3D Display Technologies for Air Traffic Control
A Pilot Study

QMW Computer Science Technical Report No.690
12 December 1994

Mark A. Brown
Advanced Computing Environments (ACE) Laboratory
Department of Computer Science
Queen Mary & Westfield College
University of London
Mile End Road
London E1 4NS

Tel: 0171 975 5200 (domestic)
+44 171 975 5200 (overseas)
Email: `eeyore@dcs.qmw.ac.uk`

Abstract

This report details the pilot experiment for the author's investigation into three-dimensional display technologies for real-time spatial command/control applications, concentrating on air traffic control. Nine ex-air traffic controllers participated in the study, carrying out three tasks: a parameter reading task to determine the effects of parallel or perspective projection on the reading of relative azimuth and horizontal distance in a pseudo-3D display, a recall task to test the effects of display type on memory of a static scenario, and a conflict detection task to test the effect of display type on controller awareness. The purpose was to validate experimental techniques and display prototype designs to be used in the main study. The results and conclusions, and some of the implications for the main study, are detailed here.

Keywords: air traffic control, three-dimensional displays, virtual reality applications.

Contents

1 Introduction

In recent years, large increases in computing power and the falling cost of that power have enabled sophisticated three-dimensional graphics to be produced on relatively cheap workstations instead of the large, expensive machines used hitherto. With the resulting increased availability of 3D graphics and the associated technology of virtual reality (VR), there has been a surge of interest in using three-dimensional displays in applications where the cost, reliability or size of computers of the required power were previously prohibitive.

In order to be adopted, however, a new display technology must demonstrate significant benefits over that which it seeks to replace. The advantage of a 3D display is that for 3D spatial visualisations, it gives an integrated, holistic presentation of all three spatial dimensions, offering potentially reduced cognitive workload required to form a mental model of the spatial relationships, and a correspondingly higher "situational awareness". Immersive displays further offer a sense of "presence" which can influence how the viewer interprets or interacts with the data. However, compared to current two-dimensional displays, 3D displays tend to be more expensive (in terms of software, hardware and computing power required to produce them) and introduce additional factors in their design and implementation to those already applicable to 2D displays. Human cognitive factors are amongst these: new technologies tend to introduce new types of errors, and human performance has to be taken into account in order to match the display to the application effectively. Poor design may lead to a new display which gives only equal or even inferior performance to the original, negating any potential benefits.

One possible application of 3D displays, whether they are immersive VR or conventional "through-the-window" (TTW) displays, is monitoring, command and/or control of trajectories of objects moving in three dimensions in safety-critical/real-time applications. These could be real objects, such as aeroplanes, submarines and ships, or abstract entities, such as the trajectories of process control parameters in state-space. The safety-critical and real-time aspects are important since they place constraints on the presentation of the information *vis-à-vis* speed of assimilation of a situation and quality of perceived data. In an ordinary scientific visualisation application, data are usually retrieved from mass storage or are computer-generated on demand and the scientist can examine the dataset at leisure, exploring it from a variety of points. In a real-time application where critical decisions may have to be reached in a short period of time, this is a luxury which cannot be afforded; the required information must be easy and quick to read and difficult to misinterpret.

This research project aims to explore the application of three-dimensional displays to this area, using Air Traffic Control (ATC) as a specific example. 3D displays have been investigated in relation to ATC over a number of years, but so far none of these has achieved operational or even training use. This could be partly due to the hitherto large cost and technical complexity involved in 3D displays, but it could also be due to other factors, possibly human cognitive and other psychological characteristics and limitations. It is therefore an application which merits specific investigation and makes an ideal vehicle for this research into 3D displays for more generic command/control tasks.

This research evaluates a number of different display types for three-dimensional spatial command/control tasks, to help to assess the suitability of these displays for such applications, studying air traffic control as a specific example.

1. 2D plan-position indicator (PPI) (the reference).
2. 3D TTW pseudo-3D displays (monocular depth cues only).
3. 3D TTW stereoscopic 3D display (non head-tracked, using LCD shutter glasses to implement stereopsis).
4. An immersive VR display.

