Workstations: A Natural Experiment’. In: arXiv: `2103.13198 [cs.HC]`.

Anderson, J., Jennifer Hill, Paul Parkin and A. Garrison (2007). *Productivity, Screens, and Aspect Ratios; A Comparison of Single, Traditional Aspect, Dual, Traditional Aspect, and Single, Widescreen Aspect Computer Displays over Simulated Office Tasks across Performance and Usability.* Research rep. University of Utah. URL: `https://collections.lib.utah.edu/ark:/87278/s69w3dmf`.

Andrews, Christopher, Alex Endert and Chris North (2010). ‘Space to Think: Large, High-Resolution Displays for Sensemaking’. In: *Proceedings of the 28th international conference on Human factors in computing systems - CHI ’10*. ACM Press. DOI: `10.1145/1753326.1753336`.

ASUSTeK Computer Inc. (2021). *Zenbook Pro Duo UX581.* URL: `https://web.archive.org/web/20210907162719/https://www.asus.com/nz/Laptops/For-Home/ZenBook/ZenBook-Pro-Duo-UX581/`.

Autodesk Inc (2021). *VR Software for Virtual Reality Design.* URL: `https://web.archive.org/web/20210825224301/https://www.autodesk.com/solutions/virtual-reality`.

AwesomeWM Project (2021). *awesome window manager.* URL: `https://web.archive.org/web/20210829080503/https://awesomewm.org/`.

Back, Tycho T. de, Angelica M. Tinga, Phong Nguyen and Max M. Louwerse (2020). 'Benefits of immersive collaborative learning in CAVE-based virtual reality'. In: *International Journal of Educational Technology in Higher Education* 17.1. DOI: `10.1186/s41239-020-00228-9`.

Bahiss, Khadija (2021). 'A serious game to provide social skills training for people with Asperger's Syndrome/High-Functioning Autism'. PhD thesis. The University of Waikato. URL: `https://hdl.handle.net/10289/14222`.

Bai, Nan, Yuan Tian, Ye Liu, Zhengxi Yuan, Zhuoling Xiao and Jun Zhou (2020). 'A High-Precision and Low-Cost IMU-Based Indoor Pedestrian Positioning Technique'. In: *IEEE Sensors Journal* 20.12, pp. 6716–6726. DOI: `10.1109/jsen.2020.2976102`.

Ballenghein, Ugo and Thierry Baccino (2018). 'Referential processing during reading: concurrent recordings of eye movements and head motion'. In: *Cognitive Processing* 20.3, pp. 371–384. DOI: `10.1007/s10339-018-0894-1`.

Benligiray, Burak, Cihan Topal and Cuneyt Akinlar (2019). 'STag: A stable fiducial marker system'. In: *Image and Vision Computing* 89, pp. 158–169. DOI: `10.1016/j.imavis.2019.06.007`.

Berg, Mark de, Otfried Cheong, Marc van Kreveld and Mark Overmars (2008). 'Binary Space Partitions'. In: *Computational Geometry: Algorithms and Applications.* 3rd ed. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 259–281. ISBN: 9783540779742. DOI: `10.1007/978-3-540-77974-2_12`. URL: `https://www.worldcat.org/oclc/1264775125`.

Berger, Laurenz and Katrin Wolf (2018). 'WIM: Fast Locomotion in Virtual Reality with SpatialOrientation Gain & without Motion Sickness'. In: *Proceedings of the 17th International Conference on Mobile and Ubiquitous Multimedia.* ACM. DOI: `10.1145/3282894.3282932`.

Beymer, David, Daniel M. Russell and Peter Z. Orton (2005). 'Wide vs. narrow paragraphs: an eye tracking analysis'. In: *IFIP Conference on human-computer interaction.* Springer. Springer Berlin Heidelberg, pp. 741–752. DOI: `10.1007/11555261_59`.

Bi, Xiaojun, Seok-Hyung Bae and Ravin Balakrishnan (2010). 'Effects of interior bezels of tiled-monitor large displays on visual search, tunnel steering, and target selection'. In: *Proceedings of the 28th international conference on Human factors in computing systems - CHI '10.* ACM Press. DOI: `10.1145/1753326.1753337`.

Bi, Xiaojun and Ravin Balakrishnan (2009). 'Comparing usage of a large high-resolution display to single or dual desktop displays for daily work'. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM. DOI: `10.1145/1518701.1518855`.

BIM Holoview Ltd (2021). *BIM Holoview*. URL: `https://web.archive.org/web/20210306111134/http://bimholoview.com/`.

Biocca, Frank and Ben Delaney (1995). 'Immersive Virtual Reality Technology'. In: *Communication in the Age of Virtual Reality*. USA: Lawrence Erlbaum Associates Inc., pp. 57–124. ISBN: 9780805815498. URL: `https://www.worldcat.org/oclc/30663651`.

Blanchard, Harry E. and Asghar Iran-Nejad (1987). 'Comprehension processes and eye movement patterns in the reading of surprise-ending stories'. In: *Discourse Processes* 10.1, pp. 127–138. DOI: `10.1080/01638538709544663`.

Bly, Sara A. and Jarrett. K. Rosenberg (1986). 'A comparison of tiled and overlapping windows'. In: *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '86*. CHI '86. New York, New York, USA: ACM Press, pp. 101–106. ISBN: 0897911806. DOI: `10.1145/22627.22356`.

Booth, Robert W. and Ulrich W. Weger (2012). 'The function of regressions in reading: Backward eye movements allow rereading'. In: *Memory & Cognition* 41.1, pp. 82–97. DOI: `10.3758/s13421-012-0244-y`.

Bordynuik, John (2005). *Xerox Mice*. URL: `https://web.archive.org/web/20210316082813/http://oldmouse.com/mouse/xerox/alto.shtml`.

Bosch Sensortec GmbH (2021). *Inertial Measurement Unit BMI085*. URL: `https://web.archive.org/web/20210628133841/https://www.bosch-sensortec.com/products/motion-sensors/imus/bmi085/`.

Burge, J. and W. S. Geisler (2011). 'Optimal defocus estimation in individual natural images'. In: *Proceedings of the National Academy of Sciences* 108.40, pp. 16849–16854. DOI: `10.1073/pnas.1108491108`.

Burruss, Caleb C., Elizabeth Bjornsen and Kaitlin M. Gallagher (2021). 'Examining Potential User Experience Trade-Offs Between Common Computer Display Configurations'. In: *Human Factors: The Journal of the Human Factors and Ergonomics Society*, p. 001872082110183. DOI: `10.1177/00187208211018344`.

Buttner, Andre, Stefan M. Grunvogel and Arnulph Fuhrmann (2020). 'The influence of text rotation, font and distance on legibility in VR'. In: *2020 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*. IEEE. DOI: `10.1109/vrw50115.2020.00182`.

Carpenter, Roger H. S. (1988). *Movements of the Eyes, 2nd Rev.* Pion Limited. ISBN: 9780850861099. URL: `https://www.worldcat.org/oclc/440095334`.

Carvalho, Felipe G., Daniela G. Trevisan and Alberto Raposo (2011). ‘Toward the design of transitional interfaces: an exploratory study on a semi-immersive hybrid user interface’. In: *Virtual Reality* 16.4, pp. 271–288. DOI: `10.1007/s10055-011-0205-y`.

Caserman, Polona, Augusto Garcia-Agundez, Alvar Gámez Zerban and Stefan Göbel (2021). ‘Cybersickness in current-generation virtual reality head-mounted displays: systematic review and outlook’. In: *Virtual Reality*. DOI: `10.1007/s10055-021-00513-6`.

Chang, Joseph Chee, Nathan Hahn, Yongsung Kim, Julina Coupland, Bradley Breneisen, Hannah S. Kim, John Hwong and Aniket Kittur (2021). ‘When the Tab Comes Due:Challenges in the Cost Structure of Browser Tab Usage’. In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM. DOI: `10.1145/3411764.3445585`.

Chekhlov, Denis, Andrew P. Gee, Andrew Calway and Walterio Mayol-Cuevas (2007). ‘Ninja on a Plane: Automatic Discovery of Physical Planes for Augmented Reality Using Visual SLAM’. In: *2007 6th IEEE and ACM International Symposium on Mixed and Augmented Reality*. IEEE. DOI: `10.1109/ismar.2007.4538840`.

Choi, S. H. and H. H. Cheung (2008). ‘A versatile virtual prototyping system for rapid product development’. In: *Computers in Industry* 59.5, pp. 477–488. DOI: `10.1016/j.compind.2007.12.003`.

Chung, Susana T. L., Samuel H. Jarvis and Sing-Hang Cheung (2007). ‘The effect of dioptric blur on reading performance’. In: *Vision Research* 47.12, pp. 1584–1594. DOI: `10.1016/j.visres.2007.03.007`.

Clay, Marie M. (2000). *Running records: for classroom teachers*. English. Portsmouth, NH: Heinemann. ISBN: 9780325002996. URL: `https://www.worldcat.org/oclc/925234124`.

Cockburn, Andy and Bruce McKenzie (2002). ‘Evaluating the effectiveness of spatial memory in 2D and 3D physical and virtual environments’. In: *Proceedings of the SIGCHI conference on Human factors in computing systems Changing our world, changing ourselves - CHI ’02*. CHI ’02. New York, New York, USA: ACM Press, p. 203. ISBN: 1581134533. DOI: `10.1145/503376.503413`.

Cohen, Ellis, Edward Smith and Lee Iverson (1986). ‘Constraint-Based Tiled Windows’. In: *IEEE Computer Graphics and Applications* 6.5, pp. 35–45. ISSN: 0272-1716. DOI: `10.1109/mcg.1986.276790`.

Computer History Museum (2021). *Keyboard with touchpad*. URL: `https://www.computerhistory.org/collections/catalog/102642008`.

Cruz-Neira, C., J. Leigh, M. Papka, C. Barnes, S. M. Cohen, S. Das, R. Engelmann, R. Hudson, T. Roy, L. Siegel, C. Vasilakis, T. A. DeFanti and D. J. Sandin (1993a). 'Scientists in wonderland: A report on visualization applications in the CAVE virtual reality environment'. In: *Proceedings of 1993 IEEE research properties in virtual reality symposium.* IEEE. IEEE Comput. Soc. Press, pp. 59–66. DOI: `10.1109/VRAIS.1993.378262`.

Cruz-Neira, Carolina, Daniel J. Sandin and Thomas A. DeFanti (1993b). 'Surround-screen projection-based virtual reality'. In: *Proceedings of the 20th annual conference on Computer graphics and interactive techniques - SIGGRAPH '93.* SIGGRAPH '93. New York, New York, USA: ACM Press, pp. 135–142. ISBN: 0897916018. DOI: `10.1145/166117.166134`.

Czarny, Eric (2010). *Spectacle.* URL: `https://web.archive.org/web/20201027011414/https://github.com/eczarny/spectacle/releases/tag/0.1`.

Czerwinski, Mary, Greg Smith, Tim Regan, Brian Meyers, George Robertson and Gary Starkweather (2003). 'Toward Characterizing the Productivity Benefits of Very Large Displays'. In: *Human-Computer Interaction INTERACT '03: IFIP TC13 International Conference on Human-Computer Interaction.* Ed. by Matthias Rauterberg, Marino Menozzi, Janet Wesson, Matthias Rauterberg, Marino Menozzi and Janet Wesson, pp. 9–16. URL: `https://www.worldcat.org/oclc/989666219`.

Dam, Andries van (1997). 'Post-WIMP user interfaces'. In: *Communications of the ACM* 40.2, pp. 63–67. DOI: `10.1145/253671.253708`.

Dell Inc. (2021). *Dell USB-C Dock – WD19S.* URL: `https://web.archive.org/web/20210913034356/https://deals.dell.com/en-us/productdetail/afvz`.

Demer, Joseph L. (2006). 'Extraocular muscles'. In: *Duane's Clinical Ophthalmology. Ocular Motility and Strabismus.* Vol. 1. Lippincott Williams & Wilkins. Chap. 1. ISBN: 9780781753296. URL: `http://www.oculist.net/downaton502/prof/ebook/duanes/pages/v1/v1c001.html`.

DesktopPlus (2021). *DesktopPlus.* Ed. by elvissteinjr. URL: `https://github.com/elvissteinjr/DesktopPlus`.

Deurzen, Max van (2019). 'The anatomy of the modern window manager - a case study for X in an Agile manner'. Bachelor's thesis. Radboud University. URL: `https://www.cs.ru.nl/bachelors-theses/2019/Max_van_Deurzen___4581903___The_anatomy_of_the_modern_window_manager_-_a_case_study_for_X_in_an_Agile_manner.pdf`.

Dietz, Paul H., Benjamin Eidelson, Jonathan Westhues and Steven Bathiche (2009). ‘A practical pressure sensitive computer keyboard’. In: *Proceedings of the 22nd annual ACM symposium on User interface software and technology - UIST ’09*. ACM Press. DOI: `10.1145/1622176.1622187`.

Dingler, Tilman, Kai Kunze and Benjamin Outram (2018). ‘VR Reading UIs. Assessing Text Parameters for Reading in VR’. In: *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI EA ’18. New York, NY, USA: ACM, pp. 1–6. ISBN: 9781450356213. DOI: `10.1145/3170427.3188695`.

Dingler, Tilman, Siran Li, Niels van Berkel and Vassilis Kostakos (2020). ‘Page-Turning Techniques for Reading Interfaces in Virtual Environments’. In: *32nd Australian Conference on Human-Computer Interaction*. ACM. DOI: `10.1145/3441000.3441024`.

DiVerdi, Stephen, Daniel Nurmi and Tobias Höllerer (2003). ‘ARWin — A Desktop Augmented Reality Window Manager’. In: *Proceedings of the 2nd IEEE/ACM International Symposium on Mixed and Augmented Reality*. ISMAR ’03. Washington, DC, USA: IEEE Computer Society, pp. 298–. ISBN: 0769520065. DOI: `10.1109/ismar.2003.1240729`.

Dominic, James, Brock Tubre, Charles Ritter, Jada Houser, Colton Smith and Paige Rodeghero (2020). ‘Remote Pair Programming in Virtual Reality’. In: *2020 IEEE International Conference on Software Maintenance and Evolution (ICSME)*. IEEE. DOI: `10.1109/icsme46990.2020.00046`.

Dostal, Jakub, Per Ola Kristensson and Aaron Quigley (2013). ‘Subtle gaze-dependent techniques for visualising display changes in multi-display environments’. In: *Proceedings of the 2013 international conference on Intelligent user interfaces - IUI ’13*. ACM Press. DOI: `10.1145/2449396.2449416`.

Duyck, Marianne, Thérèse Collins and Mark Wexler (2016). ‘Masking the saccadic smear’. In: *Journal of Vision* 16.10, p. 1. DOI: `10.1167/16.10.1`.

Dyson, Mary C. and Mark Haselgrove (2001). ‘The influence of reading speed and line length on the effectiveness of reading from screen’. In: *International Journal of Human-Computer Studies* 54.4, pp. 585–612. DOI: `10.1006/ijhc.2001.0458`.

Elmenreich, Wilfried (2002). *An Introduction to Sensor Fusion*. Research rep. Vienna University of Technology. URL: `https://web.archive.org/web/20210910134100/https://www.researchgate.net/publication/267771481_An_Introduction_to_Sensor_Fusion`.

Encyclopædia Britannica, Inc (2021). *Apple II Computer*. URL: `https://web.archive.org/web/20210517080611/https://www.britannica.com/technology/Apple-II`.

Endert, Alex, Lauren Bradel, Jessica Zeitz, Christopher Andrews and Chris North (2012). 'Designing large high-resolution display workspaces'. In: *Proceedings of the International Working Conference on Advanced Visual Interfaces - AVI '12.* ACM Press. DOI: `10.1145/2254556.2254570`.

Ens, Barrett, Juan David Hincapié-Ramos and Pourang Irani (2014a). 'Ethereal planes'. In: *Proceedings of the 2nd ACM symposium on Spatial user interaction - SUI '14.* SUI '14. New York, New York, USA: ACM Press, pp. 2–12. ISBN: 9781450328203. DOI: `10.1145/2659766.2659769`.

Ens, Barrett, Eyal Ofek, Neil Bruce and Pourang Irani (2016). 'Shared Façades: Surface-Embedded Layout Management for Ad Hoc Collaboration Using Head-Worn Displays'. In: *Collaboration Meets Interactive Spaces.* Springer International Publishing, pp. 153–176. DOI: `10.1007/978-3-319-45853-3_8`.

Ens, Barrett M., Rory Finnegan and Pourang P. Irani (2014b). 'The personal cockpit: A Spatial Interface for Effective Task Switching on Head-Worn Displays'. In: *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14.* CHI '14 August 2016, pp. 3171–3180. DOI: `10.1145/2556288.2557058`.

Evans, Gabriel, Jack Miller, Mariangely Iglesias Pena, Anastacia MacAllister and Eliot Winer (2017). 'Evaluating the Microsoft HoloLens through an augmented reality assembly application'. In: *Degraded Environments: Sensing, Processing, and Display 2017.* Ed. by John (Jack) N. Sanders-Reed and Jarvis (Trey) J. Arthur III. Vol. 10197. International Society for Optics and Photonics. SPIE, pp. 282–297. ISBN: 9781510608955. URL: `https://www.worldcat.org/oclc/1001314625`.

Findlay, John M. (2008). 'Saccadic eye movement programming: sensory and attentional factors'. In: *Psychological Research* 73.2, pp. 127–135. DOI: `10.1007/s00426-008-0201-3`.

Flesch, Rudolf (1979). *How to write plain English : a book for lawyers and consumers.* New York: Harper & Row. ISBN: 9780060112783. URL: `https://www.worldcat.org/oclc/605437305`.

fluxbox, Henrik Kinnunen, Simon Bowden, Mathias Gumz and Mark Tiefenbruck (2021). *fluxbox.* URL: `https://web.archive.org/web/20210828181410/http://fluxbox.org/`.

Forlines, Clifton, Chia Shen, Daniel Wigdor and Ravin Balakrishnan (2006). 'Exploring the effects of group size and display configuration on visual search'. In: *Proceedings of the 2006 20th anniversary conference on Computer supported cooperative work - CSCW '06.* ACM Press. DOI: `10.1145/1180875.1180878`.

Galitz, Wilbert O. (2007). *The Essential Guide to User Interface Design: An Introduction to GUI Design Principles and Techniques.* John Wiley & Sons. 886 pp. ISBN: 9780470053423. URL: `https://www.worldcat.org/oclc/1105801385`.

Gallagher, Kaitlin M., Laura Cameron, Diana De Carvalho and Madison Boulé (2019). 'Does Using Multiple Computer Monitors for Office Tasks Affect User Experience?' In: *Human Factors: The Journal of the Human Factors and Ergonomics Society* 63.3, pp. 433–449. DOI: `10.1177/0018720819889533`.

Gavish, Nirit, Teresa Gutiérrez, Sabine Webel, Jorge Rodríguez, Matteo Peveri, Uli Bockholt and Franco Tecchia (2013). 'Evaluating virtual reality and augmented reality training for industrial maintenance and assembly tasks'. In: *Interactive Learning Environments* 23.6, pp. 778–798. DOI: `10.1080/10494820.2013.815221`.

Gilchrist, Iain D., Valerie Brown and John M. Findlay (1997). 'Saccades without eye movements'. In: *Nature* 390.6656, pp. 130–131. DOI: `10.1038/36478`.

González, Victor M. and Gloria Mark (2004). '"Constant, constant, multitasking craziness"'. In: *Proceedings of the 2004 conference on Human factors in computing systems - CHI '04.* ACM Press. DOI: `10.1145/985692.985707`.

Goodfellow, Michael (1986). 'WHIM, the Window Handler and Input Manager'. In: *IEEE Computer Graphics and Applications* 6.5, pp. 46–52. ISSN: 0272-1716. DOI: `10.1109/mcg.1986.276791`.

Gordon, Claire C., Cynthia L. Blackwell, Bruce Bradtmiller, Joseph L. Parham, Patricia Barrientos, Stephen P. Paquette, Brian D. Corner, Jeremy M. Carson, Joseph C. Venezia, Belva M. Rockwell et al. (2014). *2012 Anthropometric Survey of U.S. Army Personnel: Methods and Summary Statistics.* Tech. rep. NATICK/TR-15/007. Army Natick Soldier Research Development and Engineering Center MA. URL: `https://www.hsdl.org/?view&did=762624`.

Gourlay, Michael J. and Robert T. Held (2017). 'Head-Mounted-Display Tracking for Augmented and Virtual Reality'. In: *Information Display* 33.1, pp. 6–10. DOI: `10.1002/j.2637-496x.2017.tb00962.x`.

Grudin, Jonathan (2001). 'Partitioning digital worlds'. In: *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '01.* ACM Press. DOI: `10.1145/365024.365312`.

Gugerty, Leo, Richard A. Tyrrell, Thomas R. Aten and K. Andy Edmonds (2004). 'The effects of subpixel addressing on users' performance and preferences during reading-related tasks'. In: *ACM Transactions on Applied Perception* 1.2, pp. 81–101. DOI: `10.1145/1024083.1024084`.

Gunkel, Simon N. B., Hans M. Stokking, Martin J. Prins, Nanda van der Stap, Frank B. ter Haar and Omar A. Niamut (2018). ‘Virtual reality conferencing’. In: *Proceedings of the 9th ACM Multimedia Systems Conference*. ACM. DOI: `10.1145/3204949.3208115`.

Haigh, Thomas and Paul E. Ceruzzi (2021). *A New History of Modern Computing*. Cambridge, Massachusetts: The MIT Press. ISBN: 9780262542906. URL: `https://www.worldcat.org/oclc/1204265186`.

Hembrooke, Helene and Geri Gay (2003). ‘The laptop and the lecture: The effects of multitasking in learning environments’. In: *Journal of Computing in Higher Education* 15.1, pp. 46–64. DOI: `10.1007/bf02940852`.

Henderson Jr, D. Austin and Stuart Card (1986). ‘Rooms: the use of multiple virtual workspaces to reduce space contention in a window-based graphical user interface’. In: *ACM Transactions on Graphics (TOG)* 5.3, pp. 211–243. ISSN: 0730-0301. DOI: `10.1145/24054.24056`.

Hinckley, Ken, Randy Pausch, John C. Goble and Neal F. Kassell (1994). ‘A survey of design issues in spatial input’. In: *Proceedings of the 7th annual ACM symposium on User interface software and technology - UIST ’94*. ACM Press. DOI: `10.1145/192426.192501`.

Hogue, A., M. R. Jenkin and R. S. Allison (2004). ‘An optical-inertial tracking system for fully-enclosed VR displays’. In: *First Canadian Conference on Computer and Robot Vision, 2004. Proceedings.* IEEE. DOI: `10.1109/cccrv.2004.1301417`.

Hutchings, Dugald Ralph, Greg Smith, Brian Meyers, Mary Czerwinski and George Robertson (2004). ‘Display Space Usage and Window Management Operation Comparisons between Single Monitor and Multiple Monitor Users’. In: *Proceedings of the Working Conference on Advanced Visual Interfaces (AVI ’04)*, pp. 32–39. DOI: `10.1145/989863.989867`.

Hutchings, Dugald Ralph and John Stasko (2004a). ‘Revisiting Display Space Management: Understanding Current Practice to Inform Next-generation Design’. In: *Proceedings of Graphics Interface 2004*. Vol. 62. GI ’04. London, Ontario, Canada, pp. 127–134. ISBN: 9781568812274. URL: `https://www.worldcat.org/oclc/60173658`.

Hutchings, Dugald Ralph and John Stasko (2004b). ‘Shrinking window operations for expanding display space’. In: *Proceedings of the working conference on Advanced visual interfaces - AVI ’04*, p. 350. DOI: `10.1145/989863.989922`.

Immersed Inc. (2021). *Immersed*. URL: `https://immersed.com/`.

Inhoff, Albrecht W., Seth N. Greenberg, Matthew Solomon and Chin-An Wang (2009). ‘Word integration and regression programming during reading: A test of the E-Z reader 10 model.’ In: *Journal of Experimental Psychology: Human Perception and Performance* 35.5, pp. 1571–1584. DOI: `10.1037/a0014250`.

Inhoff, Albrecht W., Ulrich W. Weger and Ralph Radach (2005). ‘Sources of information for the programming of short-and long-range regressions during reading’. In: *Cognitive processes in eye guidance*. Ed. by Geoffrey D. M. Underwood. Oxford University Press. Chap. 2, pp. 33–52. ISBN: 9780198566809. URL: `https://www.worldcat.org/oclc/783310149`.

Inman, Chris, Vivian H. Wright and Julia A. Hartman (2010). ‘Use of Second Life in K-12 and higher education: A review of research.’ In: *Journal of Interactive Online Learning* 9.1, pp. 44–63. ISSN: 1541-4914.

Jacob, Robert J. K., Audrey Girouard, Leanne M. Hirshfield, Michael S. Horn, Orit Shaer, Erin Treacy Solovey and Jamie Zigelbaum (2008). ‘Reality-based interaction’. In: *Proceeding of the twenty-sixth annual CHI conference on Human factors in computing systems - CHI ’08*. ACM Press. DOI: `10.1145/1357054.1357089`.

Jainta, Stephanie, Anne Dehnert, Sven P. Heinrich and Wolfgang Jaschinski (2011). ‘Binocular Coordination during Reading of Blurred and Nonblurred Text’. In: *Investigative Opthalmology & Visual Science* 52.13, p. 9416. DOI: `10.1167/iovs.11-8237`.

Jakobsen, Mikkel Rønne and Kasper Hornbæk (2010). ‘Piles, tabs and overlaps in navigation among documents’. In: *Proceedings of the 6th Nordic Conference on Human-Computer Interaction Extending Boundaries - NordiCHI ’10*. ACM Press. DOI: `10.1145/1868914.1868945`.

Jalaliniya, Shahram, Diako Mardanbeigi, Thomas Pederson and Dan Witzner Hansen (2014). ‘Head and eye movement as pointing modalities for eyewear computers’. In: *Proceedings - 11th International Conference on Wearable and Implantable Body Sensor Networks Workshops, BSN Workshops 2014*. ISBN: 9781479949328. DOI: `10.1109/BSN.Workshops.2014.14`.