Evaluation will consists of both subjective assessment and quantitative assessment of part-task tests:

1. An azimuth angle and relative distance reading task.
2. A memory recall task.
3. A conflict detection task.

The displays will be evaluated using full-time current Air Traffic Control Officers (ATCOs), who have a training bias towards the existing plan-view presentation, and university students, who have no such bias but who are untrained in air traffic control.

It is hoped that the results will yield a better understanding of the strengths and limitations of 3D displays for this type of application.

This report describes the pilot study for this research programme; its methods, results and conclusions. The purpose of the pilot study was to validate techniques for a full experimental study under design and to gain subjective opinion data on the display formats developed for the experiments. This is to be followed by the main study which will be based on its results.

Chapter 2 introduces air traffic control in the context of this research.

Chapter 3 describes the objectives and methodology for each of the experiments in the study.

Chapter 4 describes the display formats developed for the pilot study.

Chapter 5 then goes on to describe the design of the experiments, and the methods. Chapter 6 then discusses results.

Finally, Chapter 7 makes conclusions and recommendations from the pilot experiment.

2 Overview of Air Traffic Control

2.1 Introduction

This chapter looks at the general task of radar control of aircraft as it relates to this research, viewing ATC as a flow control problem based on studies into Computer Supported Cooperative Working (CSCW) in ATC by Harper *et alia* [HHS91, HH93].

There are many different facets of the air traffic controller's task, and different tasks in different areas. For example, approach control is a different task from en-route or ground control. A full description of the air traffic control system would be quite beyond the scope of this document; the interested reader is referred to Graves [Gra89] and Duke [Duk92] for excellent introductions to air traffic control in the United Kingdom, and to Nolan [Nol90] for more technical information about how air traffic control in the United States.

2.2 ATC as Flow Control

The prime task in air traffic control, according to the Manual of Air Traffic Services, is:

> to maintain the safe, orderly and expeditious flow of traffic (1)

or, as the title of [HH93] quotes one ATCO:

> Send 'em all to the same place and expect us to stop them hitting. (2)

Air traffic control is basically an exercise in flow control. Each ATCO is responsible for a particular *sector*, which is a three-dimensional slice of airspace. Aircraft enter the sector at various positions and times, and must be guided to their exit points (e.g. an aerodrome or a point at the boundary of an adjacent sector) whilst observing (1) above: *safe* means that aircraft must be kept from colliding with each other or the terrain (except in the controlled case of landing!), *orderly* implies that the flow should be organised with aircraft following airways and traffic patterns rather than *ad hoc* (with the side-effect that it restricts aircraft to narrow air corridors instead of making full use of the available space, hence (2) above), and *expeditious* means that aircraft should be guided to their destinations as quickly and as efficiently as possible, within the constraints of safety and orderliness.

The ATCO has a number of tools in his/her job. The most important of these are:

- R/T channel to communicate with aircraft.
- Land lines to other controllers.
- Flight Progress Strips (FPSs).
- PPI radar display (real-time display of aircraft position).

The ways in which the FPS and radar are used is described below.

2.2.1 Flight Progress Strips

Perhaps surprisingly to the uninitiated, the flight progress strip is more important than the radar. ATC may be done using flight strips and procedural separation rules with no real-time air traffic display (as indeed was done in the past[1]). The FPS comprises a strip of card corresponding to an aircraft on which are printed pertinent data: aircraft callsign, type, route code, speed, time expected at waypoint, current altitude etc. The ATCO has these arranged in racks in front of him/her near to the radar screen. The strips may be organised in the racks by waypoint, and under each waypoint by time or some other system of the ATCO's choosing. The racks allow for FPSs to be "cocked out" for whatever reason—usually as an *aide memoire* to draw attention to a strip if an aircraft requires special monitoring or other action, for example.

The FPSs are annotated by the controller as instructions are issued to the aircraft and so reflect the state of the aircraft at any given time. (For example, when instructing an aircraft to descend, the ATCO will cross out the current level and write a descending arrow on the strip with the cleared level. This will be further amended when the aircraft levels off.)

2.2.2 PPI Radar Display

In addition to the FPSs, another major tool in ATC is the plan-position indicator radar display, i.e. one showing azimuth and distance of targets from the radar. At its most basic, this comprises a circular monochrome CRT with a long-persistence phosphor. As the radar transmitter is scanned 360° in azimuth, it sends out a radio pulses in a beam which are reflected from objects (e.g. terrain, clouds and aircraft) back to a receiver[2]. The time delay between a pulse being transmitted and an echo from a target being

[1]The US system had "shrimp boats" on a board which were manually updated by operators as aircraft reported passing waypoints, but there was no way of verifying whether the reported position was correct or obtaining a display of current position except at the reporting points [Nol90].

[2]The same antenna is usually used for both transmission and receiving: the radar sends out a pulse and then listens for a while through the same antenna for echos before sending the next pulse.

received gives the target's distance from the radar, and its bearing from the radar is simply the azimuth of the antenna at the time. Radar returns are drawn on the CRT at the appropriate azimuth and at distances from the centre of the display depending on the time delay; thus a plan view picture of radar targets around the radar head is built up. Because a radar sweep can take several seconds, the long persistence phosphor is used to retain the image over several sweeps, and this allows moving targets to be seen against a static background (since their positions change between successive sweeps); they are shown as a bright primary radar return (the last fix) with a trail of previous returns of diminishing intensity behind them. The distance between the trailing returns enables the speed of the target to be estimated, and their direction enables the heading and any horizontal manœuvering of the target to be seen.

A raw radar display is often very cluttered so Moving Target Indicator (MTI) processing is used to remove static targets. This involves measuring the Doppler frequency shift of the echoes, which gives the component of velocity of the targets directly away from or towards the radar; static targets can thus be filtered out[3].

As an aid to orientation, a *video map* showing airway structures, waypoints, coastlines etc. may be superimposed on the display. Concentric *range rings* may also be superimposed to enable range from the radar head to be estimated.

Instead of displaying MTI processed raw returns, the position data from the radar data processing (RDP) computer can be used to show aircraft position on a *synthetic display*, using purely computer-generated symbology. Such information may be shown on a raster display with a short-persistence phosphor, and multiple colours may then be used.

The primary radar gives azimuth and range of an aircraft from the radar, but not its identification or height. This information is obtained using a *secondary surveillance radar* (SSR), usually located alongside the primary radar, and *transponders* on each aircraft. The SSR sends out an interrogation pulse which is picked up by an aircraft's transponder. This sends back a *mode-A* or *squawk* code (a four-digit octal number which is set by the pilot) and, if the transponder is an altitude-encoding unit, a *mode-C* code which gives barometric height in hundreds of feet above the reference pressure level set by the pilot. The transponder information may then be displayed on the radar PPI next to the corresponding 'blip' as an alphanumeric *datablock* or *tag*. Commercial aircraft generally have flight plans filed in the ATC flight

[3]However, if the target is moving tangential to a circle centered at the radar head, the velocity component along a radial from the radar will be zero and so MTI will remove the target as long as it continues along a tangent. Since this would involve travelling in a circle around the radar, in practice no aircraft return is filtered out by MTI for more than a few successive sweeps.

data processing (FDP) computer. In this case, the mode-A code will have been allocated to the aircraft by ATC and will be associated with the flight's file in the FDP system. This enables the flight number to be displayed on the radar instead of the squawk code, and allows the FDP computer to track the progress of the flight and to issue flight progress strips to controllers shortly before the aircraft arrives at their sectors.

2.2.3 Planning: The Picture

ATCOs solve the ATC flow problem by planning; looking at the current and future state of the traffic under their control, and organising a strategy to deal with it. The FPSs and radar are used as tools in the formation and execution of this plan.