Jankowski, Jacek, Krystian Samp, Izabela Irzynska, Marek Jozwowicz and Stefan Decker (2010). ‘Integrating text with video and 3D graphics’. In: *Proceedings of the 28th international conference on Human factors in computing systems - CHI ’10*. ACM Press, pp. 1321–1330. DOI: `10.1145/1753326.1753524`.

Jetter, Hans-Christian, Harald Reiterer and Florian Geyer (2013). ‘Blended Interaction: understanding natural human–computer interaction in post-WIMP interactive spaces’. In: *Personal and Ubiquitous Computing* 18.5, pp. 1139–1158. DOI: `10.1007/s00779-013-0725-4`.

Jeuris, Steven, Paolo Tell, Steven Houben and Jakob E. Bardram (2018). 'The Hidden Cost of Window Management'. In: arXiv: `1810.04673 [cs.HC]`.

Jo, Jaemin, Sehi L'Yi, Bongshin Lee and Jinwook Seo (2017). 'TouchPivot'. In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM. DOI: `10.1145/3025453.3025752`.

Johnson, Micah K. and Hany Farid (2006). 'Exposing digital forgeries through chromatic aberration'. In: *Proceeding of the 8th workshop on Multimedia and security - MM&Sec '06*. ACM Press. DOI: `10.1145/1161366.1161376`.

Juarez, Alex, Willem Schonenberg and Christoph Bartneck (2010). 'Implementing a low-cost CAVE system using the CryEngine2'. In: *Entertainment Computing* 1.3-4, pp. 157–164. DOI: `10.1016/j.entcom.2010.10.001`.

Kar, Gourab, Amy Vu, Begoña Juliá Nehme and Alan Hedge (2015). 'Effects of Mouse, Trackpad and 3d Motion and Gesture Control on Performance, Posture, and Comfort'. In: *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 59.1, pp. 327–331. DOI: `10.1177/1541931215591068`.

Kida, K., M. Ihara, S. Shiwa and S. Ishibashi (2000). 'Motion tracking method for the CAVE$^{\text{TM}}$ system'. In: *WCC 2000 - ICSP 2000. 2000 5th International Conference on Signal Processing Proceedings. 16th World Computer Congress 2000*. IEEE. DOI: `10.1109/icosp.2000.891647`.

Kim, Eunjee and Gwanseob Shin (2018). 'Head Rotation and Muscle Activity When Conducting Document Editing Tasks with a Head-Mounted Display'. In: *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 62.1, pp. 952–955. DOI: `10.1177/1541931218621219`.

Kim, Eunjee and Gwanseob Shin (2021). 'User discomfort while using a virtual reality headset as a personal viewing system for text-intensive office tasks'. In: *Ergonomics* 64.7, pp. 891–899. DOI: `10.1080/00140139.2020.1869320`.

Klinke, Hermann, Christoph Krieger and Sebastian Pickl (2014). *Analysis of the influence of screen size and resolution on work efficiency*. en. DOI: `10.18419/OPUS-3489`.

Kojić, Tanja, Danish Ali, Robert Greinacher, Sebastian Möller and Jan-Niklas Voigt-Antons (2020). 'User Experience of Reading in Virtual Reality – Finding Values for Text Distance, Size and Contrast'. In: arXiv: `2004.01545 [cs.MM]`.

Kraft, Caleb and Make Community LLC (2016). *Exclusive: See the Secret Prototypes We Found in Valve's VR Lab*. URL: `https://web.archive.org/web/20201108125335/https://makezine.com/2016/06/21/exclusive-see-the-secret-prototypes-we-found-in-valves-vr-lab/`.

Kulkarni, Minchu, Kshitij Kapoor, Deva Surya Vivek Madala, Sanchit Bansal and Sudheendra Hangal (2019). 'Compartmentalizing web browsing with Sailboat'. In: *Proceedings of the 10th Indian Conference on Human-Computer Interaction*. ACM. DOI: `10.1145/3364183.3364187`.

Kyritsis, Markos, Stephen R. Gulliver and Eva Feredoes (2016). 'Environmental factors and features that influence visual search in a 3D WIMP interface'. In: *International Journal of Human-Computer Studies* 92-93, pp. 30–43. DOI: `10.1016/j.ijhcs.2016.04.009`.

Kytö, Mikko, Barrett Ens, Thammathip Piumsomboon, Gun A. Lee and Mark Billinghurst (2018). 'Pinpointing : Precise Head - and Eye - Based Target Selection for Augmented Reality'. In: *Proc. of CHI*. ISSN: 1098-0121. DOI: `10.1145/3173574.3173655`.

Labbe, David R., Kean Kouakoua, Rachid Aissaoui, Sylvie Nadeau and Cyril Duclos (2021). 'Proprioceptive Stimulation Added to a Walking Self-Avatar Enhances the Illusory Perception of Walking in Static Participants'. In: *Frontiers in Virtual Reality* 2. DOI: `10.3389/frvir.2021.557783`.

Larsen, C. R., J. L. Soerensen, T. P. Grantcharov, T. Dalsgaard, L. Schouenborg, C. Ottosen, T. V. Schroeder and B. S. Ottesen (2009). 'Effect of virtual reality training on laparoscopic surgery: randomised controlled trial'. In: *BMJ* 338.may14 2, b1802–b1802. DOI: `10.1136/bmj.b1802`.

Larson, Kevin, Maarten van Dantzich, Mary Czerwinski and George Robertson (2000). 'Text in 3D'. In: *CHI '00 extended abstracts on Human factors in computing systems - CHI '00*. ACM Press. DOI: `10.1145/633292.633374`.

Laukkanen, Joona Antero (2011). 'A scalable and tiling multi-monitor aware window manager'. In: *Proceedings of the 2011 annual conference extended abstracts on Human factors in computing systems - CHI EA '11*. ACM Press. DOI: `10.1145/1979742.1979515`.

Lee, Haejin (2017). *Taking the Leap from 2D to 3D*. URL: `https://web.archive.org/web/20201111161901/https://medium.com/microsoft-design/taking-the-leap-from-2d-to-3d-4199ccc777a1`.

Leykin, A. and M. Tuceryan (2004). 'Automatic Determination of Text Readability over Textured Backgrounds for Augmented Reality Systems'. In: *Third IEEE and ACM International Symposium on Mixed and Augmented Reality*. IEEE. DOI: `10.1109/ismar.2004.22`.

Ling, Chen, Alex Stegman, Chintan Barhbaya and Randa Shehab (2016). 'Are Two Better Than One? A Comparison Between Single- and Dual-Monitor Work Stations in Productivity and User's Windows Management Style'. In: *International Journal of Human–Computer Interaction* 33.3, pp. 190–198. DOI: `10.1080/10447318.2016.1231392`.

Liu, Min, Tomoko Traphagan, Jin Huh, Young Inh Koh, Cilok Choi and Allison McGregor (2008). 'Designing Websites for ESL Learners: A Usability Testing Study'. In: *Calico Journal* 25.2, pp. 207–240.

Luckey, Palmer (2019). *I can't use Rift S, and neither can you.* URL: `https://web.archive.org/web/20210816180230/https://palmerluckey.com/i-cant-use-rift-s-and-neither-can-you/`.

Luu, Wilson, Barbara Zangerl, Michael Kalloniatis and Juno Kim (2021). 'Effects of stereopsis on vection, presence and cybersickness in head-mounted display (HMD) virtual reality'. In: *Scientific Reports* 11.1. DOI: `10.1038/s41598-021-89751-x`.

Mann, Steve (2002). 'Mediated reality with implementations for everyday life'. In: *Presence Connect* 1, p. 2002.

Mann, Steve, Tom Furness, Yu Yuan, Jay Iorio and Zixin Wang (2018). 'All Reality: Virtual, Augmented, Mixed (X), Mediated (X,Y), and Multimediated Reality'. In: arXiv: `1804.08386 [cs.HC]`.

Mark, Gloria, Daniela Gudith and Ulrich Klocke (2008). 'The cost of interrupted work'. In: *Proceeding of the twenty-sixth annual CHI conference on Human factors in computing systems - CHI '08*. ACM Press. DOI: `10.1145/1357054.1357072`.

Marks, Stefan, Javier E. Estevez and Andy M. Connor (2014). 'Towards the Holodeck'. In: *Proceedings of the 29th International Conference on Image and Vision Computing New Zealand - IVCNZ '14*. ACM Press. DOI: `10.1145/2683405.2683424`.

McCleary, David S. (2009). *The Optician Training Manual.* English. First. Temecula, CA: Santa Rosa Publishing, p. 278. ISBN: 9780615193816. URL: `https://www.worldcat.org/oclc/695901751`.

McGill, Mark, Aidan Kehoe, Euan Freeman and Stephen Brewster (2020). 'Expanding the Bounds of Seated Virtual Workspaces'. In: *ACM Transactions on Computer-Human Interaction* 27.3, pp. 1–40. DOI: `10.1145/3380959`.

McKenzie, Chris, Adam Glazier and Google LLC (2017). *Designing Screen Interfaces for VR.* URL: `https://www.youtube.com/watch?v=ES9jArHRFHQ`.

Mechdyne Corporation (2021a). *CAVE Immersive Virtual Reality.* URL: `https://web.archive.org/web/20210628002614/https://www.mechdyne.com/av-vr-solutions/solutions/virtual-augmented-reality/cave/`.

Mechdyne Corporation (2021b). *CAVE2 Immersive Collaboration Space.* URL: `https://web.archive.org/web/20210227182728/https://www.mechdyne.com/av-vr-solutions/solutions/virtual-augmented-reality/cave2/`.

Microsoft Corporation (2001). *ClearType.* URL: `http://research.microsoft.com/en-us/projects/ClearType/`.

Microsoft Corporation (2013). *Ways to capture the screen.* URL: `https://web.archive.org/web/20200719210943/https://docs.microsoft.com/en-us/archive/blogs/dsui_team/ways-to-capture-the-screen`.

Microsoft Corporation (2018a). *Desktop Window Manager.* URL: `https://web.archive.org/web/20210519230830/https://docs.microsoft.com/en-us/windows/win32/dwm/dwm-overview`.

Microsoft Corporation (2018b). *DWM Thumbnail Overview.* Ed. by Michael Satran, Mike Jacobs and Steven White. URL: `https://web.archive.org/web/20201112034933/docs.microsoft.com/en-us/windows/win32/dwm/thumbnail-ovw`.

Microsoft Corporation (2020). *Start designing and prototyping – Mixed Reality.* URL: `https://web.archive.org/web/20210524141704/https://docs.microsoft.com/en-us/windows/mixed-reality/design/design`.

Microsoft Corporation (2021a). *Coarse relocalization.* URL: `https://web.archive.org/web/20210525191022/https://docs.microsoft.com/en-us/azure/spatial-anchors/concepts/coarse-reloc`.

Microsoft Corporation (2021b). *How to use multiple monitors in Windows 10.* URL: `https://support.microsoft.com/en-us/windows/how-to-use-multiple-monitors-in-windows-10-329c6962-5a4d-b481-7baa-bec9671f728a`.

Microsoft Corporation (2021c). *How to use the taskbar in Windows 10.* URL: `%7Bhttps://support.microsoft.com/en-us/windows/how-to-use-the-taskbar-in-windows-11-0657a50f-0cc7-dbfd-ae6b-05020b195b07%7D`.

Microsoft Corporation (2021d). *PowerToys FancyZones utility for Windows 10.* URL: `https://web.archive.org/web/20210901123341/https://docs.microsoft.com/en-us/windows/powertoys/fancyzones`.

Milgram, Paul, Haruo Takemura, Akira Utsumi and Fumio Kishino (1995). 'Augmented reality: a class of displays on the reality-virtuality continuum'. In: *Telemanipulator and Telepresence Technologies.* Ed. by Hari Das. SPIE. DOI: `10.1117/12.197321`.

Mizage, LLC (2019). *WinDivvy.* URL: `https://web.archive.org/web/20210817160956/https://mizage.com/windivvy/`.

Mobile Pixels (2021). *DUEX Collection.* URL: `https://web.archive.org/web/20210811051128/https://www.mobilepixels.us/pages/duex-series`.

Moeslund, Thomas B. and Erik Granum (2001). 'A Survey of Computer Vision-Based Human Motion Capture'. In: *Computer Vision and Image Understanding* 81.3, pp. 231–268. DOI: `10.1006/cviu.2000.0897`.

Montello, Daniel R, David Waller, Mary Hegarty and Anthony E Richardson (2004). ‘Spatial memory of real environments, virtual environments, and maps’. In: *Human spatial memory.* Psychology Press. Chap. 11, pp. 271–306. ISBN: 9781410609984. URL: `https://www.worldcat.org/oclc/55079964`.

Murphy, Glen (2011). ‘Post-PC devices: A summary of early iPad technology adoption in tertiary environments’. In: *E-Journal of Business Education and Scholarship of Teaching* 5.1, pp. 18–32. ISSN: 1835-9132. URL: `https://eprints.qut.edu.au/44085/`.

Murray, Jeff W. (2017). *Building Virtual Reality with Unity and Steam VR.* Milton: Chapman and Hall/CRC. ISBN: 9781315305455. URL: `https://www.worldcat.org/oclc/1225366468`.

Myers, Brad A. (1988). ‘A taxonomy of window manager user interfaces’. In: *IEEE Computer Graphics and Applications* 8.5, pp. 65–84. DOI: `10.1109/38.7762`.

Ni, Tao, Doug A. Bowman and Jian Chen (2006). ‘Increased display size and resolution improve task performance in information-rich virtual environments’. In: *GI ’06: Proceedings of Graphics Interface 2006.* Canadian Information Processing Society. Toronto, Ontario, Canada: Mississauga Canadian Information Processing Society, pp. 139–146. ISBN: 9781568813080. DOI: `10.5555/1143079.1143102`. URL: `https://www.worldcat.org/oclc/255543153`.

Nimbarte, Ashish D., Rabab T. Alabdulmohsen, Steve E. Guffey and John R. Etherton (2013). ‘The Impact of Use of Dual Monitor Screens on 3D Head–Neck Posture and Activity of Neck Muscles’. In: *IIE Transactions on Occupational Ergonomics and Human Factors* 1.1, pp. 38–49. DOI: `10.1080/21577323.2012.667051`.

Noyes, Jan (1983). ‘The QWERTY keyboard: a review’. In: *International Journal of Man-Machine Studies* 18.3, pp. 265–281. DOI: `10.1016/s0020-7373(83)80010-8`.

Owens, Justin W., Jennifer Teves, Bobby Nguyen, Amanda Smith, Mandy C. Phelps and Barbara S. Chaparro (2012). ‘Examination of Dual vs. Single Monitor Use during Common Office Tasks’. In: *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 56.1, pp. 1506–1510. DOI: `10.1177/1071181312561299`.

Park, Sungryul, Donghee Choi, Jihhyeon Yi, Songil Lee, Ja Eun Lee, Byeonghwa Choi, Seungbae Lee and Gyouhyung Kyung (2017). ‘Effects of display curvature, display zone, and task duration on legibility and visual fatigue during visual search task’. In: *Applied Ergonomics* 60, pp. 183–193. DOI: `10.1016/j.apergo.2016.11.012`.

Parker, Adam J. and Timothy J. Slattery (2019). 'Word frequency, predictability, and return-sweep saccades: Towards the modeling of eye movements during paragraph reading.' In: *Journal of Experimental Psychology: Human Perception and Performance* 45.12, pp. 1614–1633. DOI: `10.1037/xhp0000694`.

Parker, Adam J., Timothy J. Slattery and Julie A. Kirkby (2019). 'Return-sweep saccades during reading in adults and children'. In: *Vision Research* 155, pp. 35–43. DOI: `10.1016/j.visres.2018.12.007`.

Parks, Marshall M. (2006). 'Eye Movements and Positions'. In: *Duane's Clinical Ophthalmology. Ocular Motility and Strabismus.* Vol. 1. Lippincott Williams & Wilkins. Chap. 2. ISBN: 9780781753296. URL: `http://www.oculist.net/downaton502/prof/ebook/duanes/pages/v1/v1c002.html`.

Paterson, Donald G. and Miles A. Tinker (1942). 'Influence of line width on eye movements for six-point type.' In: *Journal of Educational Psychology* 33.7, p. 552.

Pavanatto, Leonardo, Chris North, Doug A. Bowman, Carmen Badea and Richard Stoakley (2021). 'Do we still need physical monitors? An evaluation of the usability of AR virtual monitors for productivity work'. In: *2021 IEEE Virtual Reality and 3D User Interfaces (VR).* IEEE. DOI: `10.1109/vr50410.2021.00103`.

Phenix, Samantha J. (2017). 'Display Technology Trends in Digital Signage Displays'. In: *SID Symposium Digest of Technical Papers* 48.1, pp. 608–611. DOI: `10.1002/sdtp.11712`.

Pirchheim, Christian, Manuela Waldner and Dieter Schmalstieg (2009). 'Deskotheque: Improved Spatial Awareness in Multi-Display Environments'. In: *2009 IEEE Virtual Reality Conference.* Lafayette, LA, USA: IEEE, pp. 123–126. DOI: `10.1109/vr.2009.4811010`.

Proudlock, Frank Antony, Himanshu Shekhar and Irene Gottlob (2003). 'Coordination of Eye and Head Movements during Reading'. In: *Investigative ophthalmology & visual science* 44.7, pp. 2991–2998. DOI: `10.1167/iovs.02-1315`.

Purves, Dale, George J. Augustine, David Fitzpatrick, William C. Hall, Anthony-Samuel LaMantia, Richard D. Mooney, Michael L. Platt and Leonard E. White (2018). *Neuroscience.* 6th ed. New York: Oxford University Press. ISBN: 9781605353807. URL: `https://www.worldcat.org/oclc/1026967719`.

Pygmal Technologies (2017). *SPACE.* URL: `%7Bhttps://web.archive.org/web/20170512051822/https://getspace.io/%7D`.

Rakver, Martin (2020). *VR Display Comparison – LCD vs OLED, (Sub)Pixels, FOV & PPD.* URL: `https://web.archive.org/web/20210201203925/https://smartglasseshub.com/vr-headset-display-comparison/`.

Raptis, Dimitrios, Nikolaos Tselios, Jesper Kjeldskov and Mikael B. Skov (2013). 'Does size matter?' In: *Proceedings of the 15th international conference on Human-computer interaction with mobile devices and services - MobileHCI '13*. ACM Press. DOI: `10.1145/2493190.2493204`.

Rau, Pei-Luen Patrick, Jian Zheng, Zhi Guo and Jiaqi Li (2018). 'Speed reading on virtual reality and augmented reality'. In: *Computers & Education* 125, pp. 240–245. DOI: `10.1016/j.compedu.2018.06.016`.

Rayner, Keith (1998). 'Eye movements in reading and information processing: 20 years of research.' In: *Psychological bulletin* 124.3, pp. 372–422. ISSN: 1939-1455. DOI: `10.1037/0033-2909.124.3.372`.

Rayner, Keith, Tessa Warren, Barbara J. Juhasz and Simon P. Liversedge (2004). 'The Effect of Plausibility on Eye Movements in Reading.' In: *Journal of Experimental Psychology: Learning, Memory, and Cognition* 30.6, pp. 1290–1301. DOI: `10.1037/0278-7393.30.6.1290`.

Razer Inc (2021). *Razer Huntsman V2 Analog*. URL: `https://web.archive.org/web/20210514193957/https://www.razer.com/gaming-keyboards/razer-huntsman-v2-analog/RZ03-03610200-R3U1`.

Reichlen, B. A. (1993). 'Sparcchair: A one hundred million pixel display'. In: *Proceedings of IEEE Virtual Reality Annual International Symposium*. Los Alamitos, CA, USA: IEEE, pp. 300–307. ISBN: 0780313631. DOI: `10.1109/VRAIS.1993.380765`.

Ringel, Meredith (2003). 'When one isn't enough'. In: *CHI '03 extended abstracts on Human factors in computing systems - CHI '03*. CHI EA '03. New York, New York, USA: ACM Press, p. 762. ISBN: 1581136374. DOI: `10.1145/765891.765976`.

Ritger, Andy (2006). *Using the ExistingXFree86/X.Org LoadableDriver Framework to Achievea Composited X Desktop*. Tech. rep. NVIDIA Corporation. URL: `https://web.archive.org/web/20210308033624/https://download.nvidia.com/developer/presentations/2006/xdevconf/compositing-with-current-framework.pdf`.

Riva, Giuseppe and Luciano Gamberini (2000). 'Virtual Reality in Telemedicine'. In: *Telemedicine Journal and e-Health* 6.3, pp. 327–340. DOI: `10.1089/153056200750040183`.

Robertson, G., M. Czerwinski, P. Baudisch, B. Meyers, D. Robbins, G. Smith and Desney Tan (2005). 'The Large-Display User Experience'. In: *IEEE Computer Graphics and Applications* 25.4, pp. 44–51. DOI: `10.1109/mcg.2005.88`.

Robertson, George, Mary Czerwinski, Kevin Larson, Daniel C. Robbins, David Thiel and Maarten van Dantzich (1998). 'Data mountain'. In: *Proceedings of the 11th annual ACM symposium on User interface software and technology - UIST '98*. New York, New York, USA: ACM Press, pp. 153–162. ISBN: 1581130341. DOI: `10.1145/288392.288596`.

Robertson, George, Maarten van Dantzich, Daniel Robbins, Mary Czerwinski, Ken Hinckley, Kirsten Risden, David Thiel and Vadim Gorokhovsky (2000). 'The Task Gallery: A 3D Window Manager'. In: *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '00*. CHI '00. New York, New York, USA: ACM Press, pp. 494–501. ISBN: 1581132166. DOI: `10.1145/332040.332482`.

Robinett, Warren and Jannick P. Rolland (1993). 'A Computational Model for the Stereoscopic Optics of a Head-Mounted Display'. In: *Virtual Reality Systems*. Elsevier, pp. 51–75. DOI: `10.1016/b978-0-12-227748-1.50013-5`.

Rubinstein, Joshua S., David E. Meyer and Jeffrey E. Evans (2001). 'Executive control of cognitive processes in task switching.' In: *Journal of Experimental Psychology: Human Perception and Performance* 27.4, pp. 763–797. DOI: `10.1037/0096-1523.27.4.763`.

Rzayev, Rufat, Polina Ugnivenko, Sarah Graf, Valentin Schwind and Niels Henze (2021). 'Reading in VR: The Effect of Text Presentation Type and Location'. In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM. DOI: `10.1145/3411764.3445606`.

Sadiq, Mazhar and Antti Pirhonen (2014). 'Usability in ERP (Single and Multiple Document Interface) Application Environments'. In: *International Journal of Business, Humanities and Technology* 4.4, pp. 75–80. ISSN: 2162-1357. URL: `https://jyx.jyu.fi/handle/123456789/45324`.

Sarma, Kalluri R. and Tayo Akinwande (1996). 'Flat Panel Displays for Portable Systems'. In: *Technologies for Wireless Computing*. Ed. by Anantha P. Chandrakasan and Robert W. Brodersen. Vol. 13. 2. Springer, pp. 165–190. ISBN: 9781461286332. DOI: `10.1007/978-1-4613-1453-0_6`.

Savidis, Anthony and Andreas Maragudakis (2013). 'Mixing Geometrically Diverse Window Managers'. In: *Advances in Visual Computing*. Springer Berlin Heidelberg, pp. 714–724. DOI: `10.1007/978-3-642-41939-3_70`.

Sharma, Vibhu Saujanya, Rohit Mehra, Vikrant Kaulgud and Sanjay Podder (2019). 'An Extended Reality Approach for Creating Immersive Software Project Workspaces'. In: *2019 IEEE/ACM 12th International Workshop on Cooperative and Human Aspects of Software Engineering (CHASE)*. IEEE. DOI: `10.1109/chase.2019.00013`.

Sheedy, Jim, Yu-Chi Tai, Manoj Subbaram, Sowjanya Gowrisankaran and John Hayes (2008). 'ClearType sub-pixel text rendering: Preference, legibility and reading performance'. In: *Displays*. Health and Safety Aspects of Visual Displays 29.2, pp. 138–151. ISSN: 0141-9382. DOI: `10.1016/j.displa.2007.09.016`.

Shibata, T., J. Kim, D. M. Hoffman and M. S. Banks (2011). 'The zone of comfort: Predicting visual discomfort with stereo displays'. In: *Journal of Vision* 11.8, pp. 11–11. DOI: `10.1167/11.8.11`.

Shin, Y., H. W. Lim, M. H. Kang, M. Seong, H. Cho and J. H. Kim (2016). 'Normal range of eye movement and its relationship to age'. In: *Acta Ophthalmologica* 94. DOI: `10.1111/j.1755-3768.2016.0499`.

Shupp, Lauren, Christopher Andrews, Margaret Dickey-Kurdziolek, Beth Yost and Chris North (2009). 'Shaping the Display of the Future: The Effects of Display Size and Curvature on User Performance and Insights'. In: *Human–Computer Interaction* 24.1-2, pp. 230–272. DOI: `10.1080/07370020902739429`.

Shupp, Lauren, Robert Ball, Beth Yost, John Booker and Chris North (2006). 'Evaluation of viewport size and curvature of large, high-resolution displays.' In: *GI '06: Proceedings of Graphics Interface 2006*. Canadian Information Processing Society. Toronto, Ontario, Canada: Mississauga Canadian Information Processing Society, pp. 123–130. ISBN: 9781568813080. DOI: `10.5555/1143079.1143100`. URL: `https://www.worldcat.org/oclc/255543153`.

Skarbez, Richard, Missie Smith and Mary C. Whitton (2021). 'Revisiting Milgram and Kishino's Reality-Virtuality Continuum'. In: *Frontiers in Virtual Reality* 2. DOI: `10.3389/frvir.2021.647997`.

Slater, Mel, Martin Usoh and Anthony Steed (1995). 'Taking steps: The Influence of a Walking Technique on Presence in Virtual Reality'. In: *ACM Transactions on Computer-Human Interaction* 2.3, pp. 201–219. DOI: `10.1145/210079.210084`.

Stadler, Alexandra and Haik Lorenz (2007). 'Graphical User Interfaces'. In: *Wiley Encyclopedia of Computer Science and Engineering*, pp. 1464–1476. DOI: `10.1002/9780470050118.ecse171`.