Strategic (long-term) planning involves organising the overall flow and identifying any problem traffic. There is usually a characteristic pattern of traffic associated with each sector (e.g. outbounds from terminal area may be routed east and north, inbounds may be routed from the west and south) and usually an aircraft can be slotted into it. Occasionally, there are problem aircraft which do not conform to this pattern (e.g. an aircraft turning between airways or an overflying aircraft passing through an outbound traffic stream) and the ATCO must plan accordingly. Strategic planning tends to be done with the flight strips. Receiving a strip shortly before an aircraft is due to enter his/her sector, the ATCO may note route code, altitude and time expected at a point. Checking against other flight strips in the pertinent racks usually reveals whether or not there is likely to be a problem — e.g. if two aircraft on the same route are due over a point at approximately the same time, or if the aircraft is at a level already occupied by another aircraft, or whether the aircraft's intended route will cross other traffic. Noticing a potential problem the ATCO can plan to deal with it, and may cock potentially conflicting strips out of the rack as a reminder to monitor those aircraft more closely.

The radar tends to be used to monitor the execution of the plan and the current state of the traffic, and for tactical (short-term) planning, e.g. adjusting headings or levels to avoid other traffic. For example, if an aircraft must descend through the level of an aircraft crossing its path, the ATCO may decide initially to descend it to just above the level of the crossing traffic, and monitor the situation on the radar. As the aircraft get closer together, the radar picture will help the ATCO to decide whether it is safe to continue the aircraft down before the crossing or to delay its descent until clear of the other traffic.

At times of heavy workload, the emphasis of usage tends to be on the radar; at other times, the emphasis appears to be more on the use of the FPS.

The ATCO's mental model of the plan, the state of its execution and the disposition of traffic is known in the ATC vernacular as the *picture*, and this is a very important part of the task. Losing the picture is one of the air traffic controller's worst fears.

Radar is also used for overall monitoring. Humans do not behave perfectly and there may be a lapse on the part of a controller or pilot which brings aircraft into direct conflict that will not necessarily be spotted on the flight strips. An "altitude bust" is a good example of this — an aircraft converging head-on with another aircraft may be ordered to level off just below it, but due to pilot inattention may "bust" through its altitude, creating a conflict. For this, the only means of controller detection (apart from reports of an airmiss or an aluminium shower) is the use of the radar (and, nowadays, automated aids, although these are by no means perfect).

3 The Pilot Study

3.1 Introduction

The aim of this research is to evaluate four types of display for the purposes of 3D spatial command/control tasks, with application to ATC:

1. 2D PPI (the reference).
2. Pseudo-3D TTW display.
3. Stereoscopic 3D TTW display (non-head-tracked).
4. 3D immersive VR display.

Previous reports by the author [Bro94b, Bro94a, Bro94c] discussed background theory, hypotheses to evaluate these displays, and initial experimental designs. The purpose of the pilot study was to validate the experiment in preparation for the main study. The experiment comprised three sub-experiments or tasks:

1. Spatial parameter reading task, to determine whether or not perspective or parallel 3D projections are more suitable for ATC applications.
2. Recall task, to investigate the effects of display type on speed of assimilation of a scenario, and accuracy of recall.
3. Conflict detection task, to determine the effect of display type on controller awareness.

In addition, subjective feedback was obtained on the design of the display formats and the tasks.

This chapter describes briefly the background theory behind each task, and its objectives.

3.2 Parameter Reading Experiment

3.2.1 Introduction

The purpose of this experiment was to determine the influence of projection type on the reading of horizontal separation and azimuth angle between two objects in a three-dimensional display.

During preliminary discussions with ATCOs relating to the 3D presentation of air traffic information, the ability to judge horizontal separation and azimuth angle between aircraft clearly was cited as being of prime importance, since separation minima must be strictly maintained. The ambiguity of presentation of horizontal and vertical information is dependent on the characteristics of the display.

In a 3D display, there is an inherent distortion of distances and angles, the degree of distortion depending on the angle between dimension to be measured and the line of sight into the display, and display parameters such as geometric field of view (GFOV), virtual camera elevation etc. Figures 3.1(a)–(c) illustrate this by showing how three vectors are projected in 3D. In these figures, COP is the Centre of Projection (the position of the virtual camera), VPN is the View Plane Normal (line of sight vector) and VP is the View Plane onto which the objects in the scene are projected. Distances perpendicular to the line of sight are represented unambiguously (Fig. 3.1(a)), whereas distances along the line of sight are lost (Fig. 3.1(b)). If the 2D PPI is considered as being a 3D display viewed from directly above at infinity, distances in the plane of the ground are represented unambiguously, so horizontal separations and bearings are undistorted, whilst height information is lost altogether, and needs to be represented by some other means (in the case of an ATC display, by a digital readout). In a 3D display with an arbitrary line of sight, there will generally be some distortion to both horizontal and vertical information (Fig. 3.1(c)).

There are two main projections used to show a three-dimensional scene on a 2D surface: perspective (or polar) and parallel. Perspective texture gradients and linear perspective exist in perspective projection displays: there is a diminishing of displayed object size with distance along the VPN (i.e. the z-depth) with the perspective projection. In the parallel projection, linear perspective is absent: the displayed object size is invariant with z-depth (see Figure 3.2).

Linear perspective is an import depth cue in a three-dimensional display, but may interfere with the ability to judge distances and angles. This raises the question of whether or not perspective or parallel projections are preferable for ATC.

Previous studies have investigated the relationship between perception of azimuth and elevation angles and distances between objects in a 3D display, and the display viewing parameters in perspective projections (e.g. Ellis *et alia* [ETGS91], Barfield *et alia* [BLR90]). However, there is not much literature about such studies in a parallel projection display (but a parallel projection is just a special case of perspective projection, with the COP at infinity). Some previous investigation also had subjects making judgments of angle between objects with a regular grid background; in the displays here, circular range rings are used instead, so removing a convenient local reference to the grid axes.

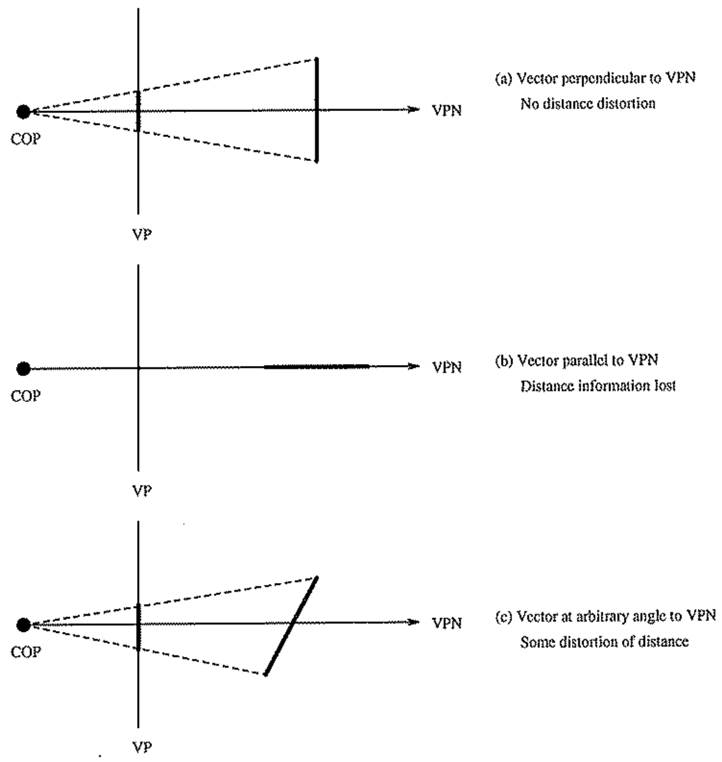

Figure 3.1: Distortions of Dimensions in a 3D Display

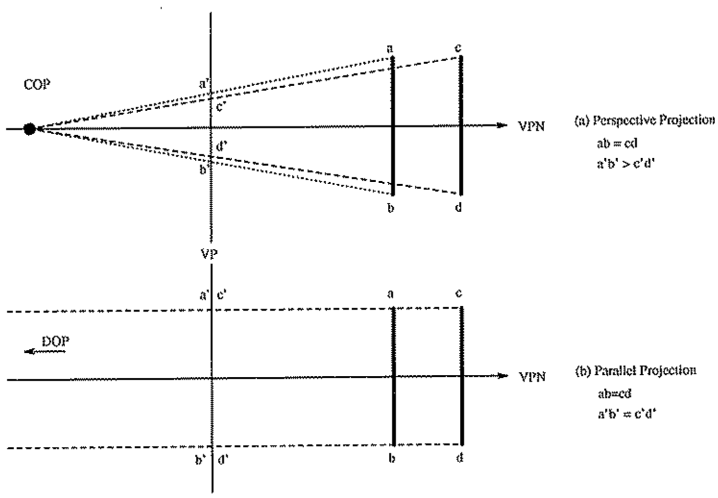

Figure 3.2: Parallel and Perspective Projection