Statcounter Global Stats (2021). *Desktop vs Mobile vs Tablet Market Share New Zealand*. URL: `https://gs.statcounter.com/platform-market-share/desktop-mobile-tablet/new-zealand/`.

Steuer, Jonathan (1992). 'Defining Virtual Reality: Dimensions Determining Telepresence'. In: *Journal of Communication* 42.4, pp. 73–93. DOI: `10.1111/j.1460-2466.1992.tb00812.x`.

Svidt, Kjeld and Bjarne Bjerg (2002). 'Visualisation of CFD Results in a Virtual Reality Environment'. English. In: *ROOMVENT 2002*. Ed. by A. K. Melikov and P. V. Nielsen. Copenhagen, Denmark: Technical University of Denmark (DTU), pp. 77–80. ISBN: 9788798911708. URL: `https://www.worldcat.org/oclc/471534325`.

Tan, Desney S. and Mary Czerwinski (2003). 'Effects of Visual Separation and Physical Discontinuities when Distributing Information across Multiple Displays.' In: *Human-Computer Interaction INTERACT '03: IFIP TC13 International Conference on Human-Computer Interaction*. Ed. by Matthias Rauterberg, Marino Menozzi, Janet Wesson, Matthias Rauterberg, Marino Menozzi and Janet Wesson, pp. 252–255. URL: `https://www.worldcat.org/oclc/989666219`.

Tan, Desney S., Darren Gergle, Peter Scupelli and Randy Pausch (2003). 'With similar visual angles, larger displays improve spatial performance'. In: *Proceedings of the conference on Human factors in computing systems - '03*. ACM Press. DOI: `10.1145/642611.642650`.

Tan, Desney S., Brian Meyers and Mary Czerwinski (2004). 'WinCuts: Manipulating Arbitrary Window Regions for More Effective Use of Screen Space'. In: *Extended abstracts of the 2004 conference on Human factors and computing systems - CHI '04*. ACM Press. DOI: `10.1145/985921.986106`.

The Parliament of New Zealand (2013). *Financial Reporting Act 2013*. Ed. by Ministry of Business, Innovation and Employment (MBIE). URL: `https://web.archive.org/web/20210127081205/https://www.legislation.govt.nz/act/public/2013/0101/latest/DLM4632829.html`.

The Trustees of Indiana University (2021). *For Mac OS X, what are Quartz and Quartz Extreme?* URL: `https://kb.iu.edu/d/ajeb`.

Thibos, Larry N., Arthur Bradley, Tao Liu and Norberto López-Gil (2013). 'Spherical Aberration and the Sign of Defocus'. In: *Optometry and Vision Science* 90.11, pp. 1284–1291. DOI: `10.1097/opx.0000000000000040`.

Tomitsch, Martin (2004). 'Trends and Evolution of Window Interfaces - A Discussion on Interfaces and Interaction Techniques for Window management'. Vienna University of Technology.

Tortora, Gerard J. and Bryan H. Derrickson (2009). *Principles of Anatomy and Physiology*. 12th ed. Hoboken, NJ: John Wiley & Sons. ISBN: 9780470084717. URL: `https://www.worldcat.org/oclc/971996505`.

Truemper, Jacob M., Hong Sheng, Michael G. Hilgers, Richard H. Hall, Morris Kalliny and Basanta Tandon (2008). 'Usability in multiple monitor displays'. In: *ACM SIGMIS Database: the DATABASE for Advances in Information Systems* 39.4, pp. 74–89. DOI: `10.1145/1453794.1453802`.

VALVE Corporation (2021). *Steam Hardware & Software Survey.* URL: `https://store.steampowered.com/hwsurvey/`.

Virtual Desktop, Inc. (2021). *Virtual Desktop.* URL: `https://www.vrdesktop.net/`.

Visbox Inc (2021). *VisCube M4, M5.* URL: `https://web.archive.org/web/20210302233332/http://www.visbox.com/products/cave/viscube-m4/`.

Vitu, Françoise and George W. McConkie (2000). ‘Regressive Saccades and Word Perception in Adult Reading’. In: *Reading as a perceptual process.* Ed. by Alan Kennedy, Ralph Radach, Dieter Heller and Joël Pynte, pp. 301–326. DOI: `10.1016/B978-008043642-5/50015-2`.

VRdudeBRO (2016). *Oculus Rift DK2 CV1 Side By Side Lens View.* URL: `https://www.youtube.com/watch?v=5qxAk3QQQDQ`.

Wagner, Julie, Wendy E. Mackay and Stéphane Huot (2012). ‘Left-over Windows Cause Window Clutter... But What Causes Left-over Windows?’ In: *Proceedings of the 2012 Conference on Ergonomie et Interaction homme-machine - Ergo’IHM ’12.* ACM Press. DOI: `10.1145/2652574.2653442`.

Wallace, James R., Daniel Vogel and Edward Lank (2014). ‘Effect of Bezel Presence and Width on Visual Search’. In: *Proceedings of The International Symposium on Pervasive Displays (PerDis’14).* ISBN: 9781450329521. DOI: `10.1145/2611009.2611019`.

Wallace, Jim, Vicki Ha, Ryder Ziola and Kori Inkpen (2006). ‘Swordfish’. In: *CHI ’06 Extended Abstracts on Human Factors in Computing Systems.* ACM. DOI: `10.1145/1125451.1125724`.

Want, Roy & Intel Research (2009). ‘Through Tinted Eyeglasses’. In: *IEEE Pervasive Computing* 8.3, pp. 2–4. DOI: `10.1109/mprv.2009.58`.

Warr, Andrew, Ed H. Chi, Helen Harris, Alexander Kuscher, Jenn Chen, Robert Flack and Nicholas Jitkoff (2016). ‘Window Shopping: A Study of Desktop Window Switching’. In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems.* ACM. DOI: `10.1145/2858036.2858526`.

Wei, Chunxue, Difeng Yu and Tilman Dingler (2020). ‘Reading on 3D Surfaces in Virtual Environments’. In: *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR).* IEEE. DOI: `10.1109/vr46266.2020.00095`.

White, Sheida, Jing Chen and Barbara Forsyth (2010). ‘Reading-Related Literacy Activities of American Adults: Time Spent, Task Types, and Cognitive Skills Used’. en. In: *Journal of Literacy Research* 42.3, pp. 276–307. ISSN: 1554-8430. DOI: `10.1080/1086296X.2010.503552`.

Wimmer, Raphael and Fabian Hennecke (2010). ‘Everything is a window: Utilizing the window manager for multi-touch interaction’. In: *Proc. EICS* 10.

Wirtz, Hanno, Jan Rüth, Martin Serror, Jó Ágila Bitsch Link and Klaus Wehrle (2014). ‘Opportunistic interaction in the challenged internet of things’. In: *Proceedings of the 9th ACM MobiCom workshop on Challenged networks - CHANTS ’14*. ACM Press. DOI: `10.1145/2645672.2645679`.

Wittmann, Frieder, Olivier Lambercy and Roger Gassert (2019). ‘Magnetometer-Based Drift Correction During Rest in IMU Arm Motion Tracking’. In: *Sensors* 19.6, p. 1312. DOI: `10.3390/s19061312`.

XMonad Developers (2021). *xmonad - the tiling window manager that rocks.* URL: `https://web.archive.org/web/20210813180852/https://xmonad.org/`.

Yang, Shun-nan, Yu-chi Tai, Hannu Laukkanen and James E. Sheedy (2011). ‘Effects of ocular transverse chromatic aberration on peripheral word identification’. In: *Vision Research* 51.21-22, pp. 2273–2281. DOI: `10.1016/j.visres.2011.08.022`.

Yang, Ungyeon and Nam-Gyu Kim (2021). ‘Near-body Interaction Enhancement with Distance Perception Matching in Immersive Virtual Environment’. In: *Journal of Multimedia Information System* 8.2, pp. 111–120. DOI: `10.33851/jmis.2021.8.2.111`.

Zeidler, Clemens, Christof Lutteroth and Gerald Weber (2013). ‘An Evaluation of Stacking and Tiling Features within the Traditional Desktop Metaphor’. In: *Human-Computer Interaction – INTERACT 2013*. Springer Berlin Heidelberg, pp. 702–719. DOI: `10.1007/978-3-642-40483-2_49`.

Zhao, He and Zheyao Wang (2012). ‘Motion Measurement Using Inertial Sensors, Ultrasonic Sensors, and Magnetometers With Extended Kalman Filter for Data Fusion’. In: *IEEE Sensors Journal* 12.5, pp. 943–953. DOI: `10.1109/jsen.2011.2166066`.

Zujovic, Lazar, Vladislav Kecojevic and Dragan Bogunovic (2021). ‘Interactive mobile equipment safety task-training in surface mining’. In: *International Journal of Mining Science and Technology* 31.4, pp. 743–751. DOI: `10.1016/j.ijmst.2021.05.011`.

Glossary

3 degrees of freedom (3DoF)

See: Degree of freedom

6 degrees of freedom (6DoF)

See: Degree of freedom

Accurate words per minute (AWPM)

A measure of reading performance combining accuracy of reading and speed.

Augmented reality (AR)

A virtual environment in which the primarily perceived environment is a view of the real world, with supplementary overlaid data being virtual in nature.

CAVE automatic virtual environment (CAVE)

A virtual reality system utilising projected images onto a combination of walls, floor, and ceiling of a purposebuilt room. Users of these systems have their position and orientation within the room tracked, with displayed content updated to reflect movement.

Cross reality (CR)

A virtual environment in which the primarily perceived environment is virtual, with supplementary overlaid data composed of or indicative of realworld details.

Degree of Freedom

The number of ways an object or body may move. In the context of virtual reality (VR), the term 3 degrees of freedom (3DoF) is used to describe a system that permits rotational movements about the x, y, and z axes. The term 6 degrees of freedom (6DoF) is similarly used to describe a system with the same rotations properties as seen with a 3DoF, with the addition of permitting movement along each of the x, y, and z axes.

Head-mounted displays (HMDs) devices that support 3DoF are capable of tracking where the device is pointing, but not where it is positioned in space. HMDs devices that support 6DoF are capable of tracking where the device is pointing as well as where it is positioned in space — though position tracking may only function within defined areas.

See: 3DoF & 6DoF

Field of view (FOV)

A descriptor of the restriction imposed on a perceiver's visual field imposed by a device. A horizontal FOV $< 210°$ would result in a loss of peripheral vision, and $< 114°$ would impact on the binocular vision field.

Flesch reading ease (FRE)

A readability score in range [0..100]derived from the number of syllables, words, and sentences in a text sample. Higher scores represent easily understandable text, lower scores more complex text.

Head-mounted display (HMD)

Display or displays positioned ahead of the user's eye or eyes. Displays are secured to the user's head, and may include additional sensors for tracking head orientation or position.

Head-up display (HUD)

Informational content displayed in a fixed-egocentric manner, appearing within the field of view of a user.

Immersive virtual reality (IVR)

Complete replacement of the real environment with virtual content, with the content providing a new reality for the perceiver. This is accomplished through the use of vision obscuring head-mounted goggles containing opaque displays, with the wearer unable to see the realworld environment. Users of such devices would see only the virtual environment.

In-engine unit (IEU)

A unit of measurement with no direct mapping to real-world dimensions used within the developed virtual environment for dimensioning and positioning virtual elements.

Information-rich virtual environment (IRVE)

A virtual environment containing graphical and spatial data augmented with abstract or symbolic information related to the displayed content.

Inter-pupillary distance (IPD)

Distance between pupil centres.

Meridian

Imaginary great circle of a sphere that pass through the poles of the sphere. The length of meridian lines is equal to the circumference of the sphere, and each meridian is perpendicular to every parallel line. In the context of developed virtual environment, the poles of a display sphere are found along the y axis, with the prime meridian intersecting the $+z$ axis. Meridians are identified by azimuth angles given with respect to the prime meridian, which is assigned the azimuth angle of 0.

Mixed reality (MR)

A virtual environment in which the primarily perceived environment includes virtual and real components in varying amounts, with supplementary overlaid data being virtual in nature or composed of or indicative of realworld details. Can be described as a combination of Cross reality (CR), Virtual reality (VR), Augmented reality (AR), and embodied virtuality (EV) (Want, 2009).

Oculus Rift CV1 (CV1)

An immersive virtual reality (IVR) head-mounted display (HMD) developed by Oculus Inc, released in 2016. Features 6 degrees of freedom (6DoF) rotational and positional tracking, accomplished using external “constellation” sensors that track infrared LEDs positioned on the headset. Video is presented with a barrel distortion on two OLED pentile displays providing a resolution of 1080×1200 per eye, with the distortion corrected by lenses between the display and user’s eyes. Lens position can be adjusted horizontally to accommodate different inter-pupillary distance (IPD) settings.

Oculus Rift DK1 (DK1)

An immersive virtual reality (IVR) head-mounted display (HMD) developed by Oculus Inc, released in 2013. Features 3 degrees of freedom (3DoF) rotational tracking using on-board gyroscopes, accelerometers, and magnetometers. Video is presented with a barrel distortion on an RGB display providing a resolution of 640×800 per eye, with the distortion corrected by lenses fitted between the display and user's eyes. Lenses are user-replaceable and the lens–display distance could be adjusted to allow for simple dioptric correction.

Oculus Rift DK2 (DK2)

An immersive virtual reality (IVR) head-mounted display (HMD) developed by Oculus Inc, released in 2014. Features 6 degrees of freedom (6DoF) rotational and positional tracking using a combination of on-board gyroscopes and accelerometers, combined with an external infrared (IR) camera tracking infrared LEDs on the headset.

Video is presented with a barrel distortion on an OLED pentile display providing a resolution of 960×1080 per eye, with the distortion corrected by lenses fitted between the display and user's eyes. Lenses are user-replaceable and the lens–display distance can be adjusted to allow for simple dioptric correction.

Parallel

Imaginary circles across a sphere, perpendicular with all meridian lines, and parallel to all other parallels. The length of parallels varies, with only the equator parallel having a length equal to the circumference of the sphere. The equator parallel is positioned equidistant between the poles of the sphere. Parallels are identified by inclination angles given with respect to the equator parallel, which is given the inclination angle of 0.

Pixel per centimetre (Ppcm)

A measure of pixel density for computer displays, specifying the number of pixels that can fit inside one centimetre of space. A metric alternative to the pixels per inch (ppi) measure.

Pixels per degree (PPD)

The number of pixels that can be seen within a degree of vision.

Virtual

Adjective describing an object or concept that does not exist. An object would be described as virtual if it were generated by a computer system.

Virtual reality (VR)

A virtual environment in which the primarily perceived environment is virtual, with supplementary overlaid data being virtual in nature.

See also: Immersive virtual reality (IVR)

Appendix A
Workplace Observation

This appendix contains details of the ethical consent application submitted to support the user study reported in Chapter 3, as well as material supplementary to the same study.

- Ethical Approval Letter from the Human Research Ethics Committee of the Faculty of Computing and Mathematical Sciences at the University of Waikato, dated 17 June 2013. (Appendix A.1)
- Participant Information Sheet, which each study participant received and had explained during participant recruitment. (Appendix A.2)
- Research Consent Form, which each participant signed prior to beginning the study. (Appendix A.3)
- Periodic Collection, information collected by the researcher during hourly data collection sessions. Arrangement of 'full-screen' indicates that primary method of arranging windows was through maximising windows. Arrangement of 'tiled' indicates that the primary method of arranging windows was through manually or automatically sizing and positioning windows to avoid overlap. Arrangement of 'split' indicates that a mixture of arrangement strategies was in use, with neither full-screen nor tiled being over represented. (Appendix A.4)
- Closing Interview, notes that the researcher used to guide exit interviews with participants.(Appendix A.5)

A.1 Ethical Approval Letter

Computing and Mathematical Sciences
Rorohiko me ngā Pūtaiao Pāngarau
The University of Waikato
Private Bag 3105
Hamilton
New Zealand

Phone +64 7 838 4021
www.scms.waikato.ac.nz

THE UNIVERSITY OF
WAIKATO
Te Whare Wānanga o Waikato

17 June 2013

Cameron Grout
C/- Department of Computer Science
THE UNIVERSITY OF WAIKATO

Dear Cameron

Request for approval to conduct a research project involving human participants

I have considered your request to carry out a study for your PhD research project *Desktop environments in virtual reality and the usability impacts of an adjusting workspace paradigm.* The goal is to create a primary source of data that breaks down details of how users in career fields utilize their existing desktop computers.

I note that data will be summarized and anonymized, audio recordings will be written to DVD's. No names will be collected from any of the participants; the questionnaires will have ID numbers. Once summarized the collected data will be stored in the FCMS data archive for five years

The procedure described in your request is acceptable.

The research participants' information sheet, consent form and questionnaire sheets all meet the requirements of the University's human research ethics policies and procedures.

I therefore approve your application to perform the research project.

Yours sincerely,

LA Hunt.

Lyn Hunt
Human Research Ethics Committee
Faculty of Computing and Mathematical Sciences

A.2 Participant Information Sheet

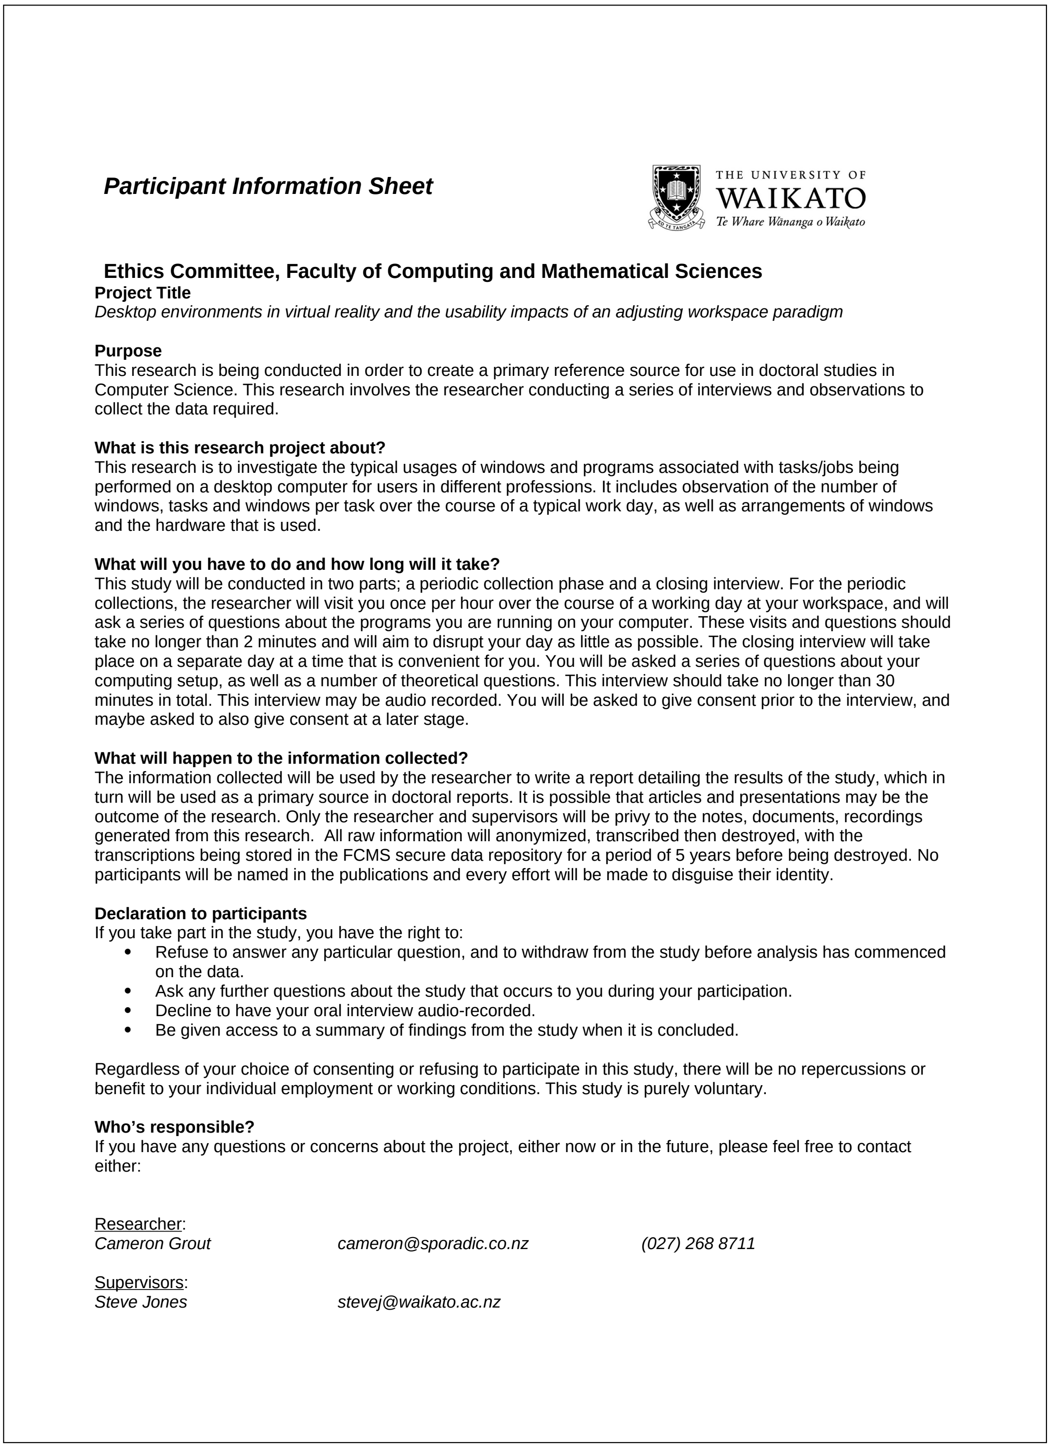

Participant Information Sheet

THE UNIVERSITY OF
WAIKATO
Te Whare Wānanga o Waikato

Ethics Committee, Faculty of Computing and Mathematical Sciences

Project Title
Desktop environments in virtual reality and the usability impacts of an adjusting workspace paradigm

Purpose
This research is being conducted in order to create a primary reference source for use in doctoral studies in Computer Science. This research involves the researcher conducting a series of interviews and observations to collect the data required.

What is this research project about?
This research is to investigate the typical usages of windows and programs associated with tasks/jobs being performed on a desktop computer for users in different professions. It includes observation of the number of windows, tasks and windows per task over the course of a typical work day, as well as arrangements of windows and the hardware that is used.

What will you have to do and how long will it take?
This study will be conducted in two parts; a periodic collection phase and a closing interview. For the periodic collections, the researcher will visit you once per hour over the course of a working day at your workspace, and will ask a series of questions about the programs you are running on your computer. These visits and questions should take no longer than 2 minutes and will aim to disrupt your day as little as possible. The closing interview will take place on a separate day at a time that is convenient for you. You will be asked a series of questions about your computing setup, as well as a number of theoretical questions. This interview should take no longer than 30 minutes in total. This interview may be audio recorded. You will be asked to give consent prior to the interview, and maybe asked to also give consent at a later stage.

What will happen to the information collected?
The information collected will be used by the researcher to write a report detailing the results of the study, which in turn will be used as a primary source in doctoral reports. It is possible that articles and presentations may be the outcome of the research. Only the researcher and supervisors will be privy to the notes, documents, recordings generated from this research. All raw information will anonymized, transcribed then destroyed, with the transcriptions being stored in the FCMS secure data repository for a period of 5 years before being destroyed. No participants will be named in the publications and every effort will be made to disguise their identity.

Declaration to participants
If you take part in the study, you have the right to:

- Refuse to answer any particular question, and to withdraw from the study before analysis has commenced on the data.
- Ask any further questions about the study that occurs to you during your participation.
- Decline to have your oral interview audio-recorded.
- Be given access to a summary of findings from the study when it is concluded.

Regardless of your choice of consenting or refusing to participate in this study, there will be no repercussions or benefit to your individual employment or working conditions. This study is purely voluntary.

Who's responsible?
If you have any questions or concerns about the project, either now or in the future, please feel free to contact either:

Researcher:

Cameron Grout	*cameron@sporadic.co.nz*	*(027) 268 8711*

Supervisors:

Steve Jones	*stevej@waikato.ac.nz*

Bill Rogers *mcoms0108@waikato.ac.nz*
Mark Apperley *mark.apperley@waikato.ac.nz*

A.3 Research Consent Form

Research Consent Form

THE UNIVERSITY OF
WAIKATO
Te Whare Wānanga o Waikato

Ethics Committee, Faculty of Computing and Mathematical Sciences

Desktop environments in virtual reality and the usability impacts of an adjusting workspace paradigm

Consent Form for Participants

I have read the **Participant Information Sheet** for this study and have had the details of the study explained to me. My questions about the study have been answered to my satisfaction, and I understand that I may ask further questions at any time.

I also understand that I am free to withdraw from the study before the closing interview, or to decline to answer any particular questions in the study. I understand I can withdraw any information I have provided up until the researcher has commenced analysis on my data. I agree to provide information to the researchers under the conditions of confidentiality set out on the **Participant Information Sheet**.

I agree to participate in this study under the conditions set out in the **Participant Information Sheet**.

Signed: ____

Name: ____

Date: ____

Additional Consent

I agree / **do not agree** (select one) to my oral responses being audio recorded.

Signed: ____

Name: ____

Date: ____

Researcher:

Cameron Grout	*cameron@sporadic.co.nz*	*(027) 268 8711*

Supervisors:

Steve Jones	*stevej@waikato.ac.nz*

Bill Rogers *mcoms0108@waikato.ac.nz*
Mark Apperley *mark.apperley@waikato.ac.nz*

A.4 Periodic Collection

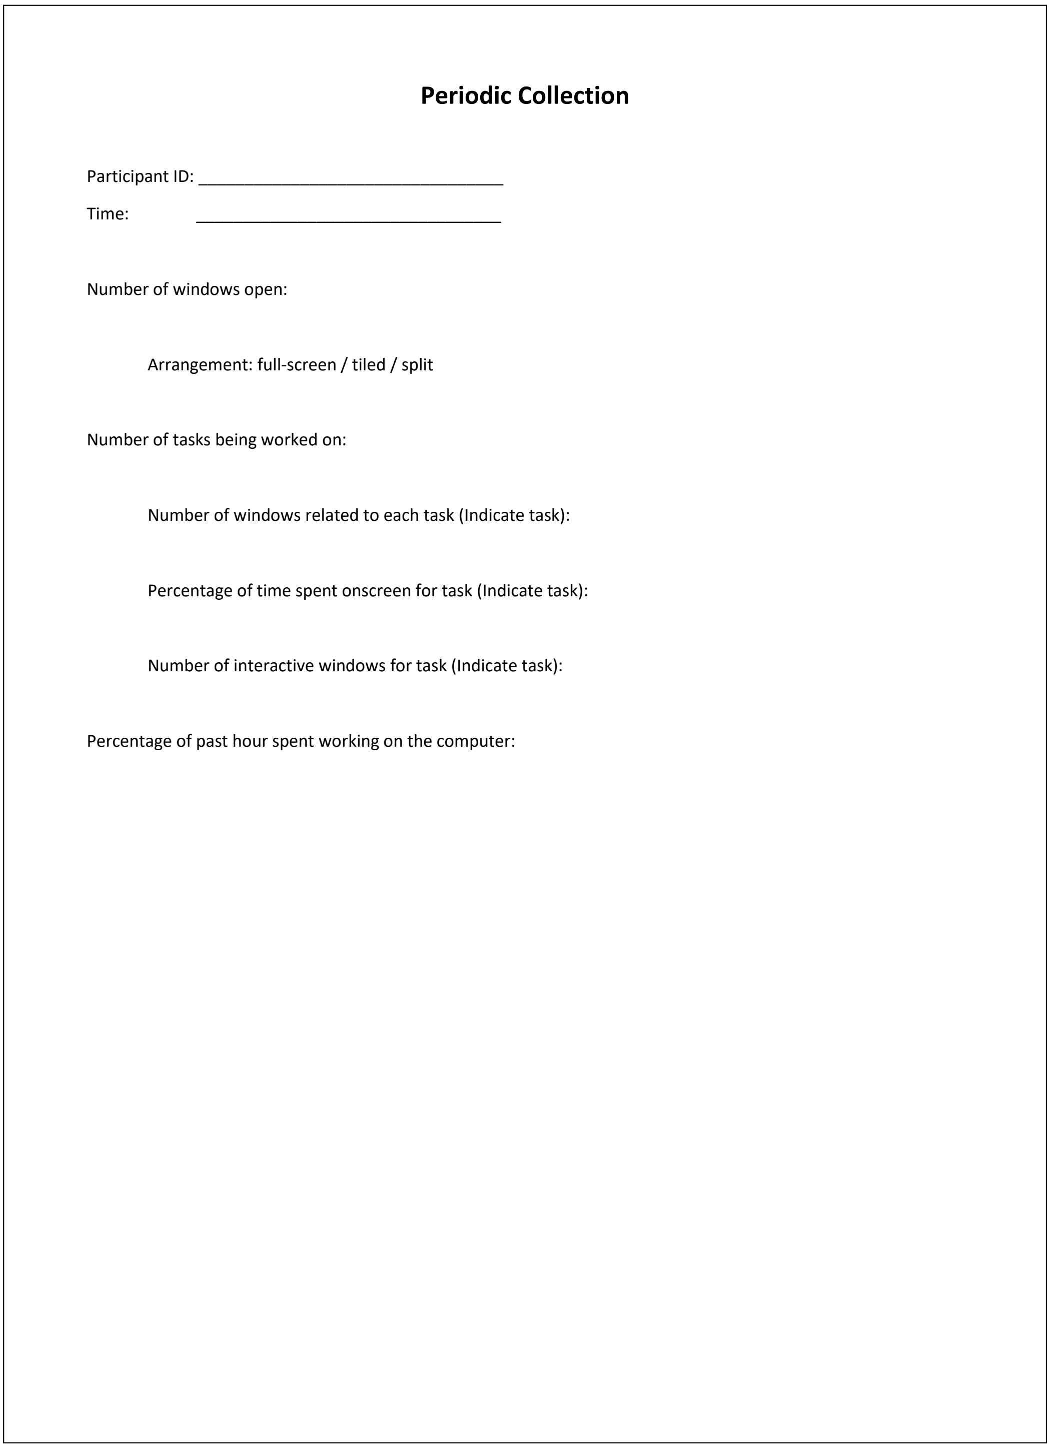

Periodic Collection

Participant ID: ________________________________

Time: ________________________________

Number of windows open:

Arrangement: full-screen / tiled / split

Number of tasks being worked on:

Number of windows related to each task (Indicate task):

Percentage of time spent onscreen for task (Indicate task):

Number of interactive windows for task (Indicate task):

Percentage of past hour spent working on the computer:

A.5 Closing Interview

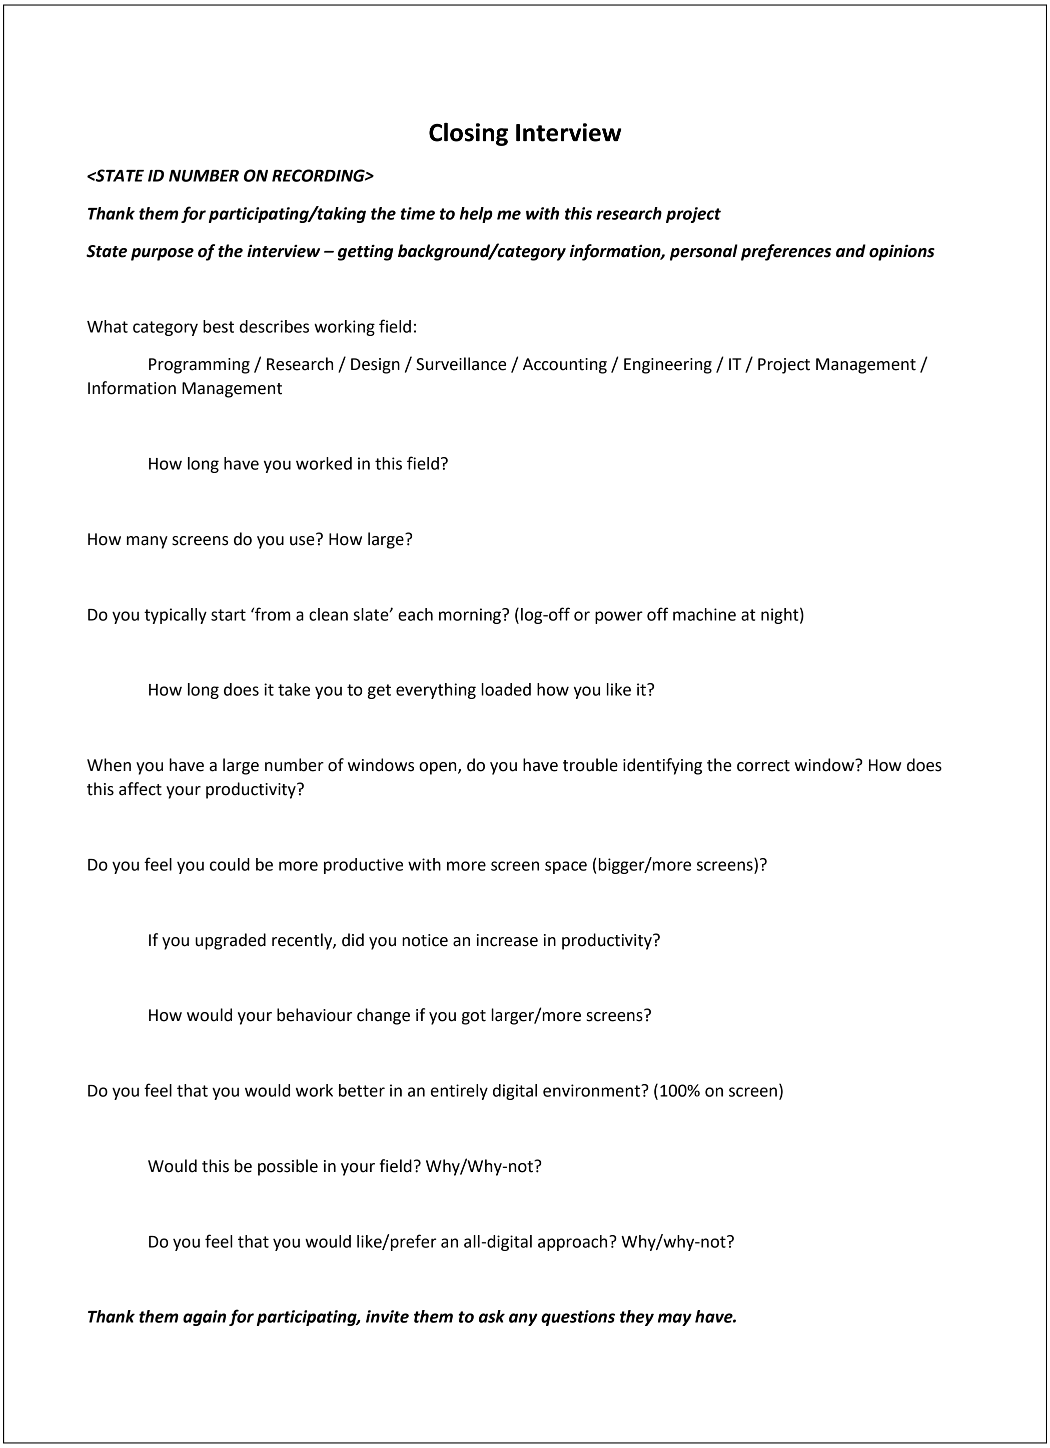

Closing Interview

<STATE ID NUMBER ON RECORDING>

Thank them for participating/taking the time to help me with this research project

State purpose of the interview – getting background/category information, personal preferences and opinions

What category best describes working field:

Programming / Research / Design / Surveillance / Accounting / Engineering / IT / Project Management / Information Management

How long have you worked in this field?

How many screens do you use? How large?

Do you typically start 'from a clean slate' each morning? (log-off or power off machine at night)

How long does it take you to get everything loaded how you like it?

When you have a large number of windows open, do you have trouble identifying the correct window? How does this affect your productivity?

Do you feel you could be more productive with more screen space (bigger/more screens)?

If you upgraded recently, did you notice an increase in productivity?

How would your behaviour change if you got larger/more screens?

Do you feel that you would work better in an entirely digital environment? (100% on screen)

Would this be possible in your field? Why/Why-not?

Do you feel that you would like/prefer an all-digital approach? Why/why-not?

Thank them again for participating, invite them to ask any questions they may have.

Appendix B
Publications

This appendix contains a verbatim copy of a paper presented at CHINZ in 2015. This paper covers results of the user study presented in Chapter 6.

Grout, Cameron, Mark Apperley, William Rogers and Steve Jones (2015). 'Reading text in an immersive head-mounted display: An investigation into displaying desktop interfaces in a 3D virtual environment'. In: *CHINZ 2015: Proceedings of the 15th New Zealand Conference on Human–Computer Interaction*, pp. 9–16. ISBN: 9781450336703. DOI: `10.1145/2808047.2808055`.

Reading text in an immersive head-mounted display

An investigation into displaying desktop interfaces in a 3D virtual environment

Cameron Grout
University of Waikato
Hamilton, New Zealand
cgrout@waikato.ac.nz

William Rogers
University of Waikato
Hamilton, New Zealand
b.rogers@waikato.ac.nz

Mark Apperley
University of Waikato
Hamilton, New Zealand
mapperle@waikato.ac.nz

Steve Jones
University of Waikato
Hamilton, New Zealand
stevej@waikato.ac.nz

ABSTRACT

This paper describes a user study analyzing the performance of reading tasks typical of a general purpose computing environment conducted in immersive virtual reality headsets. Results of this study are evaluated, and suggest that reading tasks can be performed with near equivalent performance in the virtual environment when compared to performance values obtained from baseline tasks on a traditional display.

CCS Concepts

•**Computing methodologies** → **Virtual reality;** *Perception;* •**Human-centered computing** → *User studies;* User interface design; •**Applied computing** → Personal computers and PC applications;

Keywords

Virtual reality; Fonts hinting; Rendering techniques; Readability; Application redirection

1. INTRODUCTION

In 2003, Czerwinski *et al.*[6] identified a significant performance advantage in using multiple monitors or large higher resolution displays. Through analyzing 15 participants performing tasks on large displays, they noted that participants experienced a 9% decrease in required time to complete a task when using a large display compared to the time required on a smaller display. Czerwinski noted that 14 of the 15 participants preferred the large display over smaller ones, reporting that it was easier to manage multiple windows and to switch between tasks on the larger display.

In 2013, the primary researcher conducted an on-site observation study with a medium sized company, where 15 participants were observed at hourly intervals to identify the number of applications being managed, and the organization of these windows on their screens. On average, 7 open windows were present on participants computers, with as many as 15 windows open in extreme cases. The dominant presentation method for these windows was a full-screen configuration, with free-form tiling and 'half-maximized' window arrangements being less favored methods of window presentation. While the majority of participants used multiple displays, 53% of participants indicated that despite their large display areas, they had difficulty identifying windows pertinent to their tasks, and indicated that window management and identification resources provided by their window managers were insufficient for facilitating quick task switching. When polled, 67% of participants indicated that following their last display upgrade (encompassing larger, or more displays) they experienced an increase in productivity, and believed that should their displays be upgraded again, they would experience a similar productivity boost.

An issue with traditional windowing systems and displays is that the amount of usable desktop space is finite, limited by the size and number of displays used. A number of factors limit how large this space can grow, including the price of larger displays, and hardware limitations preventing large numbers of displays from being connected to a computer. In order to help users utilize their limited space efficiently, many window managers provide facilities for organizing windows, through tiling methods, or options like 'half-maximizing' windows to occupy half of the display, and provide a number of assistive measures for identifying windows. In addition to these window management techniques, some systems provide virtual desktop options, allowing users to have multiple virtual desktops. Each of these virtual desktops display a set of windows, and can be swapped onto the physical device for display using key combinations. This allows users to partition work tasks onto different virtual spaces, organizing tasks in an efficient manner, but introduces additional cognitive load in recalling which virtual desktop contains the desired information, and requires additional action to switch desktops.

As it is clear that users could benefit from more usable screen space, alternative solutions have been proposed that utilize head-mounted displays (HMD) in order to expand the users' working space by presenting a contiguous virtual world. In 1983, Fisher *et al.*[7] presented an immersive virtual environment (IVE) that could present virtual windows

CHINZ 2015, September 03-04, 2015, Hamilton, New Zealand

DOI: http://dx/doi.org/10.1145/2808047.2808055

around the user in a 3D environment through use of binocular virtual reality HMDs. The interface was proposed as a general purpose system, allowing visualization of scenes and telepresense in addition to displaying information on virtual windows. Interaction with this system was accomplished through use of an instrumented glove, as well as speech and gesture recognition allowing interaction with virtual controls.

In 1993, Reichlen[10] demonstrated Sparcchair, a system using a custom built monochromatic monocular HMD that presents familiar applications on a large X-windows frame buffer, with the position of a viewport into this buffer determined by orientation of the users' head. Reichlen reported that users were attracted by the facility to rapidly switch task contexts in this system, particularly for small quick tasks, and it was noted that there was a distinct advantage in having an "unlimited number of windows open at once"[10], reducing mouse movements and clicks normally associated with task switching on a traditional display.

Lawton[8] identified in their 2006 article that the primary barrier to common adoption of virtual reality technologies was the lack of affordable VR devices with specifications that would make them desirable. In recent years, announcements for consumer priced virtual reality hardware with competitive specifications by Oculus, Sony, Valve and others has driven consumer interest in such devices. While the initial target market for these devices is identified as the gaming and entertainment industries, companies such as Bloomberg have begun investigating the use of these devices in areas such as finance[2]. With the release of these affordable devices, it is expected that it will begin to be commonplace to see them in homes and businesses. It is also expected that consumers will identify that through the use of such devices, digital workspaces can be expanded, which will drive demand for general purpose interfaces and environments that provide access to traditional applications using these devices.

As a first step towards building a system, we needed to identify common tasks performed by users. In the preliminary study discussed above, we observed that users were employing large numbers of windows to make a large amount of reference material available for their work. In most cases this material was textual in a flow or a tabular format.

More generally, in a diary study conducted in 2010 by White *et al.*[13] roughly 400 Americans aged 20 years and older reported "all reading, writing and computer activities"[13, p. 285] for one working and one non-working day within a week. Analysis of this data by White revealed that on average participants spent 312 minutes on working days, and 244 minutes on non-working days engaging in reading tasks taking longer than 1 minute.

Clearly in order to support general computing tasks, an interface must facilitate reading text with ease and comfort. To address this, we undertook a study of reading in an immersive virtual environment. Our goal was to determine the best way of displaying textual information.

This paper describes a user study analyzing the performance of reading tasks inside a general purpose computing environment for use within virtual reality headsets. Section 2 describes the proposed system; Section 3 & 4 describes our user study and presents the results; Sections 5 & 6 present our conclusions.

2. SPATIAL INFORMATION DISPLAY

In our work, we are building a body-stabilized virtual environment (view is "fixed relative to the users body position and varies as the user changes viewpoint orientation"[4]). The system is designed for use with the Oculus Rift series of immersive virtual reality displays, that is intended to be used in a seated configuration. This environment is written in C#, using Win32 APIs via Platform Invocation services, and DirectX through use of the SharpDX library. This environment gives users the impression of being surrounded by a virtual sphere, upon which application windows can be placed and freely moved (See Figure 1). Application interfaces are captured and interacted with through use of application redirection techniques discussed by van Dantzich *et al.*[12], though the implemented system is significantly cruder than the reference system due to the lack of access to required source code.

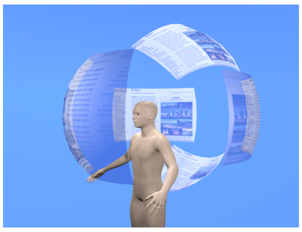

Figure 1: Visualization of the developed environment, showing curved panels

Navigation inside this environment is accomplished through head movement, detected by sensors provided with the Oculus devices. Orientation is detected via three degrees of freedom (3doF) sensors onboard each headset, which provide data collected from gyroscopic, magnetometer and accelerometer sensors at 1000Hz. Orientation data obtained from these sensors is used to modify the viewport that is seen by the user, by rotating the virtual cameras to match the real-world head rotation. The DK2 additionally provides positional data through use of an external infrared camera, which detects the location of the headset within a $72°W \times 52°H$ frustum at a range of 0.5m–2.5m from the camera. When using the DK2 headset, this positional data is used to modify the position of the virtual camera in concert with real world movements, allowing the user to move their eyes closer to an item in order to see it better amongst other uses.

Interaction in the environment and applications within is performed using a traditional mouse and keyboard interface rather than utilizing any novel system. The virtual mouse cursor takes XY movements from the physical mouse, and translates these to Euler rotations which are then applied to the virtual cursor glyph. Familiar interactions have been preserved, such as click-dragging on an application window decoration in order to move it within the environment, and resizing applications through grabbing the border or handle located in the bottom-right of the window.

In order to provide a sense of familiarity, applications are represented such that they are always oriented in space to face the world origin, with the center-top and center-bottom of the application aligning on a longitudinal line. Applications can be positioned such that they occlude one-another, obeying z-ordering rules that are familiar from traditional desktop environments. Applications displayed inside the environment can currently be displayed mapped onto a plane, or onto a distorted panel that maintains a fixed distance from the environment origin at all points on its surface.

3. USER STUDY

During informal user testing with the described environment, a number of questions were raised that needed to be addressed before continuing on with development. A number of these questions were addressed using reading performance metrics, comparing accurate reading speeds across devices in order to provide a reasonable indication of performance for the study. A reduction in reading performance would indicate a likely reduction in productivity when compared to the participant's normal performance, and vice versa.

During informal testing there had been no consensus as to which headset was preferable, despite their numerous differences (detailed in Table 1). Because the DK2 has a higher resolution, initial expectations were that, in spite of its reduced field of view (FOV) it would be preferred over the DK1, but general feedback given was that they were both equally suited for interacting with the environment. In order to determine if a performance difference was present, it was deemed necessary to perform experiments on both devices, and to compare the results.

Table 1: Oculus Rift Version Features

	DK1	**DK2**
Resolution (pixels/eye)	640x800	960x1080
FOV (Nominal)	110°	100°
Tracking	3doF	6doF
Pixel Layout	RGB	RGBG PenTile
Weight	380g	440g

During these informal tests, a number of users indicated that they felt that curved applications were harder to read because they felt unnatural. These users expressed a preference for flat panels, indicating that these felt more like reading information from a screen – something they were very familiar with – and it was believed that these flat panels allowed for faster reading due to this feeling of familiarity. This sentiment was not shared by all users however, so this indicated that an investigation into these distortion options, and their effect on performance, was desirable.

For informal testing, the distance between all applications and the environment origin had been fixed to a value determined to be suitable by the primary researcher. As this position was chosen arbitrarily, it is possible that users were required to physically move to get themselves in a position in the virtual environment where they could read comfortably. In order to provide a default position, preferred distances should be collected, allowing a well reasoned average position to be determined.

3.1 Approach

In order to address these questions and concerns, an experiment was designed utilizing a modified version of the developed environment. In place of the normal live applications, static images of an application were displayed, and the user given the ability to change the distance at which the image was placed relative to themselves. Users would then read textual content from the application image, and the application position, reading speed and reading accuracy would be recorded.

In order to ensure the size of the application window was not a confounding factor, multiple sizes of page were used, with one of each size being displayed as a flat panel, and one as a curved panel. Each participant was tested with both headsets and was shown the same number of samples in each. In order to obtain a basis for comparison, users were also presented with images of the same size displayed on a traditional computer screen, and the same statistics were recorded.

3.2 Hypotheses

The study set out to test several hypotheses, aiming to address each of the questions and concerns raised. The hypotheses were as follows:

1. Curved display panels will be preferred over flat panels for all display sizes
2. An increase in reading performance will be noted in all experiments using a curved panel when compared to that of a equivalently sized flat panel
3. Study participants will prefer the traditional computer monitor over either of the headsets, and will prefer the DK2 over the DK1
4. Reading performance will peak using the traditional monitor, and will be followed by the DK2, then finally DK1
5. Participants will identify an optimal reading distance in the virtual environment, and we will observe a near constant placement at this point for all participants

3.3 Participants

Participants for the user study were volunteers sourced from the university student population, with the majority being from the Faculty of Computing and Mathematical Sciences. 24 participants were selected from the pool of volunteers with visual acuity the deciding factor in selection. Of the participants, 16 were male and 8 were female; 2 of these participants had their sight corrected through the use of contact lenses, and 3 had congenital issues affecting a single eye that could not be corrected. This is visualized in Figure 2.

Subjects self-reported their preferences for reading content, with 23 participants indicating a strong preference for reading on-screen, and 1 participant indicating a strong preference for reading physical paper-based media. Of the participants, only 1 had extensive experience with virtual reality devices, with the remainder indicating very slight, or no experience with VR.

3.4 Equipment

Two immersive virtual reality headsets, the Oculus Rift DK1 (DK1) and the Oculus Rift DK2 (DK2), were used as the headset devices for this experiment. Both headsets supply multiple lens sets which can be swapped as required by the user. Both headsets were fitted with the ‘A’ lens

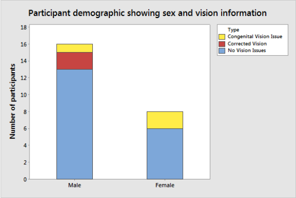

Figure 2: Study participant demographic information showing numbers of participants with vision issues

sets, intended for use when no vision issues are present, and the eye relief dials were set to their maximum extent. The DK1, which includes configurable hardware brightness and contrast setting, was configured such that these levels were set the middle of their operational levels, which gave the display a similar appearance to a traditional display. As the DK1 does not provide positional data, the positional tracking feature of the DK2 was disabled.

A Dell U2312HMi LCD monitor was used as the traditional display. This display contains a 23" panel, with a 60Hz refresh rate and native resolution of 1920x1080. The display was reset to factory settings prior to commencement of the study, with brightness and contrast set to 75% of their maximum values.

The testing computer was a Dell Optiplex 9020 with an quad-core Intel i7 4770 clocked at 3.40 GHz, with 16GB of DDR3 RAM and an NVidia GeForce GTX 750 Ti. The operating system used was Windows 7 SP1 64-bit, updated with the latest hotfixes and updates available at the time of testing. Peripherals used were a traditional 2-button mouse with scroll-wheel, and a 104-key US-layout keyboard.

The software running the experiment was a version of the application redirection system developed by the primary researcher that was modified to display static images. Interaction was managed through the keyboard arrow keys, to advance or retreat through samples, and the mouse scroll wheel, to move the panel toward or away from the participant. When the sample was changed, the sample's position in the environment was set to an unreadable distance in order to force the participants to position each panel independently. This application displayed output on the headsets used, as well as on a secondary monitor observed by the experiment supervisor, and operated in the 60-75 FPS range.

3.5 Experiment and Method

Fifteen page samples were selected from the Simple Wikipedia collection of major world cities. From each sample, a paragraph was selected that consisted of 110-130 words, with a Flesch–Kincaid[1] reading score of between 60 & 65 (as measured by Microsoft Word, 2013). Each of the selected sample paragraphs was modified to remove superfluous visual noise, such as inline reference marks, and to remove words which would be difficult for someone unfamiliar with them to pronounce. 5 samples were then displayed at each of the resolutions specified in Table 2 using Google Chrome running on Windows 7 with factory default settings, then a screenshot of the sample was taken. The collection of sample images were then split into 3 sequences as specified in Table 3.

Table 2: Sample sizes and aliases

Resolution	Aspect Ratio	Name
800x600	4:3	**S**mall
1366x768	16:10	**M**edium
1920x1080	16:9	**L**arge

Table 3: Sample collections

Sequence #	S	M	L
Seq 0	2	2	2
Seq 1	1	1	1
Seq 2	2	2	2

Each participant was tested on two headsets, and one traditional computer screen. The order of these tests was assigned in an alternating fashion to participants as DK1, Monitor, DK2 or DK2, Monitor, DK1. The baseline test using the monitor was always the second device, as it gave the participant's eyes a chance to rest in between headset runs. A sequence of sample images was assigned to each headset device in an alternating fashion, with the monitor device always being assigned Seq 1. Within each sequence, the order of display was randomized between participants and devices. For each sequence displayed on a headset, one of each sample size was displayed using a curved distortion, and the other with no distortion (flat).