The purpose of this experiment was therefore to make a quantitative comparison of readings of horizontal distance and azimuth between two aircraft between parallel and perspective 3D projections. As alluded to above, there are a number of variables which affect the reading of these parameters; however, this study is only concerned with the *overall* effects of projection type on the reading of parameters, and not on the interactions between the nuisance variables. This is taken into account by choosing random values for some of these nuisance variables (including the off-axis angle and the z-depth (see Figure 3.2.1) and fixing others (such as the relative elevation angle, by putting the targets at the same height).

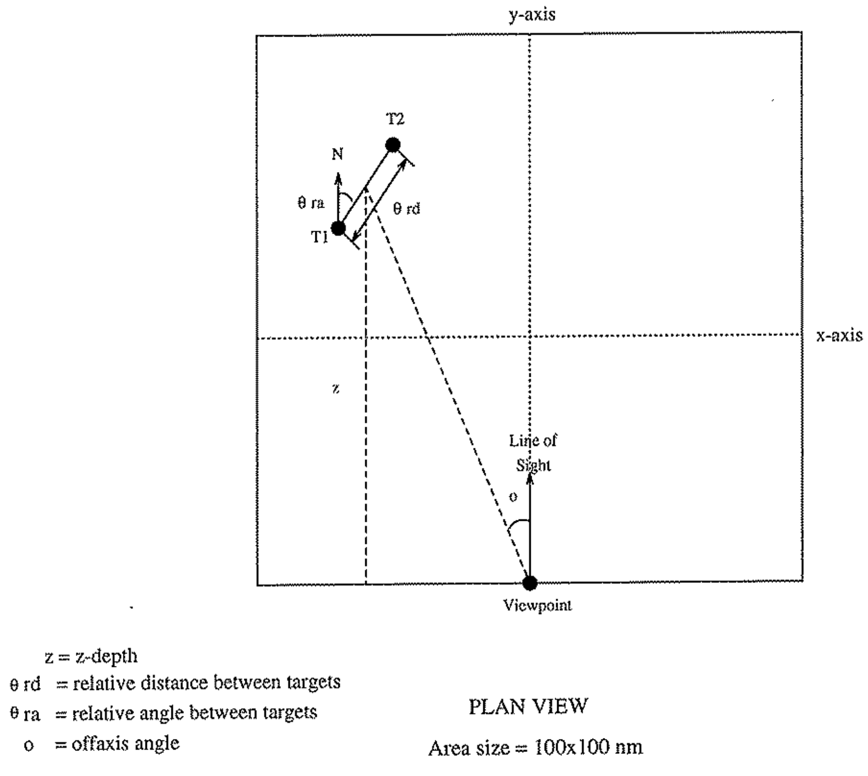

Figure 3.3: Geometry of Parameter Reading Experiment

3.2.2 Null Hypothesis

There is no effect between projection type and the reading of relative azimuth angle and horizontal distance between targets in a 3D display.

3.2.3 Methodology

In the following discussions, the symbols below are used (see Figure 3.3 for an illustration of some of the symbols):

(x,y) world coordinates of the bisector of a vector joining the two targets T1 and T2. (x = distance east of the centre of the map, y = distance north of the centre of the map).

$(T1_x, T1_y)$ world coordinates of target 1.

$(T2_x, T2_y)$ world coordinates of target 2.

θ_{ra} true azimuth from T1 to T2.

θ_{rd} true horizontal distance between T1 and T2.

λ_{ra} observed azimuth from T1 to T2.

λ_{rd} observed horizontal distance from T1 to T2.