For experimental runs using the traditional computer screen, participants were seated facing the display, and positioned such that their eyes were approximately 750mm away from the display. The height and angle of the display were then adjusted to ensure that the participant's line of sight was perpendicular with the center of the display. For experimental runs using a headset, the participant was fitted with the appropriate device such that it was fixed firmly but comfortably to their head. They were then seated in front of a desk, and provided with a keyboard and mouse. A sample image was then displayed centered on the screen with black borders filling in any additional space.

For experiment runs using headsets, participants were presented with a sample image in the environment, loaded at an unreadable distance. The participant was then prompted to move the panel using the mouse scroll wheel until the sample was at the far edge of a comfortable reading distance. After positioning the panel, the panel position was logged and the participant was prompted to release the controls.

Following the setup for each experimental run, participants were shown each sample in sequence. For each displayed sample, the participant was prompted to begin reading the sample at a specified point in the text until asked to stop. As the participant read, the range and pattern of their head movements were recorded by the experiment supervisor, and their responses were recorded using audio recording devices. This procedure was repeated until all applicable samples had been read, then a questionnaire was

---

[1] A readability score derived from the Flesch readability formula, defined as
$RE = 206.835 - (0.846 \times NSYLL) - (1.015 \times \frac{W}{S})$
"Where NYSLL is the average number of syllables per 100 words, and $\frac{W}{S}$ is the average number of words per sentence"[3].

provided for the participant to fill in.

These questionnaires asked the participants if they encountered any difficulties reading the sample images, and if so, were prompted to indicate their perceived cause of the difficulty. For experiment runs using the headsets, participants were also asked if they had a preferred display method (curved, or flat), to indicate if they felt comfortable using the devices, and to rate any discomfort they had felt. Following the completion of all 3 experiment runs, the participants were asked to fill in one additional questionnaire, which asked them to rank the experimental devices in order of preference, and give reasoning for their decision.

3.6 Data Treatment & Analysis

Following the experiment, recorded participant responses were edited to remove comments to the experimental supervisor, as well as general comments. These included, but were not limited to, comments like 'this is an odd sensation' regarding reading in the headset, "I didn't know that" when reading novel information. Comments that were not deliberate asides and were a part of the thinking process such as "I think" after reading a word, and filler words and noises like "Umm" and "Ah", were not removed from the recordings. Any removals from the recordings were performed against the edges of the waveform that comprised the comment for removal. Any surrounding silence or words were left in place. The length of time required to read each sample was recorded from the point in the recording where the participant began to speak, until the end of the last word of the sample.

These modified recordings were then analyzed using running record techniques as described by Clay[5] to determine the accuracy of the reading. With this analysis, errors, self corrections and repetitions were recorded based on the rules described for the method. An error was recorded for each missed, substituted or inserted word. Neither contractions or initialisms were penalized, allowing "it's" and "it is" substitutions, as well as "US" and "United States" expansions or contractions. Proper noun errors were excluded from the error count, so long as the error could be attributed to unfamiliarity with the word (eg, "Chesapeake" said as "Cheese Peak"). If a participant made an error, then corrected themselves, the error was not counted, but was instead recorded as a self correction.

Accuracy of reading was calculated for each participant and sample using the word count of the sample, and the number of errors made by the participant while reading that sample using the equation:

$$AR_{sample} = 100 - \frac{errors}{wordcount} \times 100$$

In addition to calculating the accuracy of reading, an *Accurate Words per Minute* (AWPM) value was calculated with the wordcount of the sample, the number of errors made while reading the sample, and the time required to read the sample (TTR). This was evaluated as:

$$AWPM_{sample} = \frac{wordcount - errors}{TTR_{seconds}} \times 60$$

Following the calculations for AWPM for all participants and samples, the AWPM of each sample read while using a headset was compared against that participants own baseline AWPM calculated from their sample of the same size read from the traditional monitor. This was accomplished with the equation:

$$AWPM_{diff} = \frac{AWPM_{baseline} - AWPM_{sample}}{AWPM_{baseline}}$$

Information written by the testing application to log files was processed to determine the horizontal angle subtended for a lowercase letter 'm' for each sample at the location chosen by the participant. The letter 'm' was chosen for evaluation, as it was indicated as one of the more difficult characters to identify during headset experiments. Including color fringing caused by font-hinting, the lower case 'm' in the sample images occupied 12 × 8 pixels before being displayed in the headsets. The subtended angle was calculated by determining the angle between two 3-D vectors positioned at the center-left and center-right of the character bounding box. For flat panels, both best-case and worst-case angles were calculated, while for curved panels only a single value was needed as the nature of the curve distortion provides a constant angular character pitch across its entire area. Best case samples, as well as the one sample for curved samples, were taken with the assumption that the character to be measured was centered at the $(0, 0, +Z)$ position vector, where Z is the in-engine distance the panel was positioned at by the participant. Worst-case calculations for the flat samples were taken with the character to be measured positioned at the top-left extreme of the sample area. The character positions as measured on a flat panel can be visualized in Figure 3.

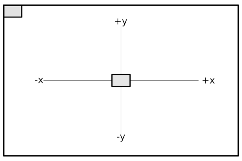

Figure 3: Locations of character samples used for calculating horizontal angle subtended relative to the sample in question. The box in the top-left shows the worst-case measurement, while the box in the center shows the best-case

4. RESULTS

From the final questionnaire completed at the end of the user study, the device preferences indicated by participants were collected, and presented in Table 4. From this data, a clear pattern can be seen, with the traditional monitor being indicated as the most preferred device by 92% of participants. 75% of participants indicated that the DK2 was their second choice of device, and finally 79% indicated that the DK1 was their least preferred. From the feedback given in this questionnaire it was found that most participants preferred the traditional monitor because of its familiarity and clarity, indicating that the text was crisp and easy to read, and that the amount of practice they had at reading from these types of displays meant there was nothing they had to grow accustomed to in order to complete the reading task. A minority of participants indicated a preference for the DK1 over the DK2, indicating that while reading was more difficult using the DK1 as the amount of head movements required was large, the clarity of text inside the DK1 was better than that of the DK2. With reference to hypothesis #3, we can conclude that the hypothesis is true, in that the greater proportion of the participant set indicated the device preference order hypothesized.

Table 4: Overall device preferences

Device	1^{st}	2^{nd}	3^{rd}
Monitor	22	2	0
DK2	1	18	5
DK1	1	4	19

Table 5: Panel distortion preferences

Distortion	DK1	DK2
Flat	2	2
Curved	16	19
No Pref.	5	3

From evaluating the responses given in questionnaires following experiment runs on each device, the preferred distortions indicated by participants for each headset were collected, and are shown in Table 5. As can be seen from the table, the curved distortion method was most preferred amongst the participants, with 67% indicating it as their preferred method on the DK1 headset, and 79% indicating the same on the DK2. 2 participants indicated that they preferred flat panels on both headsets, indicating that while the curved panels were easier to read when the samples were large, the curve felt unnatural and was likened to 'reading off the side of a building'. All participants indicated that when reading larger samples the curved distortion made reading easier, but noted that with smaller samples there was no notable difference in reading ease between curved and flat panels. With respect to hypothesis #1, we can conclude that for large samples, the curved distortion was preferred by the majority of participants on both headsets, but there was ambivalence about the distortion method with smaller samples.

Calculated differential AWPM values were evaluated using a paired t-test over the results for flat and curved samples of the same size. This information is displayed in Table 6, and is shown graphically in Figure 4. From examining the mean difference values between samples, it can be seen that for the small and medium samples on the DK2 and the small samples on the DK1 that the difference in performance for these samples were not statistically significantly different ($p \gg 0.05$). However for the large samples read on both headsets as well as the medium samples read on the DK1, large mean difference values are present, indicating that there was significant performance differences ($p = 0.000$) between the samples when they were presented in a flat or curved manner. This allows us to address hypothesis #2, and state that while the hypothesis holds true for large samples, it does not hold for smaller ones.

Referring to Figure 4, it can be seen that when comparing samples of the same size and distortion method across devices, the DK2 consistently has a smaller percentage difference from the baseline when compared to the performance on the DK1. With reference to the data contained in Table 6, it can be seen that small samples on both devices and either distortion method fall within standard error of the baseline value, indicating that for small samples, both devices are near equivalent in performance to the baseline. This is not the case however with medium samples, where the flat panel read on the DK1 plummets below the baseline performance. The remaining medium samples fall within standard error of the baseline performance, with the samples read on the DK2 having means lying closer to the baseline and with smaller standard error values than that of the DK1 curved sample. Large samples present a large spread of mean percentage differences, with the large curved sample read on the DK2 still falling near the baseline within standard error, and the remaining samples falling below the baseline completely. With respect to hypotheses #4, we can conclude that for small samples, there is no statistically significant performance difference across the 2 headsets and traditional monitor. However, as the size of samples increases, the hypothesized pattern begins to emerge with all devices falling below the baseline in combined performance across distortions.

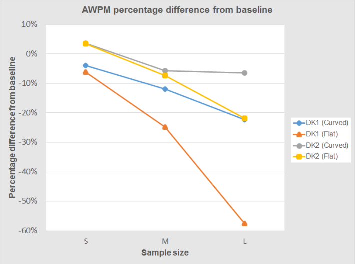

Figure 4: Mean AWPM percentage differences for all sample sizes and devices

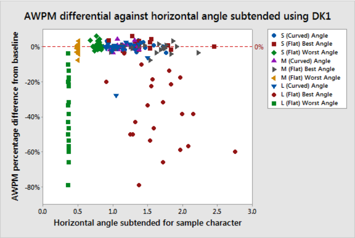

Figure 5: Horizontal angle subtended for a lowercase letter 'm' against participant running record accuracy compared to baseline for samples using the DK1 headset

For each participant and sample, the best and worst case subtended angles were calculated and compared against the differential reading accuracy of the sample from the baseline. This information can be visualized for the DK1 experiments

Table 6: Paired T-tests for distortions pairs, over the percentage difference from baseline of same sample size

Device	Size	μ(Flat)	μ(Curve)	μ(Difference)	95% CI	t	p
DK1	L	-57.59% ± 14.32	-22.25% ± 17.36	-35.33%	-43.15% to -27.52%	t(21) = -9.41	0.000
DK1	M	-24.87% ± 13.66	-12.00% ± 15.43	-12.87%	-18.45% to -7.30%	t(22) = -4.79	0.000
DK1	S	-6.12% ± 15.20	-3.90% ± 20.36	-2.22%	-9.50% to 5.05%	t(22) = -0.63	0.533
DK2	L	-21.88% ± 14.58	-6.51% ± 9.63	-15.37%	-21.93% to -8.82%	t(23) = -4.85	0.000
DK2	M	-7.39% ± 10.36	-5.70% ± 7.76	-1.69%	-6.25% to -2.86%	t(22) = -0.77	0.449
DK2	S	3.51% ± 11.85	3.57% ± 16.89	-0.06%	-7.12% to 7.00%	t(23) = -0.02	0.987

in Figure 5, and for the DK2 experiments in Figure 7.

A broad cluster of points can be seen in each graph centered around the 0% differential accuracy line, showing that while there was little difference in reading accuracy for most samples, there was notable variation in placements of samples inside the environment. For the best case values only, this cluster represents a mean best case value of $1.35° \pm 0.33$ for the DK1 samples, and $0.97° \pm 0.26$ for the DK2.

For the large flat panels on the DK1, it can be seen that although the best case angles are spread over a wide range and are larger than the best case mean, the differential accuracy decreased. This effect can be explained by the narrow worst case angles that accompany the best case options, which shows that while there was a large degree of angular variation in the best case, in the worst case the subtended angle is near constant. The flat medium panel worst case markers can be seen to the right of the large markers, and also show a near constant subtended angle, but with a greatly reduced differential accuracy score. This indicates that as the worst case angle gets narrower, the overall performance of the sample will decrease. This effect can be seen in the DK2 results also, with the large flat panel showing a narrow angle, with reduced accuracy, however the differential accuracy values are notably better than those expressed for the DK1 sample. This difference in differential accuracies despite the same sized samples occupying near identical angles across devices, suggests that the device specifications affect the reading performance, likely through differences in pixel density.

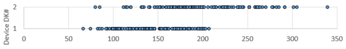

Figure 6: Calculated distance (mm) from the participant to the chosen panel position

With regards to hypothesis #5, we can conclude that while participants did not identify a single optimal location for placement of panels, they did identify a range of locations that would be suitable (see Figure: 6). Large flat panels presented as outliers with significant variance in placement, and an equally notable variation in accuracy differentials, suggesting that the size of these panels exceeds what would be considered reasonable to display at the headset resolution with no panel distortion.

5. EVALUATION

Over the duration of the user study, a number of issues were indicated by participants that fell outside of the scope the study was intending to address. The most prominent of these was a perceived 'fuzziness' of text, predominantly experienced on samples in the DK2. The most common description of this problem was seeing a shadow of color surrounding the text, causing the characters to be difficult to read because of the increased graphical noise. The cause of this on the DK2 appears to relate to Microsoft's ClearType, which applies subpixel rendering techniques to text in order to improve the look of text on LCD RGB displays[1]. As the samples that were displayed were taken as screenshot images from a system utilizing ClearType, fonts were decorated with color fringing (see Figure 8) that would not normally be readily noticeable on an LCD display due to the known arrangement of subpixels that the ClearType engine targets. The display in the DK2 however, is an RGBG PenTile display with subpixels rendered in a diamond pattern. As this layout is different from that which the ClearType engine is expecting to be rendering to, the effects are not what was intended, and a blurriness is observed.

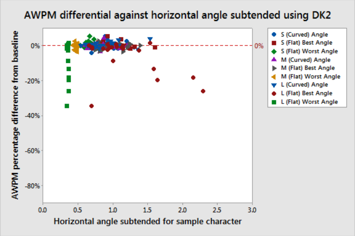

Figure 7: Horizontal angle subtended for a lowercase letter 'm' against participant running record accuracy compared to baseline for samples using the DK2 headset

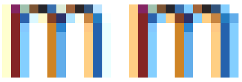

Figure 8: Letter 'm' extracted from a sample image (left), zoomed to display the ClearType color fringing. The same 'm' with faint colors amplified is displayed on the right

In addition to the color shadowing issue observed, some participants noted that they had difficulty reading when their head was stationary, but no problems reading when moving their heads. As the image samples being read were taken at a high resolution, these had to be scaled down to display on a smaller number of physical pixels inside the headset. This texture filtering process produces interesting results in conjunction with font-hinting and the kerning of certain characters, producing words that cannot be easily read. However, when moving the headset, the view is constantly being redrawn, and the texture filtering effects recalculated, leading to differing images being produced on each frame. This can be see in Figure 9, where the word 'while' is seen through the lenses of the DK2 over 9 consecutive frames. Issues can clearly be seen with displaying a white pixel column between the 'i' and 'l' characters, with some frames blurring these two characters together, and the 'h' and 'e' characters appear with many of their defining features obscured on some frames.

Figure 9: The word 'while' as seen over 9 frames through the lenses of the DK2

In order to address these issues in future experiments, trials with different fonts and the use of ClearType will be conducted. Sheedy *et al.*[11] note that while the effects of ClearType font hinting is preferred by readers, there is no improvement in text legibility, reading speed or comfort with it turned on or off. Due to this, it is likely that future experiments will be conducted with it off in order to remove a confounding factor.

Additional issues noted by some participants were varying degrees of double vision and chromic aberration, which were alleviated by adjusting the Inter-Pupilliary Distance (IPD) value. These IPD values were calculated using a utility that is a part of the Oculus Runtime, but this tool appears to suffer from imprecision, which becomes an issue when using the devices. McCleary[9] notes that this double vision 'prism' effect can be avoided with accurate measurements of the IPD using a digital pupilometer, and for future experiments we will attempt to utilize one of these in order to avoid these issues.

6. CONCLUSIONS & FUTURE WORK

Our results show that it is possible for users to perform traditional reading tasks inside a immersive virtual environment with near-baseline performance under ideal circumstances. We identified that rendering application windows onto flat panels significantly impaired reading performance on larger panels, due to the effects of individual characters occupying a narrower subtended angle of vision when rendered near the edge of the panel when compared to characters displayed at the center. This, coupled with the discovery that curved distortion had no notable effect on performance when compared to small flat panels, suggests that future work should focus on utilizing curved panels exclusively. It was found that there were significant performance differences when comparing the results for either headset, with the DK2 performing better than the same sample on the DK1, which when combined with the positional tracking functionality the DK2 provides over the DK1, allows us to pursue the DK2 as our primary development target.

As development of the testing application continues, information gleaned from this study will direct the development, and will allow us to make decisions based on quantitative data. Future investigations will be performed, investigating issues identified by this study, with an emphasis on font-hinting techniques, and their suitability for use inside these environments.

7. REFERENCES

[1] Microsoft Research, ClearType. http://research.microsoft.com/en-us/projects/ClearType/.

[2] 2014. Virtual reality headset Oculus Rift meets the Bloomberg terminal. http://qz.com/218129.

[3] J. C. Alderson. *Assessing Reading*. Cambridge Language Assessment Series. Cambridge University Press, 2000.

[4] M. Billinghurst, J. Bowskill, N. Dyer, and J. Morphett. An evaluation of wearable information spaces. In *Virtual Reality Annual International Symposium, 1998. Proceedings., IEEE 1998*, pages 20–27, 1998.

[5] M. M. Clay. *Running Records for Classroom Teachers*. 2000.

[6] M. Czerwinski, G. Smith, T. Regan, B. Meyers, G. Robertson, and G. Starkweather. Toward Characterizing the Productivity Benefits of Very Large Displays. *Interact*, 3:pages 9–16, 2003.

[7] S. S. Fisher, M. McGreevy, J. Humphries, and W. Robinett. Virtual environment display system. In *Proceedings of the 1986 workshop on Interactive 3D graphics*, I3D '86, pages 77–87, New York, NY, USA, 1987. ACM.

[8] G. Lawton. Making virtual reality more accessible. *Computer*, 39(6):pages 12–15, 2006.

[9] D. S. McCleary. *The Optician Training Manual*. Santa Rosa Publishing, Temecula, CA, 1 edition edition, Aug. 2009.

[10] B. A. Reichlen. Sparcchair: A one hundred million pixel display. In *Virtual Reality Annual International Symposium*, volume 0, pages 300–307, Los Alamitos, CA, USA, 1993. IEEE Computer Society.

[11] J. Sheedy, Y.-C. Tai, M. Subbaram, S. Gowrisankaran, and J. Hayes. ClearType sub-pixel text rendering: Preference, legibility and reading performance. *Displays*, 29(2):pages 138–151, Mar. 2008.

[12] M. van Dantzich, V. Gorokhovsky, and G. Robertson. Application redirection: hosting Windows applications in 3d. In *Proceedings of the 1999 workshop on new paradigms in information visualization and manipulation in conjunction with the eighth ACM internation conference on Information and knowledge management*, NPIVM '99, pages 87–91, New York, NY, USA, 1999. ACM.

[13] S. White, J. Chen, and B. Forsyth. Reading-Related Literacy Activities of American Adults: Time Spent, Task Types, and Cognitive Skills Used. *Journal of Literacy Research*, 42(3):pages 276–307, Sept. 2010.

Appendix C
Reading Study

This appendix contains details of the ethical consent application submitted to support the user study reported in Chapter 6, as well as material supplementary to the study.

- Ethical Approval Letter from the Human Research Ethics Committee of the Faculty of Computing and Mathematical Sciences at the University of Waikato, dated 30 January 2015. (Appendix C.1)
- Updated Ethical Approval Letter from the Human Research Ethics Committee of the Faculty of Computing and Mathematical Sciences at the University of Waikato, dated 13 February 2015, noting the addition of audio recording to the experiment. (Appendix C.2)
- Participant Information Sheet, which each study participant received and had explained during participant recruitment. (Appendix C.3)
- Research Consent Form, which each participant signed prior to beginning the study. (Appendix C.4)
- Result Sheet, form completed by the researcher during and following participant device trials. Head movement magnitude is recorded on a 0–5 scale. Time taken is calculated during evaluation of recorded results. (Appendix C.5)
- Post-test questionnaire, completed by participants following each device trial undertaken. Discomfort scale is 0–5 for normal feedback, with additional option 6 to indicate that the discomfort resulted in an aborted experiment. (Appendix C.6)
- Exit interview, completed by participants at the conclusion of the study. (Appendix C.7)
- Study Texts, text samples used in the study, extracted from the page in which they were displayed. Following sample coding, the flesch reading ease (FRE) score and sample word count is given. (Appendix C.8)

C.1 Ethical Approval Letter

Faculty of Computing and
Mathematical Sciences
Rorohiko me ngā Pūtaiao Pāngarau
The University of Waikato
Private Bag 3105
Hamilton
New Zealand

Phone +64 7 838 4322
www.fcms.waikato.ac.nz

THE UNIVERSITY OF
WAIKATO
Te Whare Wānanga o Waikato

30 January 2015

Cameron Grout
C/- Department of Computer Science
THE UNIVERSITY OF WAIKATO

Dear Cameron

Application for approval under the Ethical Conduct in Human Research and Related Activities Regulations

I have considered your request to add in to your previous applicaton for audio recording for your research project involving human participants entitled “Extended workspaces through virtual reality”. The purpose is to identify a number of factors and settings for virtual reality desktop enviroments that will be used in later studies.

I note that you will be alerting prospective participants that minor symptoms of motion sickness may occur while using virtual reality headsets, and in that occurs you will halt the study. Participants will be free to withdraw at any stage.

The procedure described in your request is acceptable. Participants involved in the study will not be identified in any resulting publications. At the conclusion of the project the notes will be destroyed and a copy of the anonymized results will be submitted to the FCMS Data Archive repository for 5 years following the conlusion of the data analysis.

The Participant Information Sheet, Research Consent Form and questionnairs comply with the requirements of the University's human research ethics policies and procedures.

I therefore approve your application to perform the research project.

Yours sincerely

LA Hunt.

Lyn Hunt
Human Research Ethics Committee
Faculty of Computing and Mathematical Sciences

C.2 Updated Ethical Approval Letter

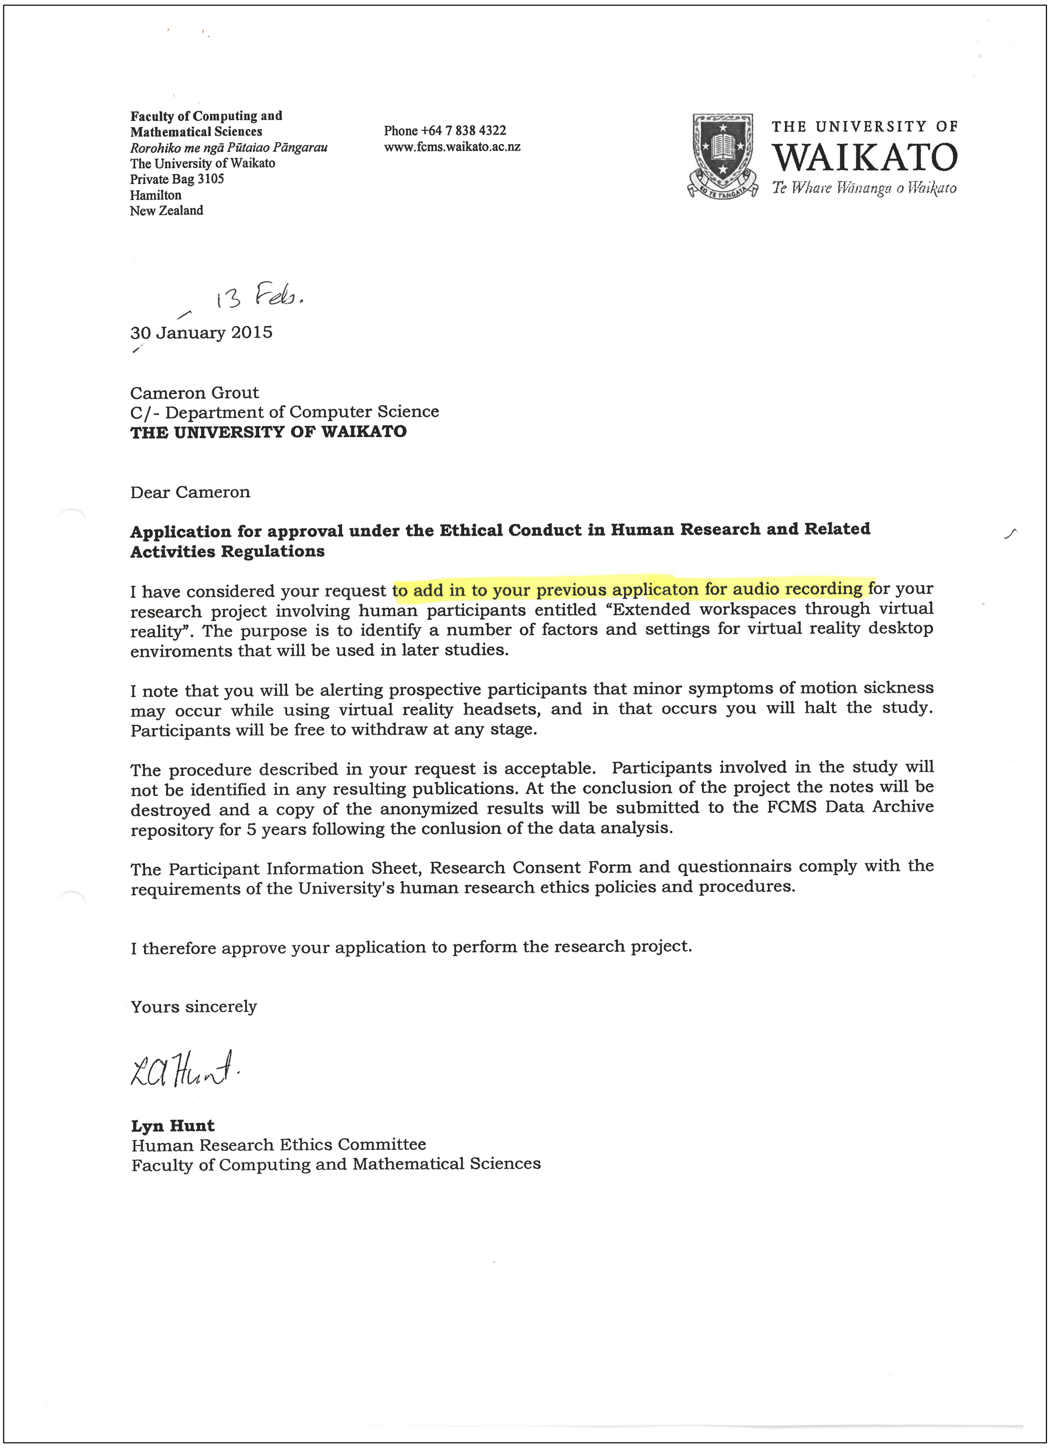

Faculty of Computing and
Mathematical Sciences
Rorohiko me ngā Pūtaiao Pāngarau
The University of Waikato
Private Bag 3105
Hamilton
New Zealand

Phone +64 7 838 4322
www.fcms.waikato.ac.nz

THE UNIVERSITY OF
WAIKATO
Te Whare Wānanga o Waikato

13 Feb.

30 January 2015

Cameron Grout
C/- Department of Computer Science
THE UNIVERSITY OF WAIKATO

Dear Cameron

Application for approval under the Ethical Conduct in Human Research and Related Activities Regulations

I have considered your request to add in to your previous applicaton for audio recording for your research project involving human participants entitled "Extended workspaces through virtual reality". The purpose is to identify a number of factors and settings for virtual reality desktop enviroments that will be used in later studies.

I note that you will be alerting prospective participants that minor symptoms of motion sickness may occur while using virtual reality headsets, and in that occurs you will halt the study. Participants will be free to withdraw at any stage.

The procedure described in your request is acceptable. Participants involved in the study will not be identified in any resulting publications. At the conclusion of the project the notes will be destroyed and a copy of the anonymized results will be submitted to the FCMS Data Archive repository for 5 years following the conlusion of the data analysis.

The Participant Information Sheet, Research Consent Form and questionnairs comply with the requirements of the University's human research ethics policies and procedures.

I therefore approve your application to perform the research project.

Yours sincerely

LA Hunt.

Lyn Hunt
Human Research Ethics Committee
Faculty of Computing and Mathematical Sciences

C.3 Participant Information Sheet

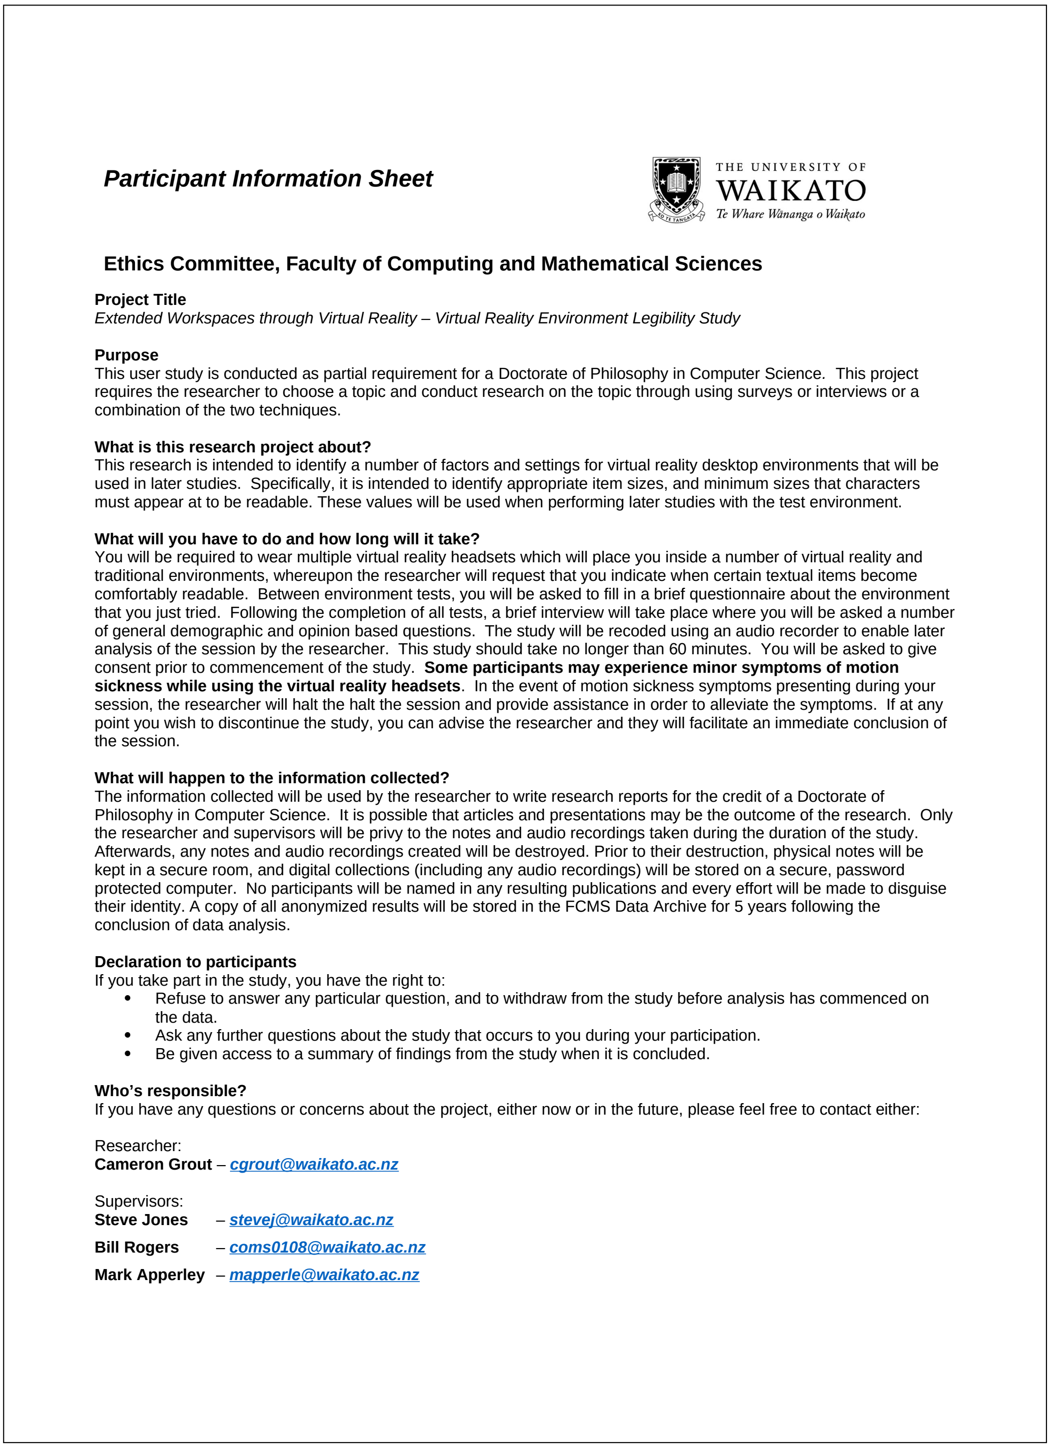

Participant Information Sheet

THE UNIVERSITY OF
WAIKATO
Te Whare Wānanga o Waikato

Ethics Committee, Faculty of Computing and Mathematical Sciences

Project Title
Extended Workspaces through Virtual Reality – Virtual Reality Environment Legibility Study

Purpose
This user study is conducted as partial requirement for a Doctorate of Philosophy in Computer Science. This project requires the researcher to choose a topic and conduct research on the topic through using surveys or interviews or a combination of the two techniques.

What is this research project about?
This research is intended to identify a number of factors and settings for virtual reality desktop environments that will be used in later studies. Specifically, it is intended to identify appropriate item sizes, and minimum sizes that characters must appear at to be readable. These values will be used when performing later studies with the test environment.

What will you have to do and how long will it take?
You will be required to wear multiple virtual reality headsets which will place you inside a number of virtual reality and traditional environments, whereupon the researcher will request that you indicate when certain textual items become comfortably readable. Between environment tests, you will be asked to fill in a brief questionnaire about the environment that you just tried. Following the completion of all tests, a brief interview will take place where you will be asked a number of general demographic and opinion based questions. The study will be recoded using an audio recorder to enable later analysis of the session by the researcher. This study should take no longer than 60 minutes. You will be asked to give consent prior to commencement of the study. **Some participants may experience minor symptoms of motion sickness while using the virtual reality headsets**. In the event of motion sickness symptoms presenting during your session, the researcher will halt the halt the session and provide assistance in order to alleviate the symptoms. If at any point you wish to discontinue the study, you can advise the researcher and they will facilitate an immediate conclusion of the session.

What will happen to the information collected?
The information collected will be used by the researcher to write research reports for the credit of a Doctorate of Philosophy in Computer Science. It is possible that articles and presentations may be the outcome of the research. Only the researcher and supervisors will be privy to the notes and audio recordings taken during the duration of the study. Afterwards, any notes and audio recordings created will be destroyed. Prior to their destruction, physical notes will be kept in a secure room, and digital collections (including any audio recordings) will be stored on a secure, password protected computer. No participants will be named in any resulting publications and every effort will be made to disguise their identity. A copy of all anonymized results will be stored in the FCMS Data Archive for 5 years following the conclusion of data analysis.

Declaration to participants
If you take part in the study, you have the right to:

- Refuse to answer any particular question, and to withdraw from the study before analysis has commenced on the data.
- Ask any further questions about the study that occurs to you during your participation.
- Be given access to a summary of findings from the study when it is concluded.

Who's responsible?
If you have any questions or concerns about the project, either now or in the future, please feel free to contact either:

Researcher:
Cameron Grout – ***cgrout@waikato.ac.nz***

Supervisors:
Steve Jones – ***stevej@waikato.ac.nz***
Bill Rogers – ***coms0108@waikato.ac.nz***
Mark Apperley – ***mapperle@waikato.ac.nz***

C.4 Research Consent Form

Research Consent Form

THE UNIVERSITY OF
WAIKATO
Te Whare Wānanga o Waikato

Ethics Committee, Faculty of Computing and Mathematical Sciences

Extended Workspaces through Virtual Reality – Virtual Reality Environment Legibility Study

Consent Form for Participants

I have read the **Participant Information Sheet** for this study and have had the details of the study explained to me. My questions about the study have been answered to my satisfaction, and I understand that I may ask further questions at any time.

I also understand that I am free to withdraw from the study before analysis has commenced on the data, or to decline to answer any particular questions in the study. I understand I can withdraw any information I have provided up until the researcher has commenced analysis on my data. I agree to provide information to the researchers under the conditions of confidentiality set out on the **Participant Information Sheet**.

I agree to participate in this study under the conditions set out in the **Participant Information Sheet**.

Signed: ____

Name: ____

Date: ____

Researcher contact information:

Cameron Grout – ***cgrout@waikato.ac.nz***

Supervisor’s contact information:

Steve Jones – ***stevej@waikato.ac.nz***
Bill Rogers – ***coms0108@waikato.ac.nz***
Mark Apperley – ***mapperle@waikato.ac.nz***

C.5 Result Sheet

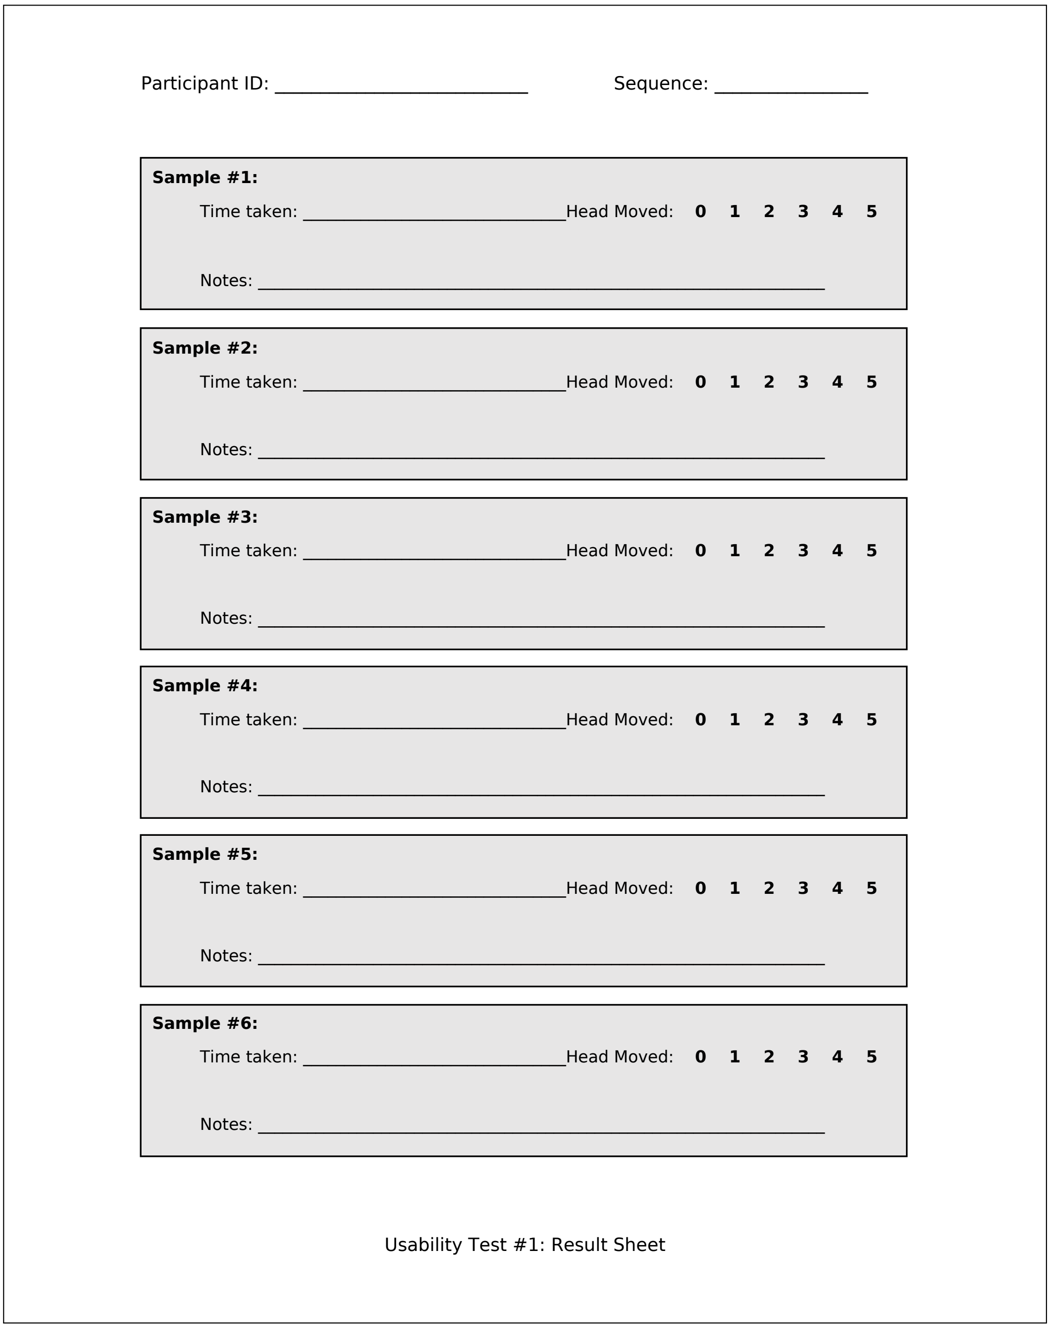

Participant ID: ______________________________ Sequence: __________________

Sample #1:

Time taken: ________________________________ Head Moved: **0 1 2 3 4 5**

Notes: ____

Sample #2:

Time taken: ________________________________ Head Moved: **0 1 2 3 4 5**

Notes: ____

Sample #3:

Time taken: ________________________________ Head Moved: **0 1 2 3 4 5**

Notes: ____

Sample #4:

Time taken: ________________________________ Head Moved: **0 1 2 3 4 5**

Notes: ____

Sample #5:

Time taken: ________________________________ Head Moved: **0 1 2 3 4 5**

Notes: ____

Sample #6:

Time taken: ________________________________ Head Moved: **0 1 2 3 4 5**

Notes: ____

Usability Test #1: Result Sheet

C.6 Post-test questionnaire

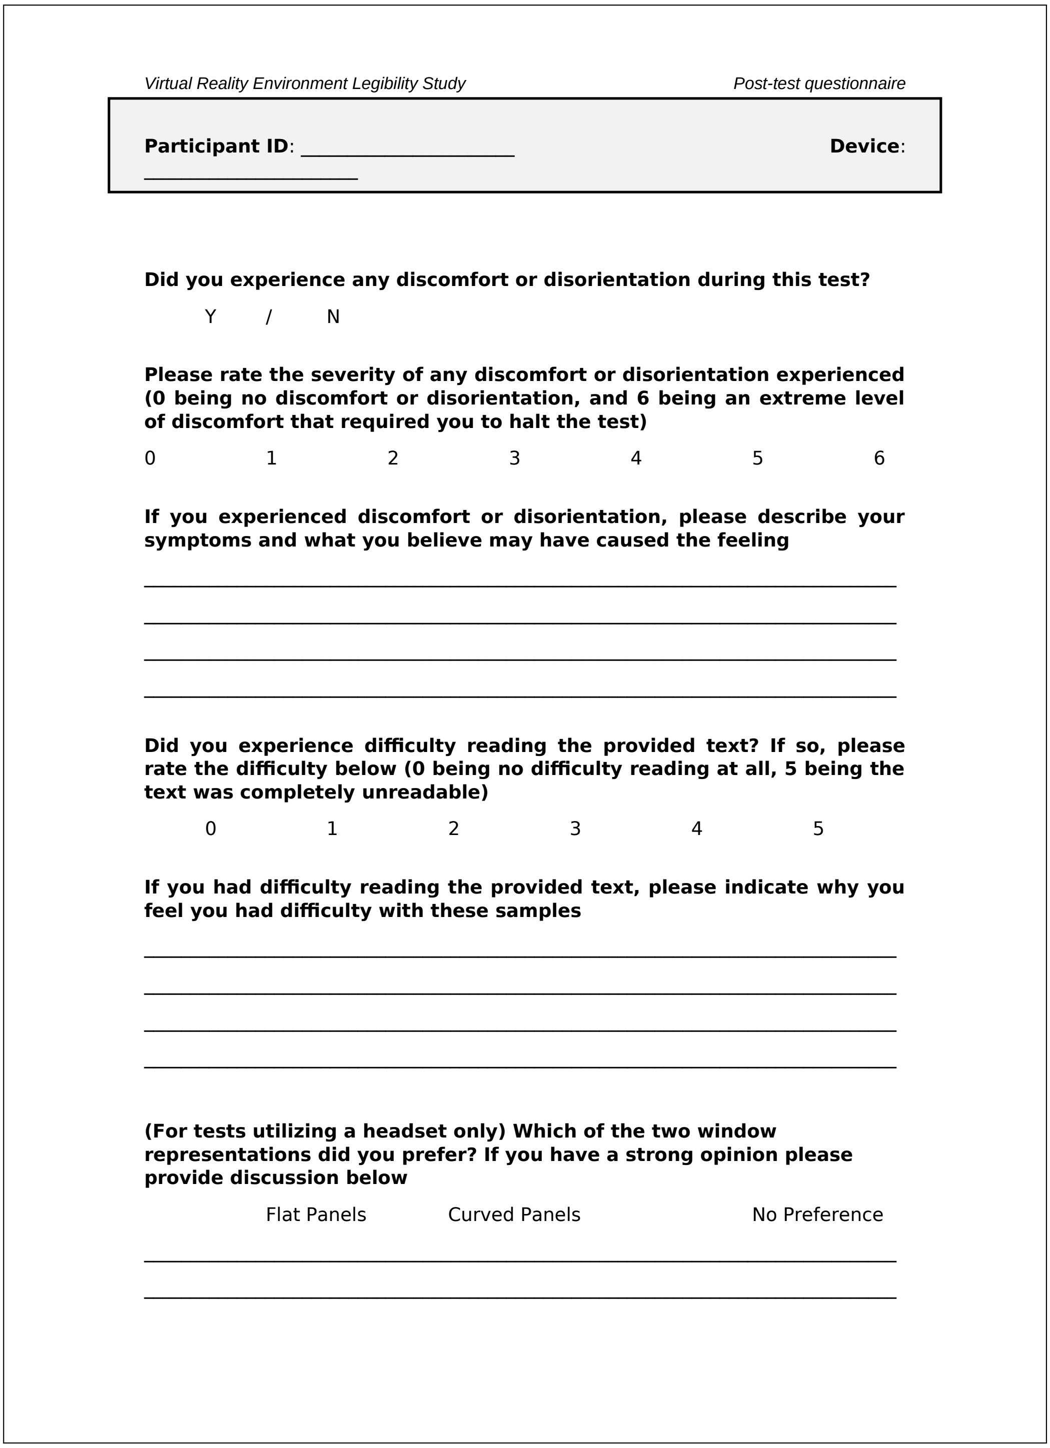

Virtual Reality Environment Legibility Study *Post-test questionnaire*

Participant ID: ____________________ **Device**: ____________________

Did you experience any discomfort or disorientation during this test?

Y / N

Please rate the severity of any discomfort or disorientation experienced (0 being no discomfort or disorientation, and 6 being an extreme level of discomfort that required you to halt the test)

0 1 2 3 4 5 6

If you experienced discomfort or disorientation, please describe your symptoms and what you believe may have caused the feeling

____

____

____

____

Did you experience difficulty reading the provided text? If so, please rate the difficulty below (0 being no difficulty reading at all, 5 being the text was completely unreadable)

0 1 2 3 4 5

If you had difficulty reading the provided text, please indicate why you feel you had difficulty with these samples

____

____

____

____

(For tests utilizing a headset only) Which of the two window representations did you prefer? If you have a strong opinion please provide discussion below

Flat Panels Curved Panels No Preference

____

____

C.7 Exit interview

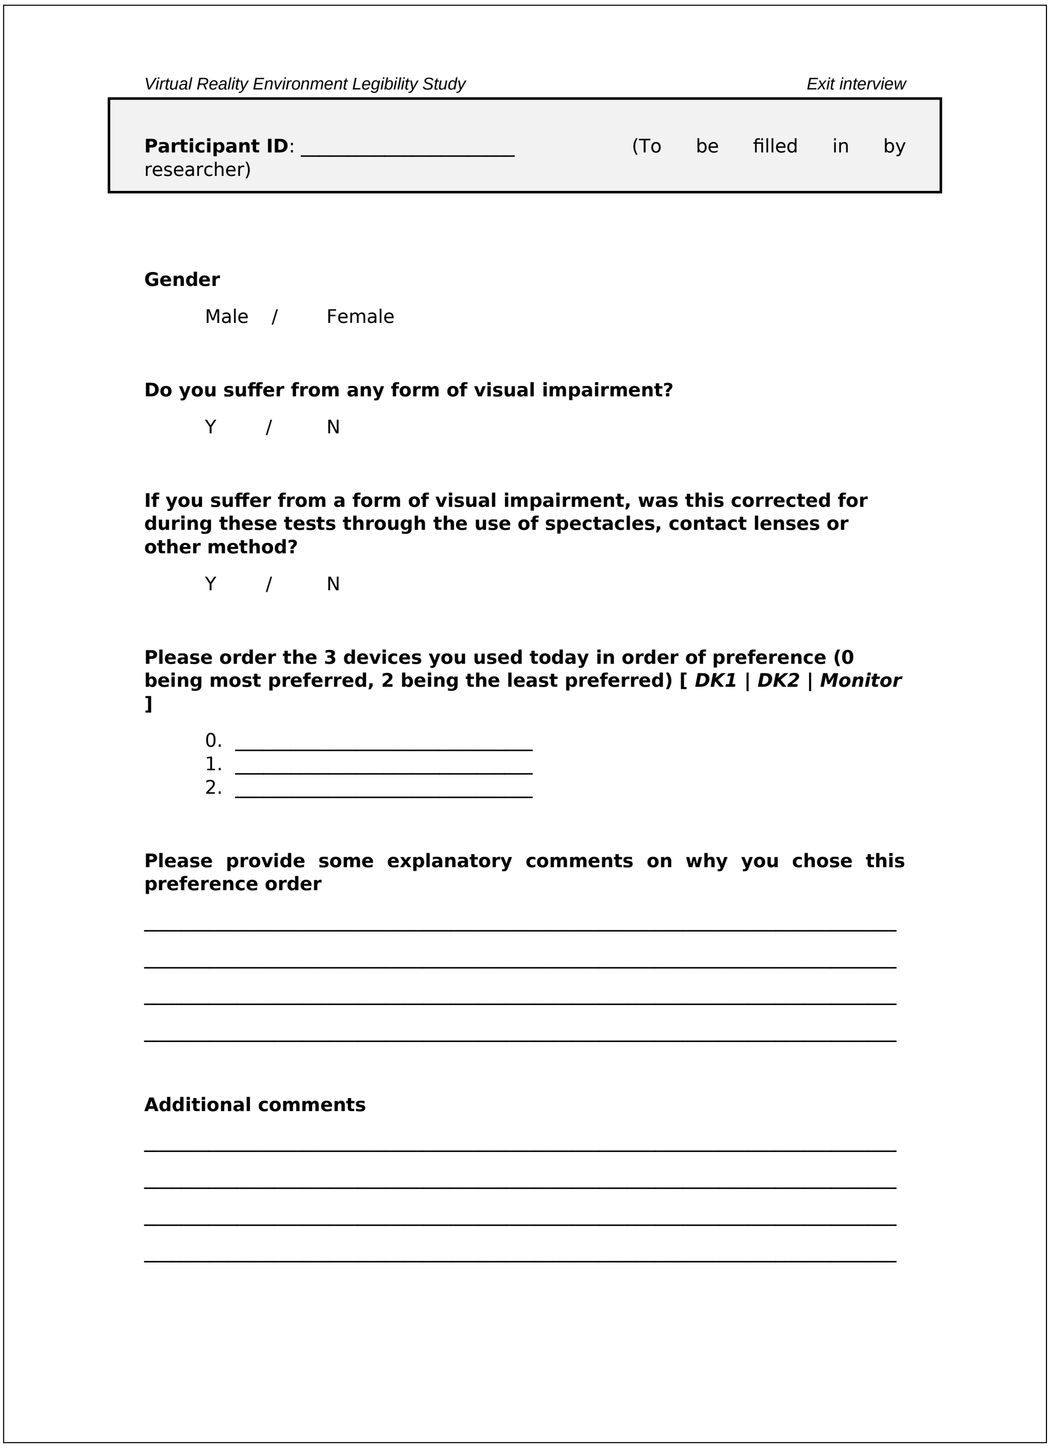

Virtual Reality Environment Legibility Study *Exit interview*

Participant ID: ______________________ (To be filled in by researcher)

Gender

Male / Female

Do you suffer from any form of visual impairment?