ϵ_{ra} angle error: $\epsilon_{ra} = \lambda_{ra} - \theta_{ra}$.

ϵ_{rd} distance error: $\epsilon_{rd} = \lambda_{rd} - \theta_{rd}$.

Subjects were shown either a parallel or perspective projection. Twenty stimuli, each composed of a pair of targets, were administered in random order. Subjects were asked in each case to estimate the bearing from target 1 to target 2, and distance between targets.

The stimuli were generated randomly. The ranges of the randomly-generated parameters are shown in Table 3.1. The values of the stimuli used are tabulated in Table B.1 of Appendix B.

$$\begin{array}{ccccc} -45\text{nm} & \leq & x & \leq & +45\text{nm} \\ -45\text{nm} & \leq & y & \leq & +45\text{nm} \\ -180^\circ & \leq & \theta_{ra} & \leq & +180^\circ \\ 2\text{nm} & \leq & \theta_{rd} & \leq & 20\text{nm} \end{array}$$

Table 3.1: Ranges of randomly-generated parameters

3.3 Recall Experiment

3.3.1 Introduction

The purpose of this experiment was to assess the effect of display type on the assimilation into memory and recall of a static traffic scenario.

Memory of spatial relationships may be more persistent in one display format than another, and this could have advantages. For example, if the radar service is lost for a duration, immediately after loss controllers must attempt to position aircraft in order to avoid loss of separation (usually by attempting vertical separation by altitude changes) and then either fall back to procedural control with flight strips or hold aircraft at points until radar service is resumed. A display which gives a more durable memory trace might have advantages in this type of situation.

If a scenario could be shown to be assimilated into a mental model more rapidly in a 3D presentation than in a 2D presentation, then the 3D presentation might have advantages in situations where rapid assessment of a scenario is necessary; for example at the resumption of radar service, or at change-over between controllers at the end of a shift.

This experiment set out to make a preliminary investigation of this area.

A similar experiment was carried out by Burnett and Barfield [BB91]. In their experiment, Burnett & Barfield compared performance for reconstructing a scenario between two-dimensional and three-dimensional perspective air traffic control displays, for two different densities of traffic. Prior to the task, air traffic controllers were required to memorise FPS information and then watch the assocaiated scenario over 70s of animation. They were then asked to reconstruct the final frame of the scenario, giving position, callsigns, altitude, speed and heading of aircraft. Two traffic densities were used: light (7 aircraft) and heavy (17 aircraft).

It was found that performance for the reconstruction task in the both traffic scenarios was about the same for both display types; however, in the perspective display format, aircraft placement was found to be consistently 3cm north of the actual aircraft location, whereas horizontal placement was accurate within $\frac{1}{2}$cm of actual aircraft position.

3.3.2 Working Memory and Chunking

Human memory is a complex phenomenon with many different characteristics. Psychologists therefore often find it convenient to model different memory characteristics as different types of memory. The type of memory under consideration here is known as *working memory*. As described by Wickens & Flach [WF88], working memory is employed when a person hears a number and must enter it on a keyboard, or when he/she must recall the relative positions of blips on a radar display after a brief scan. These examples illustrate two different *codes* in working memory; verbal information is normally retained using an acoustic-phonetic rehearsal, and spatial information is normally retained in working memory using a visual code. Wickens & Flach also state that there is evidence to suggest that visual codes are less easily rehearsed than verbal codes.

The number of items that can be retained in working memory is quite small, the most oft-quoted figure being 7±2 items. If more items are presented, then to be recalled they must be consolidated into related clusters, a principle referred to as *chunking*. Chunking thus can expand the number of individual items which can be retained in working memory.

Working memory is more likely to treat dimensions of a single object as a single chunk than the same dimension of several objects. For example, in an ATC problem, altitude, airspeed, heading and size of two aircraft would be retained better than the altitude and airspeed of four aircraft, even though in each case eight items are to be held in working memory.

Regarding spatial codes, analogue pictures are the most useful mode for storage in working memory, print the least [Lea94].

3.3.3 Hypotheses

The accuracy and speed of assimilation and recall of spatial relationships are expected to be influenced by the type of display.