Y / N

If you suffer from a form of visual impairment, was this corrected for during these tests through the use of spectacles, contact lenses or other method?

Y / N

Please order the 3 devices you used today in order of preference (0 being most preferred, 2 being the least preferred) [*DK1* | *DK2* | *Monitor*]

0. ______________________________
1. ______________________________
2. ______________________________

Please provide some explanatory comments on why you chose this preference order

____
____
____
____

Additional comments

____
____
____
____

C.8 Study Texts

Sequence	Size	Number	FRE	Word Count
0	Small	0	60.1	116

Sydney is a city on the east coast of Australia. Sydney is the capital city of New South Wales. About four million people live in Sydney which makes it the biggest city in Australia. Sydney started in 1788, when the Captain Arthur Phillip brought the First Fleet to settle in Australia. The settlers were mostly convicts from crowded prisons in England and Ireland, with a group of soldiers to guard them.

In Sydney, there are many famous buildings: the Sydney Opera House, the Queen Victoria Building and the Sydney Harbour Bridge. Sydney has a large harbour and many beaches. The most famous beach is Bondi Beach, some other famous beaches are Coogee Beach and Manly Beach.

Sequence	Size	Number	FRE	Word Count
0	Small	1	64.7	127

London is by far the largest city in England, the United Kingdom and middle Europe, and is the world's largest financial centre. 8.3 million people live in London, which is located on the River Thames. It is the capital of the United Kingdom.

London was founded by the Romans in AD 43 and called Londinium. London features the world's first underground railway system, the London Underground.

For a long time, London was a small city, with all of its people lived inside walls that were built by the Romans. This area is still known as the City of London. There were many villages around the city. Gradually, more and more people came to live there. Finally, step by step, the villages joined together into one huge city.

Sequence	Size	Number	FRE	Word Count
0	Medium	0	64.5	121

Portland is the largest city in the U.S. state of Oregon. It is the county seat of Multin County. About 500 thousand people live in the city of Portland. About 1.9 million people live in the city's metropolitan area. It is found in the north part of the state, where the William River meets the Columbia River. Portland has the second most people of any city in the Northwest United States, after the city of Seattle.

Portland has many nicknames. One of these is ”The City of Roses,” because roses grow well there. Some other nicknames of Portland are ”Stumptown” from the fact that Portland was built over a forest, and ”Puddletown” from the fact that it often rains in Portland.

Sequence	Size	Number	FRE	Word Count
0	Medium	1	61.0	128

Denver is the capital and largest city in the U.S. state of Colorado. It is located at the foot of the Rocky Mountains on the South Platte River, and was founded in 1858. A nickname for Denver is The Mile High City, because Denver is very high above sea level. It is at least 1,600 meters above sea level at a point in the city's state capitol building. The dome at the top of the capitol building is covered in gold.

Denver City was founded in November 1858 during the Pikes Peak Gold Rush as a mining town in western Kansas Territory. That summer, a group of gold prospectors from Lawrence, Kansas, started a settlement on the banks of the South Platte River. They called it Montana City.

Sequence	Size	Number	FRE	Word Count
0	Large	0	64.5	129

Chicago was founded in the early 1700s by John Baptist Point. The city was founded to create a canal that allowed steamboats and sailing ships on the Great Lakes to connect to the Mississippi River. The city later became a trading center for food, crops, and fur. The city grew very fast because of how the river back then was clean and healthy to drink. In 1837, Chicago became a city. The city grew until the Great Chicago Fire happened in 1871. The fire lasted for almost a week. Almost half the city and its population were lost in the fire. After the fire, Chicago grew faster than ever.

Also after the fire happened the city's economy grew and also more people migrated here from parts of the world.

Sequence	Size	Number	FRE	Word Count
0	Large	1	61.1	125

New York City was settled by Europeans from The Netherlands in 1624. The Dutch called the whole area of New York "New Netherland" and they named a fort and town on the south end of Manhattan Island New Amsterdam, after the capital city of the Netherlands, which was to become present-day New York. The English took over the colony in 1664 during the second Dutch War. They changed the name to New York, to honor the Duke of York, who later became King James II of England. The Dutch gave up New Amsterdam without fighting. They were afraid of the English Royal Navy, so they traded the town to England for the colony of Surinam in South America, which they thought was worth more money.

Sequence	Size	Number	FRE	Word Count
1	Small	0	61.7	121

Baltimore is the biggest city in the U.S. state of Maryland. In 2010 it had about 620 thousand people living there. It is not in any county, so it is called an independent city. It is next to the Chesapeake Bay and used to be an important port for trade by ships. There is still some shipping but the Inner Harbor is now mostly famous for shopping and restaurants, and also for the National Aquarium and other museums.

Baltimore is home to the Ravens (football) and Oreos (baseball), both sports teams. Its main newspaper is the Sun. There is a place for horse races in the city called Pimco. There are many colleges and universities in Baltimore, like Johns Hopkins University.

Sequence	Size	Number	FRE	Word Count
1	Medium	0	60.2	122

Atlanta is the capital and largest city of the U.S. state of Georgia. It is one of the South's largest cities. Atlanta is known as a major business city. It is the home of Coca-Cola, CNN, AT&T, and Home Depot, as well as many other Fortune 500 companies. Coca-Cola, which is made in Atlanta, is drunk all over the world and its factory is a favorite tourist spot. CNN is also a big tourist attraction. Atlanta's airport, called the Jackson International Airport, is the busiest airport in the world. Atlanta is near the center of Georgia and is on the Chatta River.

Atlanta was built on Native American land. It was called Terminus until 1843, when the name was changed to Martha.

Sequence	Size	Number	FRE	Word Count
1	Large	0	64.7	127

The first people settled at the place where Charlotte is in 1755 when a man named Thomas Polk built a house near two Native American trading paths. More people started living in the area and in 1768 it became a town named Charlotte Town. It was named after the wife of King George III because the people wanted him to like them. But he did not, and soon he started passing laws that the people in Charlotte did not like. So, on May 20, 1775, the people in Charlotte signed a proclamation that later was called the Charlotte Declaration of Independence. They did not want to be ruled by the king anymore so eleven days later they had a meeting and made new laws for their town.

Sequence	Size	Number	FRE	Word Count
2	Small	0	63.9	130

Seattle is the largest city in the U.S. state of Washington. It is the home of the Space Needle and a monorail, both of which were built for the 1962 World's Fair. It is the American headquarters of Starbucks, and Amazon. In the 1980s and 1990s, grunge music artists like Nirvana, Pearl Jam, and others from the city became popular. It is also the setting of the TV shows Here Come the Bride, and Grey's Anatomy.

Seattle has many sports teams, including the Mariners (baseball), and the Seahawks (football). Seattle has a lot of water around it, with the Pacific Ocean to the west and Lake Washington to the east. About 600 thousand people live in the city. More than 3 million people live in the city or near it.

Sequence	Size	Number	FRE	Word Count
2	Small	1	60.4	120

Washington, D.C. is the capital city of the United States. It is not a state or in a state. The President of the United States and many major government offices are in the city.

Washington was named after the first U.S. President, George Washington. The ”D.C.” stands for ”District of Columbia”, a special area created that is not a state. At first, it was made up of a piece from Virginia south of the Potomac River and a piece from Maryland north of the Potomac River. In 1847, Virginia’s piece was given back to it, and is now and part of the city of Alexandria. Since 1847, all of Washington D.C. is on the north side of the Potomac River.

Sequence	Size	Number	FRE	Word Count
2	Medium	0	61.1	117

San Francisco is a city in the U.S. state of California. It is famous for the Golden Gate Bridge. It has 744 thousand people living in it. San Francisco is the 4th largest city in California. It is in the top part of California between the Pacific Ocean and the San Francisco Bay.

San Francisco was founded in 1776 by the Spanish people. It was called ”Good Herb”, because a lot of mint grew there. After the Mexican War, Good Mint was taken over by the U.S. In 1848 it was renamed ”San Francisco” and became a city in 1850. The city is famous for its many internet companies and being home to a large gay population.

Sequence	Size	Number	FRE	Word Count
2	Medium	1	62.1	125

Cleveland is a city in northern Ohio, U.S. It is home to over 400 thousand people. It was named for General Moses Cleaveland in 1796, but a mistake in a local newspaper left out the first letter ”a” in its name, which is why it is spelt like it is today. Its metropolitan area makes it the largest urban division in the state.

Cleveland is home to the Cleveland Orchestra.

Professional sports teams that make their home in or near Cleveland include the Indians, who play baseball; the Browns, who play football; and the Cavaliers, who play basketball. Both the Indians and the Cavaliers have their games at the Gateway District in Downtown while the Browns have a stadium on the shores of Lake Eerie.

Sequence	Size	Number	FRE	Word Count
2	Large	0	63.1	121

San Diego is the second largest city in the U.S. state of California and 8th largest in the U.S. It is at the bottom corner of California, as well as the bottom corner of the U.S. It was founded in 1769 and it is the oldest city in California.

It has nice weather most of the year. There are many military bases in and near San Diego. It has a lot of Pacific Ocean beaches. The San Diego Zoo is very famous.

San Diego is home to San Diego State University and the University of California, San Diego.

San Diego is on the border between the United States and Mexico.

San Diego is home to the San Diego Chargers football team.

Sequence	Size	Number	FRE	Word Count
2	Large	1	64.2	128

Detroit is the largest city in the state of Michigan in the U.S. It was the 10th largest city in the U.S. in the year 2000. In 2004, it fell to 11th biggest as many people have moved away. Detroit city has a population of 912 thousand people. Nearly six million people live in Detroit and the surrounding areas. The city borders Windsor, Ontario in Canada. The border between Detroit and Windsor is one of the most crossed in the world.

Detroit is a city where many cars are made and this is why it is sometimes called the ”Motor City”. Many people call it the car capital of the world. General Motors, Ford and Chrysler have their offices and many of their plants in and around Detroit.

Appendix D
Tuning Text

This appendix contains details of the ethical consent application submitted to support the user study reported in Chapter 7, as well as supplemental data collected from the same user study.

- Ethical Approval Letter from the Human Research Ethics Committee of the Faculty of Computing and Mathematical Sciences at the University of Waikato, dated 11 December 2015. (Appendix D.1)
- Participant Information Sheet, which each study participant received and had explained during participant recruitment. (Appendix D.2)
- Research Consent Form, which each participant signed prior to beginning the study. (Appendix D.3)
- Participant demographic information. (Appendix D.4)
- Text sample screenshots. (Appendix D.5)
- Raw sample rating results, grouped by sample attribute permutations. (Appendix D.6)
- Processed $\mathbb{D}_n$ and $\mathbb{D}_f$ values for participants. (Appendix D.7)

D.1 Ethical Approval Letter

Faculty of Computing and
Mathematical Sciences
Rorohiko me ngā Pūtaiao Pāngarau
The University of Waikato
Private Bag 3105
Hamilton
New Zealand

Phone +64 7 838 4322
www.fcms.waikato.ac.nz

THE UNIVERSITY OF
WAIKATO
Te Whare Wānanga o Waikato

11 December 2015

Cameron Grout
C/- Department of Computer Science
THE UNIVERSITY OF WAIKATO

Dear Cameron

Application for approval under the Ethical Conduct in Human Research and Related Activities Regulations

I have considered your application for a research project involving human participants entitled "Extended workspaces through virtual reality – virtual reality environment study". The purpose is to identify a number of factors and settings for virtual reality desktop enviroments that will be used in later studies.

I note that you will be alerting prospective participants that minor symptoms of motion sickness may occur while using virtual reality headsets, and in that occurs you will halt the study. Participants will be free to withdraw at any stage.

The procedure described in your request is acceptable. Participants involved in the study will not be identified in any resulting publications. At the conclusion of the project the notes will be destroyed and a copy of the anonymized results will be submitted to the FCMS Data Archive repository for 5 years following the conlusion of the data analysis.

The Participant Information Sheet and Research Consent Form comply with the requirements of the University's human research ethics policies and procedures.

I therefore approve your application to perform the research project.

Yours sincerely

Bernhard Pfahringer
Human Research Ethics Committee
Faculty of Computing and Mathematical Sciences

D.2 Participant Information Sheet

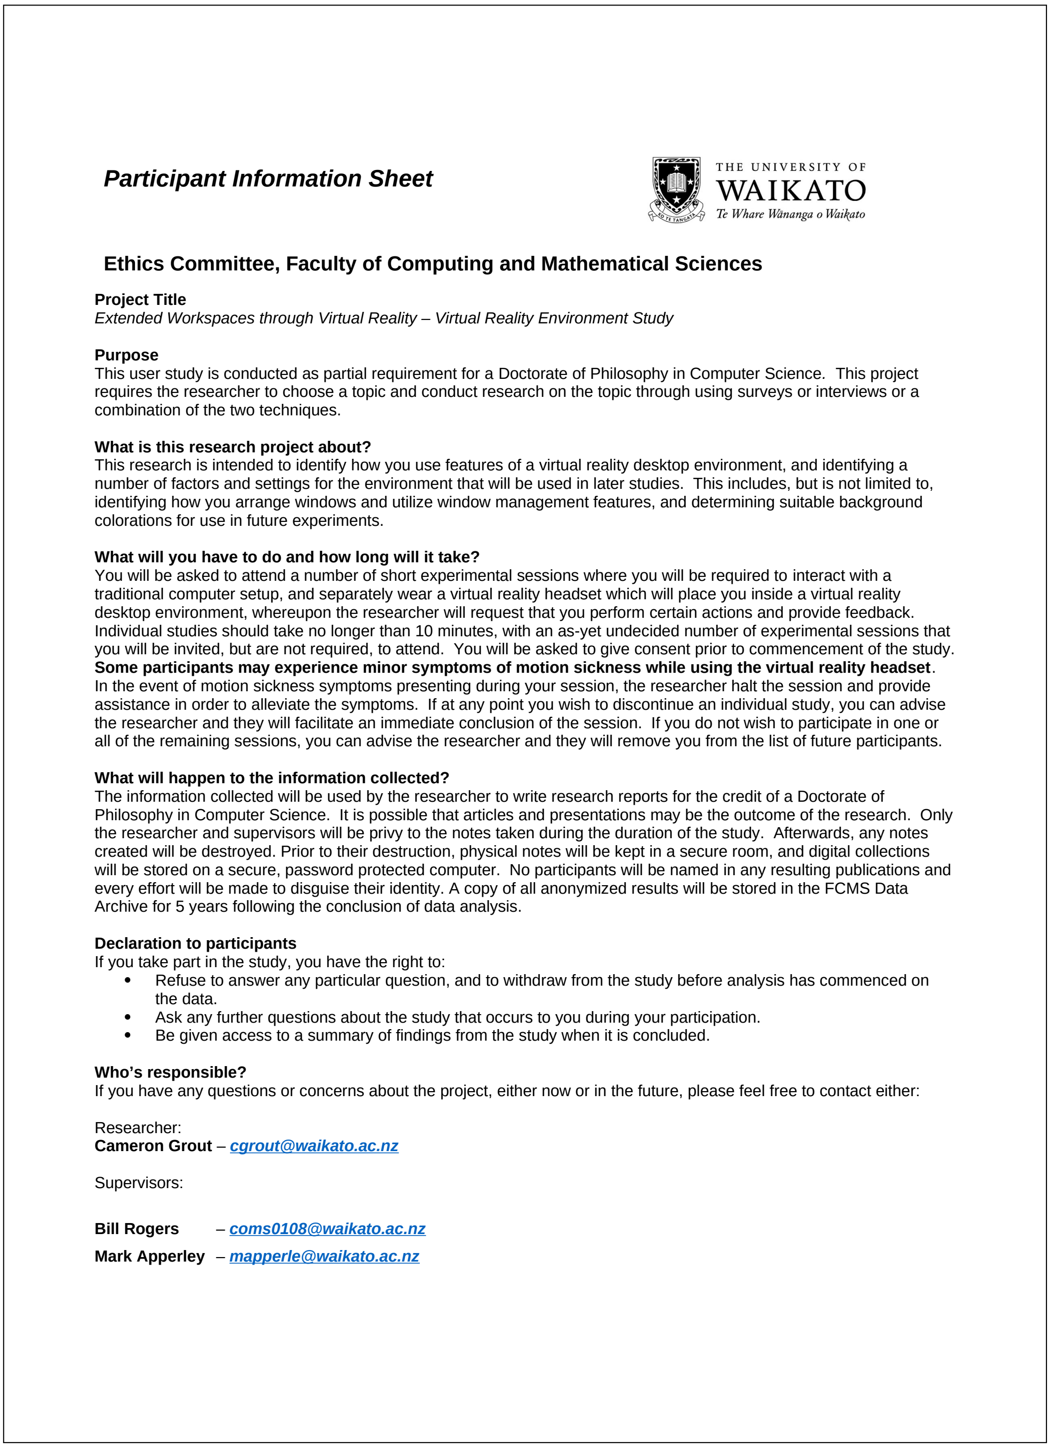

Participant Information Sheet

THE UNIVERSITY OF
WAIKATO
Te Whare Wānanga o Waikato

Ethics Committee, Faculty of Computing and Mathematical Sciences

Project Title
Extended Workspaces through Virtual Reality – Virtual Reality Environment Study

Purpose
This user study is conducted as partial requirement for a Doctorate of Philosophy in Computer Science. This project requires the researcher to choose a topic and conduct research on the topic through using surveys or interviews or a combination of the two techniques.

What is this research project about?
This research is intended to identify how you use features of a virtual reality desktop environment, and identifying a number of factors and settings for the environment that will be used in later studies. This includes, but is not limited to, identifying how you arrange windows and utilize window management features, and determining suitable background colorations for use in future experiments.

What will you have to do and how long will it take?
You will be asked to attend a number of short experimental sessions where you will be required to interact with a traditional computer setup, and separately wear a virtual reality headset which will place you inside a virtual reality desktop environment, whereupon the researcher will request that you perform certain actions and provide feedback. Individual studies should take no longer than 10 minutes, with an as-yet undecided number of experimental sessions that you will be invited, but are not required, to attend. You will be asked to give consent prior to commencement of the study. **Some participants may experience minor symptoms of motion sickness while using the virtual reality headset**. In the event of motion sickness symptoms presenting during your session, the researcher halt the session and provide assistance in order to alleviate the symptoms. If at any point you wish to discontinue an individual study, you can advise the researcher and they will facilitate an immediate conclusion of the session. If you do not wish to participate in one or all of the remaining sessions, you can advise the researcher and they will remove you from the list of future participants.

What will happen to the information collected?
The information collected will be used by the researcher to write research reports for the credit of a Doctorate of Philosophy in Computer Science. It is possible that articles and presentations may be the outcome of the research. Only the researcher and supervisors will be privy to the notes taken during the duration of the study. Afterwards, any notes created will be destroyed. Prior to their destruction, physical notes will be kept in a secure room, and digital collections will be stored on a secure, password protected computer. No participants will be named in any resulting publications and every effort will be made to disguise their identity. A copy of all anonymized results will be stored in the FCMS Data Archive for 5 years following the conclusion of data analysis.

Declaration to participants
If you take part in the study, you have the right to:

- Refuse to answer any particular question, and to withdraw from the study before analysis has commenced on the data.
- Ask any further questions about the study that occurs to you during your participation.
- Be given access to a summary of findings from the study when it is concluded.

Who's responsible?
If you have any questions or concerns about the project, either now or in the future, please feel free to contact either:

Researcher:
Cameron Grout – ***cgrout@waikato.ac.nz***

Supervisors:

Bill Rogers – ***coms0108@waikato.ac.nz***
Mark Apperley – ***mapperle@waikato.ac.nz***

D.3 Research Consent Form

Research Consent Form

THE UNIVERSITY OF WAIKATO
Te Whare Wānanga o Waikato

Ethics Committee, Faculty of Computing and Mathematical Sciences

Extended Workspaces through Virtual Reality – Virtual Reality Environment Study

Consent Form for Participants

I have read the **Participant Information Sheet** for this study and have had the details of the study explained to me. My questions about the study have been answered to my satisfaction, and I understand that I may ask further questions at any time.

I also understand that I am free to withdraw from individual study sessions or the study as a whole at any time, or to decline to answer any particular questions in the study. I understand I can withdraw any information I have provided up until the researcher has commenced analysis on my data. I agree to provide information to the researchers under the conditions of confidentiality set out on the **Participant Information Sheet**.

I agree to participate in this study under the conditions set out in the **Participant Information Sheet**.

Signed: ____

Name: ____

Date: ____

Researcher contact information:

Cameron Grout – ***cgrout@waikato.ac.nz***

Supervisor's contact information:

Bill Rogers – ***coms0108@waikato.ac.nz***
Mark Apperley – ***mapperle@waikato.ac.nz***

D.4 Participant Demographic

Table D.1: Participant demographic summary.

ID	Gender	Age	Vision
0	Female	23	
1	Female	31	
2	Female	25	Spectacles
3	Female	27	Spectacles
4	Female	26	
5	Male	22	Uncorrected
6	Female	24	
7	Female	31	
8	Male	21	
9	Male	22	
10	Male	21	
11	Female	23	
12	Other	20	
13	Male	23	
14	Male	24	Uncorrected
15	Female	22	Uncorrected
16	Male	33	Spectacles
17	Female	22	
18	Male	25	
19	Female	24	
20	Male	24	
21	Female	23	

D.5 Experiment Text Samples

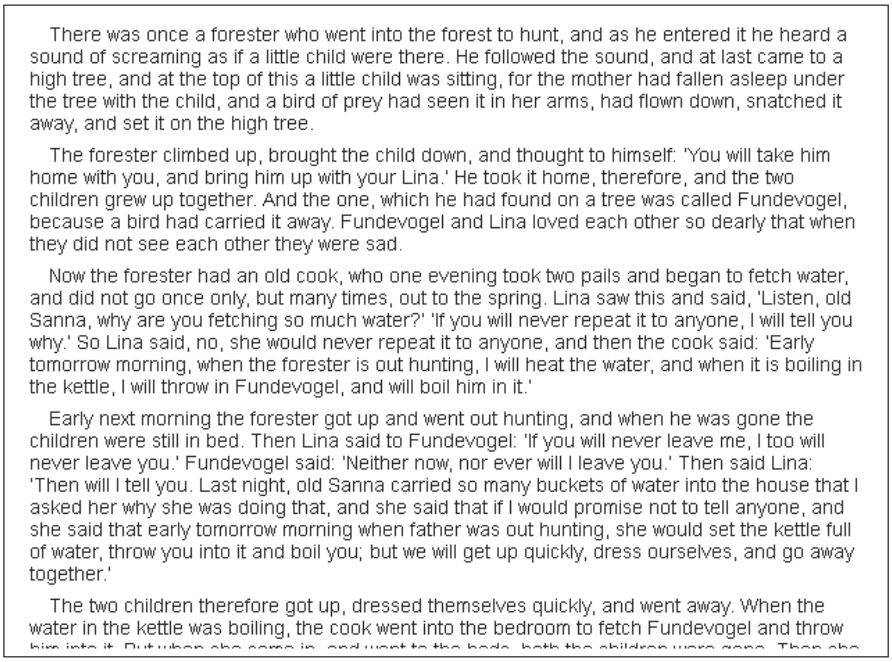

There was once a forester who went into the forest to hunt, and as he entered it he heard a sound of screaming as if a little child were there. He followed the sound, and at last came to a high tree, and at the top of this a little child was sitting, for the mother had fallen asleep under the tree with the child, and a bird of prey had seen it in her arms, had flown down, snatched it away, and set it on the high tree.

The forester climbed up, brought the child down, and thought to himself: 'You will take him home with you, and bring him up with your Lina.' He took it home, therefore, and the two children grew up together. And the one, which he had found on a tree was called Fundevogel, because a bird had carried it away. Fundevogel and Lina loved each other so dearly that when they did not see each other they were sad.

Now the forester had an old cook, who one evening took two pails and began to fetch water, and did not go once only, but many times, out to the spring. Lina saw this and said, 'Listen, old Sanna, why are you fetching so much water?' 'If you will never repeat it to anyone, I will tell you why.' So Lina said, no, she would never repeat it to anyone, and then the cook said: 'Early tomorrow morning, when the forester is out hunting, I will heat the water, and when it is boiling in the kettle, I will throw in Fundevogel, and will boil him in it.'

Early next morning the forester got up and went out hunting, and when he was gone the children were still in bed. Then Lina said to Fundevogel: 'If you will never leave me, I too will never leave you.' Fundevogel said: 'Neither now, nor ever will I leave you.' Then said Lina: 'Then will I tell you. Last night, old Sanna carried so many buckets of water into the house that I asked her why she was doing that, and she said that if I would promise not to tell anyone, and she said that early tomorrow morning when father was out hunting, she would set the kettle full of water, throw you into it and boil you; but we will get up quickly, dress ourselves, and go away together.'

The two children therefore got up, dressed themselves quickly, and went away. When the water in the kettle was boiling, the cook went into the bedroom to fetch Fundevogel and throw

Figure D.1: Example of a text sample screenshot.

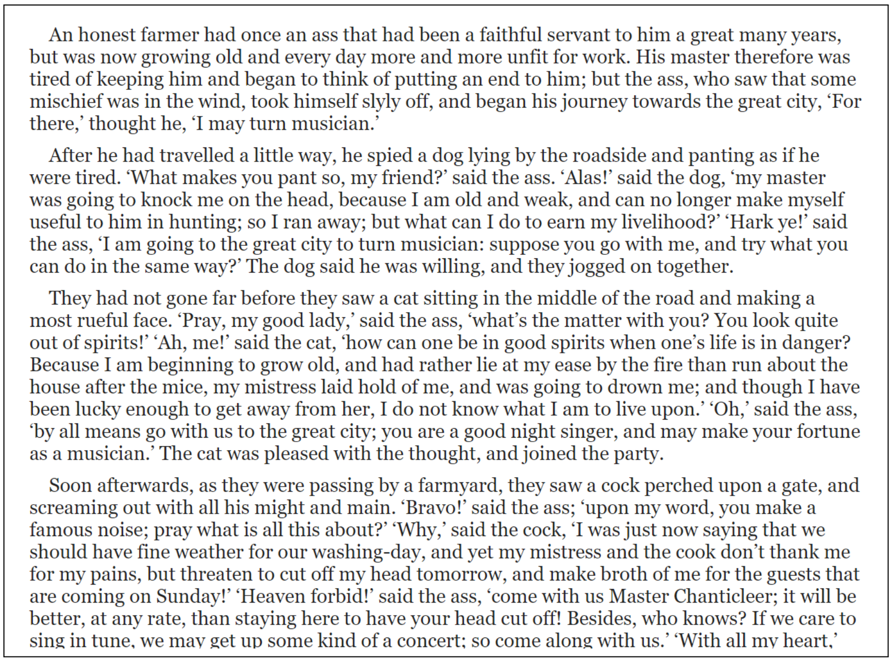

An honest farmer had once an ass that had been a faithful servant to him a great many years, but was now growing old and every day more and more unfit for work. His master therefore was tired of keeping him and began to think of putting an end to him; but the ass, who saw that some mischief was in the wind, took himself slyly off, and began his journey towards the great city, 'For there,' thought he, 'I may turn musician.'

After he had travelled a little way, he spied a dog lying by the roadside and panting as if he were tired. 'What makes you pant so, my friend?' said the ass. 'Alas!' said the dog, 'my master was going to knock me on the head, because I am old and weak, and can no longer make myself useful to him in hunting; so I ran away; but what can I do to earn my livelihood?' 'Hark ye!' said the ass, 'I am going to the great city to turn musician: suppose you go with me, and try what you can do in the same way?' The dog said he was willing, and they jogged on together.

They had not gone far before they saw a cat sitting in the middle of the road and making a most rueful face. 'Pray, my good lady,' said the ass, 'what's the matter with you? You look quite out of spirits!' 'Ah, me!' said the cat, 'how can one be in good spirits when one's life is in danger? Because I am beginning to grow old, and had rather lie at my ease by the fire than run about the house after the mice, my mistress laid hold of me, and was going to drown me; and though I have been lucky enough to get away from her, I do not know what I am to live upon.' 'Oh,' said the ass, 'by all means go with us to the great city; you are a good night singer, and may make your fortune as a musician.' The cat was pleased with the thought, and joined the party.

Soon afterwards, as they were passing by a farmyard, they saw a cock perched upon a gate, and screaming out with all his might and main. 'Bravo!' said the ass; 'upon my word, you make a famous noise; pray what is all this about?' 'Why,' said the cock, 'I was just now saying that we should have fine weather for our washing-day, and yet my mistress and the cook don't thank me for my pains, but threaten to cut off my head tomorrow, and make broth of me for the guests that are coming on Sunday!' 'Heaven forbid!' said the ass, 'come with us Master Chanticleer; it will be better, at any rate, than staying here to have your head cut off! Besides, who knows? If we care to sing in tune, we may get up some kind of a concert; so come along with us.' 'With all my heart,'

Figure D.2: Example of a text sample screenshot.

D.6 Raw Study Data

Table D.2: Raw participant readability ratings for samples displayed with Arial font, scaled to 100%, and ClearType off. Values are arranged as participant ID against sample render distance given in in-engine units (IEUs).

𝔻	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21
2.0			2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
2.1			2	2	2	2	2	2	2	2	2	2	1	2	2	2	1	2	2	2	2	2
2.2			1	2	2	2	2	2	2	2	2	2	2	2	2	1	2	2	2	2	2	2
2.3			1	2	2	2	2	2	1	2	2	2	2	2	2	1	1	2	1	1	2	2
2.4			1	2	2	2	2	2	2	2	2	2	2	2	1	2	0	2	1	1	2	1
2.5			1	2	2	2	1	2	2	2	1	2	1	2	0	1	1	2	1	1	2	1
2.6			1	2	2	2	2	2	2	1	2	2	1	2	1	1	1	2	1	1	2	1
2.7			1	2	1	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	2	1
2.8			1	2	2	2	2	2	2	1	2	2	1	2	0	0	1	2	1	1	1	1
2.9			0	2	2	2	1	1	1	1	1	2	1	2	0	0	1	2	1	1	2	1
3.0	1	2	0	1	1	1	2	2	1	1	1	1	1	1	0	0	0	2	1	1	1	0
3.1	2	2	0	1	1	1	1	1	1	1	1	1	1	1	0	0	0	1	0	0	1	0
3.2	2	1	0	1	1	0	1		1	1	1	1	1	1	0	0	0	2	0	0	1	
3.3	1	2	0	1	1	0	0		1	1	0	1	0	1	0	0	0	1	0	0	1	
3.4	0	1	0	1	1	0	0		0	0	1	1	1	1	0	0	0	1	0	0	1	
3.5	0	0	0	0	1	0	1		0	1	0	0	0	1	0	0	0	1	0	0	1	
3.6	0	0	0	1	1	0	0		0	0	0	0	0	1	0	0	0	1	0	0	1	
3.7	0	0	0	1	1	0	0		0	0	1	0	0	0	0	0	0	1	0	0	0	
3.8	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
3.9	0	0	0	0	1	0	0		0	0	0	0	0	0	0	0	0	1	0	0	0	
4.0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.1	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.2	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.3	0	0	0	0	1	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.4	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.5	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.6	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.7	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.8	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.9	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
5.0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	

Table D.3: Raw participant readability ratings for samples displayed with Arial font, scaled to 100%, and ClearType on. Values are arranged as participant ID against sample render distance given in IEUs.

𝔻	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21
2.0			2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	1	2	2	2
2.1			2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	1	2	2
2.2			2	2	2	2	2	2	2	2	2	2	2	2	2	2	1	2	1	1	2	2
2.3			1	2	2	2	2	2	2	2	2	2	2	2	2	1	1	2	2	1	2	1
2.4			2	2	2	2	2	2	2	2	2	1	2	2	1	1	1	2	1	1	2	2
2.5			1	2	2	2	2	2	2	2	2	2	2	2	1	0	1	2	1	1	2	1
2.6			1	2	2	1	1	2	2	2	2	2	1	2	1	1	1	2	1	1	2	1
2.7			0	2	2	1	2	2	2	1	2	2	1	2	0	0	1	2	1	1	2	1
2.8			1	2	2	2	2	2	2	2	1	1	1	2	0	0	0	2	1	1	2	1
2.9			0	2	2	1	2	1	2	2	1	1	1	2	0	0	0	2	1	1	2	1
3.0	2	2	0	2	2	0	2	1	2	1	1	2	1	2	0	1	0	2	1	0	1	0
3.1	2	2	0	2	2	1	2	1	2	1	1	1	1	2	0	0	0	2	1	0	1	0
3.2	1	2	0	1	2	0	1		1	1	1	1	1	1	0	0	0	1	1	0	1	
3.3	1	1	0	1	2	0	1		1	1	1	1	1	1	0	0	0	1	1	0	1	
3.4	1	1	0	2	1	0	1		1	1	0	0	1	1	0	0	0	1	0	0	1	
3.5	0	1	0	1	1	0	1		0	1	0	0	1	1	0	0	0	1	0	0	1	
3.6	0	1	0	1	1	0	0		1	0	0	1	1	1	0	0	0	1	1	0	1	
3.7	0	1	0	1	0	0	0		0	0	0	0	1	0	0	0	0	1	0	0	1	
3.8	0	1	0	1	1	0	0		1	1	0	0	0	0	0	0	0	0	0	0	1	
3.9	0	0	0	1	1	0	0		0	0	0	0	0	0	0	0	0	1	0	0	1	
4.0	0	1	0	1	1	0	0		0	0	0	0	1	0	0	0	0	0	0	0	0	
4.1	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.2	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.3	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.4	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.5	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.6	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.7	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.8	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.9	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
5.0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	

Table D.4: Raw participant readability ratings for samples displayed with Arial font, scaled to 200%, and ClearType off. Values are arranged as participant ID against sample render distance given in IEUs.

𝔻	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21
2.0			2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
2.1			1	2	2	2	2	2	2	2	2	2	1	2	2	2	2	2	2	2	2	2
2.2			1	2	2	2	2	2	2	2	2	2	1	2	1	2	1	2	2	1	2	2
2.3			1	2	2	2	2	2	2	2	2	2	1	2	1	2	1	2	2	1	1	2
2.4			1	2	2	2	2	2	2	2	2	2	1	2	1	1	1	2	1	1	2	1
2.5			1	2	2	2	2	1	2	1	2	2	1	2	1	0	1	2	2	1	1	1
2.6			0	2	2	2	2	2	2	1	2	2	1	2	0	1	1	2	1	1	2	2
2.7			0	1	2	2	1	1	2	2	2	1	1	2	0	1	0	2	1	1	2	1
2.8			0	2	1	2	2	1	1	1	1	1	1	2	0	0	0	2	1	0	1	1
2.9			0	2	1	1	1	1	1	1	1	1	1	1	0	1	0	2	1	0	1	1
3.0	0	2	0	1	2	1	1	1	1	1	1	1	1	1	0	0	0	1	1	0	1	1
3.1	1	2	0	1	1	1	0	1	1	1	1	0	1	1	0	0	0	1	1	0	1	1
3.2	0	1	0	1	1	1	1		0	1	1	1	0	1	0	0	0	1	0	0	1	
3.3	1	1	0	1	1	0	0		0	1	0	0	0	0	0	0	0	1	0	0	1	
3.4	0	1	0	1	1	0	1		0	0	0	0	0	0	0	0	0	0	1	0	1	
3.5	0	0	0	0	1	0	0		0	1	0	0	0	0	0	0	0	0	0	0	0	
3.6	0	1	0	0	1	0	0		0	1	0	0	0	0	0	0	0	0	0	0	0	
3.7	0	0	0	0	1	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
3.8	0	0	0	0	1	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
3.9	0	0	0	0	1	0	0		0	1	0	0	0	0	0	0	0	0	0	0	0	
4.0	0	1	0	0	1	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.1	0	0	0	0	1	0	0		0	1	0	0	0	0	0	0	0	0	0	0	0	
4.2	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.3	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.4	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.5	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.6	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.7	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.8	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.9	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
5.0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	

Table D.5: Raw participant readability ratings for samples displayed with Arial font, scaled to 200%, and ClearType on. Values are arranged as participant ID against sample render distance given in IEUs.

𝔻	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21
2.0			2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
2.1			2	2	2	2	2	2	2	2	2	2	2	2	2	2	1	2	1	2	2	2
2.2			2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
2.3			2	2	2	2	2	2	2	2	2	2	2	2	2	1	1	2	1	2	2	2
2.4			1	2	2	2	2	2	2	2	2	2	2	2	1	1	1	2	1	1	2	1
2.5			1	2	2	2	2	2	2	2	2	2	1	2	1	1	0	2	1	1	1	1
2.6			1	2	2	2	2	1	2	2	2	2	1	2	0	0	1	2	1	1	2	2
2.7			1	2	2	2	1	2	2	2	2	1	1	2	0	0	1	2	1	1	1	1
2.8			1	1	2	2	2	1	2	2	1	2	1	2	0	0	0	2	1	1	2	1
2.9			0	2	2	2	1	1	1	2	1	1	1	2	0	0	0	2	0	0	1	1
3.0	1	2	1	2	2	1	1	1	1	1	1	1	1	2	0	0	0	1	0	1	1	0
3.1	1	1	0	1	2	1	1	1	1	1	1	1	1	1	0	0	0	1	0	0	1	0
3.2	1	1	0	1	1	1	1		1	1	0	1	1	1	0	0	0	1	0	0	1	
3.3	1	1	0	1	1	0	0		0	1	1	0	0	1	0	0	0	1	1	0	1	
3.4	1	1	0	1	1	0	1		2	1	0	0	1	1	0	0	0	1	0	0	0	
3.5	1	1	0	0	0	0	0		0	1	1	0	0	1	0	0	0	1	0	0	1	
3.6	0	1	0	0	1	0	0		0	1	0	0	0	0	0	0	0	1	0	0	0	
3.7	0	0	0	0	0	0	0		0	1	0	0	0	0	0	0	0	0	0	0	0	
3.8	0	0	0	0	1	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
3.9	0	0	0	0	1	0	0		0	1	0	0	0	0	0	0	0	1	0	0	0	
4.0	0	0	0	0	1	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.1	0	0	0	1	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.2	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.3	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.4	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.5	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.6	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.7	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.8	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.9	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
5.0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	

Table D.6: Raw participant readability ratings for samples displayed with Georgia font, scaled to 100%, and ClearType off. Values are arranged as participant ID against sample render distance given in IEUs.

𝔻	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21
2.0			2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
2.1			2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
2.2			2	2	2	2	2	2	2	2	2	2	2	2	1	2	2	2	2	2	2	1
2.3			1	2	2	2	2	2	2	2	2	2	2	2	1	1	1	2	1	2	2	1
2.4			2	2	2	2	2	2	2	2	2	2	1	2	1	1	1	2	1	1	2	1
2.5			2	2	2	2	2	1	2	2	2	2	2	2	0	2	1	2	1	1	2	1
2.6			1	2	2	2	2	2	2	2	1	2	2	2	0	0	1	2	0	1	2	1
2.7			1	2	2	1	1	2	2	2	1	1	1	2	0	1	1	2	1	1	1	0
2.8			0	2	2	1	1	1	2	1	2	1	1	2	0	1	1	2	1	1	2	0
2.9			0	2	2	1	1	1	1	1	1	1	1	2	0	0	1	1	1	0	2	1
3.0	1	2	0	1	2	1	1	1	1	1	1	1	1	2	0	1	0	1	1	0	1	0
3.1	2	2	0	1	1	1	1	1	1	1	1	1	1	1	0	0	0	1	0	1	1	0
3.2	1	1	0	1	1	1	1		1	1	1	1	1	1	0	0	0	1	0	1	1	
3.3	1	1	0	1	1	0	1		1	1	1	0	1	0	0	0	0	1	0	0	1	
3.4	1	1	0	1	1	1	0		1	1	1	0	1	1	0	0	0	1	0	0	0	
3.5	0	1	0	1	0	1	1		1	1	0	0	0	0	0	0	0	1	0	0	0	
3.6	1	1	0	0	1	1	0		0	1	0	1	0	0	0	0	0	1	0	0	0	
3.7	0	1	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
3.8	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
3.9	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.0	0	0	0	0	0	0	0		0	0	0	0	1	0	0	0	0	0	0	0	0	
4.1	0	0	0	0	1	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.2	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.3	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.4	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.5	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.6	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.7	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.8	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.9	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
5.0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	

Table D.7: Raw participant readability ratings for samples displayed with Georgia font, scaled to 100%, and ClearType on. Values are arranged as participant ID against sample render distance given in IEUs.

𝔻	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21
2.0			2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
2.1			2	2	2	2	2	2	2	2	2	2	2	2	1	2	2	2	2	2	2	2
2.2			1	2	2	2	2	2	2	2	2	2	2	2	1	1	2	2	2	1	2	2
2.3			2	2	2	2	2	2	2	2	2	2	2	2	1	1	2	2	1	2	2	2
2.4			2	2	2	2	2	2	2	2	2	2	2	2	1	1	1	2	1	1	2	2
2.5			1	2	2	2	2	2	1	2	2	2	2	2	1	0	1	2	1	1	2	1
2.6			1	2	2	2	2	2	2	1	2	1	2	2	0	1	2	2	1	1	2	1
2.7			0	2	2	2	2	2	2	2	2	2	1	2	0	0	1	2	1	0	2	1
2.8			1	2	2	2	2	1	1	2	2	2	2	2	0	0	0	1	1	0	2	1
2.9			0	2	2	1	1	1	1	2	1	1	1	1	0	0	0	2	0	0	1	0
3.0	1	2	0	2	2	1	1	1	1	2	1	0	1	2	0	0	0	1	1	0	1	1
3.1	2	2	1	2	2	1	2	1	0	1	1	1	2	1	0	0	0	1	1	0	2	0
3.2	2	2	0	2	2	1	1		1	1	1	0	1	1	0	0	0	1	0	0	1	
3.3	0	1	0	2	1	1	1		1	1	1	1	1	1	0	0	0	1	0	0	1	
3.4	0	1	0	1	2	0	1		0	1	1	1	1	1	0	0	0	1	0	0	1	
3.5	1	1	0	1	2	0	0		0	1	1	0	1	1	0	0	0	1	0	0	1	
3.6	0	0	0	1	1	0	0		0	1	0	0	1	0	0	0	0	1	0	0	1	
3.7	0	1	0	1	1	0	1		0	0	0	0	0	0	0	0	0	1	0	0	0	
3.8	0	0	0	2	1	0	0		0	0	0	0	1	0	0	0	0	0	0	0	0	
3.9	0	0	0	1	1	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.0	0	0	0	0	1	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.1	0	0	0	1	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.2	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.3	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.4	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.5	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.6	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.7	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.8	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.9	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
5.0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	

Table D.8: Raw participant readability ratings for samples displayed with Georgia font, scaled to 200%, and ClearType off. Values are arranged as participant ID against sample render distance given in IEUs.

𝔻	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**	**16**	**17**	**18**	**19**	**20**	**21**
2.0			1	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	1	1	2	2
2.1			1	2	2	2	2	2	2	2	2	2	2	2	2	2	1	2	2	1	2	2
2.2			1	2	2	2	2	2	2	2	2	2	1	2	1	2	2	2	1	1	2	2
2.3			2	2	2	2	2	2	2	2	2	2	2	2	1	2	1	2	2	1	2	1
2.4			1	2	2	2	2	2	2	2	1	2	2	2	1	1	2	2	1	1	2	1
2.5			1	2	2	2	1	2	2	2	2	2	2	2	1	0	2	2	1	1	1	1
2.6			1	2	2	2	2	2	1	1	1	2	1	2	1	2	1	2	1	1	2	1
2.7			0	2	2	2	2	2	2	2	1	2	1	2	2	1	1	2	0	1	1	0
2.8			0	2	2	1	1	1	1	1	1	1	1	2	0	1	0	1	0	1	1	0
2.9			0	2	2	1	2	1	2	1	1	1	1	1	0	1	0	2	1	1	1	0
3.0	1	2	0	1	2	1	1	1	1	1	1	1	1	1	0	0	0	1	0	0	1	0
3.1	1	1	0	1	2	1	1	1	1	1	0	1	1	1	0	0	0	1	0	0	1	0
3.2	1	1	0	1	1	1	1		1	1	1	1	1	1	0	0	0	1	0	0	1	
3.3	1	1	0	1	1	0	1		1	1	0	1	1	1	0	0	0	1	0	0	1	
3.4	1	1	0	1	1	0	1		0	0	0	0	1	1	0	0	0	1	1	0	1	
3.5	1	0	0	1	1	0	0		0	1	0	0	1	0	0	0	0	1	0	0	1	
3.6	0	1	0	1	1	0	0		0	1	0	0	0	0	0	0	0	0	0	0	0	
3.7	0	0	0	1	1	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
3.8	0	0	0	0	1	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
3.9	0	0	0	0	1	0	0		0	0	0	0	1	0	0	0	0	0	0	0	0	
4.0	0	0	0	0	1	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.1	0	0	0	0	1	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.2	0	0	0	0	1	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.3	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.4	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.5	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.6	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.7	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.8	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.9	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
5.0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	

Table D.9: Raw participant readability ratings for samples displayed with Georgia font, scaled to 200%, and ClearType on. Values are arranged as participant ID against sample render distance given in IEUs.

𝔻	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**	**16**	**17**	**18**	**19**	**20**	**21**
2.0			2	2	2	2	2	2	2	2	2	2	2	2	2	2	1	2	1	1	2	2
2.1			2	2	2	2	2	2	2	2	2	2	2	2	1	2	1	2	2	1	2	2
2.2			1	2	2	2	2	2	2	2	2	2	2	2	1	2	1	2	1	1	2	2
2.3			1	2	2	2	2	2	2	2	2	2	1	2	1	2	1	2	1	1	2	1
2.4			1	2	2	2	2	2	2	2	2	2	2	2	0	1	1	2	0	1	2	1
2.5			1	2	2	2	2	2	2	2	2	2	1	2	1	1	1	2	1	1	2	1
2.6			1	2	2	2	1	2	2	2	2	1	1	2	0	1	0	2	1	1	2	1
2.7			1	2	2	2	1	2	2	2	1	2	1	2	0	1	0	1	1	1	1	1
2.8			0	2	2	1	1	2	2	1	1	2	1	2	0	0	0	1	0	1	1	1
2.9			0	2	2	1	1	1	1	1	1	1	1	1	0	1	0	1	1	1	2	0
3.0	2	1	0	1	2	1	1	1	2	1	1	1	1	1	0	1	1	1	0	0	1	0
3.1	1	1	0	2	2	1	1	1	2	1	1	1	1	1	0	0	0	1	0	0	1	0
3.2	1	1	0	2	1	1	1		1	1	0	0	1	1	0	0	0	1	0	0	1	
3.3	1	0	0	1	1	1	1		1	1	1	0	1	1	0	0	0	1	0	0	1	
3.4	0	1	0	1	1	0	1		1	0	0	1	0	1	0	0	1	1	1	0	1	
3.5	1	0	0	1	1	0	0		0	0	0	0	0	1	0	0	0	1	0	0	1	
3.6	0	0	0	1	1	0	0		1	0	0	0	0	0	0	0	0	0	0	0	1	
3.7	0	1	0	0	1	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
3.8	0	0	0	0	1	0	0		0	0	0	0	0	1	0	0	0	0	0	0	0	
3.9	0	0	0	1	1	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.0	0	0	0	1	1	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.1	0	0	0	0	1	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.2	0	0	0	0	1	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.3	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.4	0	0	0	0	1	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.5	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.6	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.7	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.8	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
4.9	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	
5.0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	

D.7 Processed Distance Values

Table D.10: $\mathbb{D}_n$ values for participants, given in IEUs.

ID	Arial 100% On	Arial 100% Off	Arial 200% On	Arial 200% Off	Georgia 100% On	Georgia 100% Off	Georgia 200% On	Georgia 200% Off
0	3.2	3.1			3.4	3.5	3.0	3.6
1	3.3	3.2	3.1	3.0	3.5	3.9	3.5	3.7
2	2.1	2.4	2.0	2.3	2.9	2.7	2.6	2.9
3	2.9	3.4	2.9	3.0	3.5	4.1	3.5	3.5
4	2.9	3.3	3.0	3.1	3.8	3.7	4.2	3.5
5	2.9	2.8	2.8	2.9	3.2	3.0	3.3	3.3
6	3.0	3.1	2.8	2.8	3.3	3.6	3.1	3.3
7	3.0	2.8	2.6	2.7				
8	2.8	3.1	2.7	3.4	3.4	3.5	3.2	3.3
9	2.7	2.9	2.7	2.9	3.4	3.6	3.4	3.8
10	2.8	2.7	2.7	2.7	3.3	3.4	3.3	3.2
11	2.9	3.0	2.6	2.8	3.5	3.4	3.1	3.3
12	2.4	2.5	2.0	2.4	3.3	3.8	3.2	3.3
13	2.9	3.1	2.8	3.0	3.7	3.7	3.3	3.6
14	2.3	2.3	2.1	2.3	2.5	2.7	2.6	2.6
15	2.4	2.2	2.3	2.2	2.8	2.5	2.5	2.6
16	2.2	2.1	2.1	2.2	2.4	2.8	2.7	2.5
17	3.2	3.1	2.9	2.9	3.8	3.8	3.4	3.7
18	2.2	2.3	2.5	2.2	3.1	3.4	3.2	2.9
19	2.2	2.0	2.1	2.3	3.1	3.0	2.8	2.9
20	2.9	2.9	2.7	2.8	3.7	4.0	3.5	3.4
21	2.3	2.4	2.6	2.6	3.0	3.0		3.0

Table D.11: $\mathbb{D}_f$ values for participants, given in IEUs.

ID	**Arial 100% On**	**Arial 100% Off**	**Arial 200% On**	**Arial 200% Off**	**Georgia 100% On**	**Georgia 100% Off**	**Georgia 200% On**	**Georgia 200% Off**
0	3.1	3.2		3.0	3.5	3.3	3.6	3.4
1	3.1	3.2	3.0		3.8	3.6	3.5	3.3
2	2.5	2.4	2.3	2.1	2.8	2.7	2.7	2.8
3	2.9	3.8	2.9	3.2	3.6	4.0	3.8	3.7
4	3.0	3.5	3.1	3.1	3.5	4.1	4.3	4.3
5	2.6	2.8	2.7	2.7	3.3	3.4	3.3	3.4
6	2.6	3.1	2.9	2.5	3.4	3.5	3.5	3.5
7	2.7	2.7	2.7	2.8				
8	2.8	2.7	2.9	3.1	3.6	3.1	3.4	3.5
9	2.7	3.0	2.7	2.7	3.7	3.7	3.4	3.4
10	2.8	2.8	2.5	2.6	3.5	3.6	3.1	3.2
11	2.6	2.8	2.7	2.8	3.3	3.0	3.4	3.2
12	2.6	3.1	2.5	2.4	3.5	3.7	3.6	3.4
13	3.0	3.0	2.8	2.8	3.3	3.6	3.5	3.6
14	2.1	2.0	2.7	2.0	2.5	2.6	2.8	2.4
15	2.5	2.1	2.6	2.3	2.6	2.5	2.5	2.8
16	2.2	2.6	2.5		3.0	2.8	2.8	2.6
17	2.8	2.9	2.9	2.6	3.7	3.8	3.6	3.6
18	2.2	2.2	2.3	2.1	2.6	2.9	2.7	2.4
19	2.3	2.3			2.9	2.7	3.0	3.0
20	2.9	3.1	2.6	2.9	3.4	3.7	3.6	3.7
21	2.1	2.4	2.2	2.2	2.7	2.9	2.7	2.9

I am certain of one thing and it's this: the worst thing you can do is nothing.

Terry Pratchett

— *Snuff*