Regarding speed of assimilation, in a 2D plan-view, vertical information is represented textually, whereas in the 3D display it is shown graphically, albeit subject to some distortion. It is expected that an approximate spatial mental picture will be formed faster given a 3D presentation than a 2D presentation for two reasons:

1. Analogue pictures are more useful for storage of spatial codes in working memory than text, and may thus be assimilated more quickly.

2. In the 2D PPI display, the integration into an internal 3D spatial model must be done mentally, whereas it is already presented in an integrated analogue form in a 3D display.

A three-dimensional display should therefore give a lower mental workload in forming an integrated spatial mental model than a 2D PPI display and thus be quicker to assimilate, and may have a more durable memory trace (since the relationships are remembered in analogue form which is more useful to working memory than text).

Regarding positional accuracy of the recalled image, however, it is expected that the recall of horizontal positional information in the 2D display will be more accurate than in a 3D display, since in a 2D PPI display, horizontal information (horizontal distances and azimuth angles between targets) is shown unambiguously, whereas in a 3D display it is subject to distortion unless viewed in plan. As summarised in Wickens, Todd and Seidler [WTS89]:

> ...the improved holistic awareness of space, gained by 3D representation, may be gained at the expense of analytic detail.

To summarise, it is conjectured that 3D displays will be better for conveying rapidly approximate spatial relationships because these can be seen directly as an analogue picture. However, this will be at the expense of analytic detail, such as precise horizontal and vertical position.

3.3.4 Null Hypotheses

The display format is expected to have no effect on the speed and accuracy of the recall of three-dimensional spatial information.

3.3.5 Methodology

An exhaustive investigation of what types of information are best recalled in what display format, and the interaction of parameters such as number of depth cues on speed and accuracy of recall, is quite beyond the scope of this investigation. This research programme's emphasis is looking at the feasibility of 3D displays for spatial command/control applications, so the objective here is to conduct a preliminary investigation into what sort of errors and problems one can expect from a 3D display, and whether or not 3D display technologies are suitable for this type of application.

The approach adopted was an adaptation of the work of Burnett & Barfield, described above. For this research, it was felt that their approach was too specific to ATC to relate to command/control applications in general, was too tied to current working practices, and did not concentrate on the effects of the display alone. (Recall would be partly influenced by the flight strip information, although the only variable was the display so any variations should have been due to that factor alone.). The methodology chosen here was therefore to try to memorise a *static* traffic scenario with information from the display *alone*.

Some preliminary research was conducted into how many aircraft should be in the stimulus scenarios, and how long the stimuli should be presented for. Too long a presentation with too few aircraft might remove any differences in performance between the display formats; more aircraft for a shorter a period of time might make differences in performance between the display types more apparent (since spatial relationships would have to be assimilated very quickly) but too short an exposure period with too many aircraft might swamp the subjects. Moerover, memory is not a simple phenomenon, and different types of stimulus are remembered to different degrees in different ways; the best advice received was to try to determine the number of aircraft and the time of presentation empirically. However, there was no time to investigate this area before the pilot study, so an "educated guess" had to be made.

In the end it was decided to show 15 aircraft for 60s. This was intended to place a deliberate heavy burden on the memory capacity of the subjects so that any differences in speed of assimilation might be shown, but it was thought that it might have a demoralising effect. For the main research, it would probably be advisable to follow the approach of Burnett & Barfield in selecting two traffic densities crossed with the display types.

In this experiment, two static scenarios containing 15 aircraft each were presented to subjects for 60s; the first for training, the second as the main run. Subjects were allocated between two different display formats: 2D PPI or pseudo-3D TTW. For each scenario, subjects were asked to try to memorise the positions and altitudes of each aircraft. After 60s, the scenario was removed and the subjects were given a piece of paper on which was printed the same area as presented in plan view, complete with range rings and video map outlines. Subjects were then asked to try to mark the positions of the aircraft on the piece of paper and to indicate their heights were possible, either by writing a number and/or by drawing the length of the "drop lines" (vertical lines connecting each aircraft to the ground, representing height) for each aircraft.

The reconstruction was to be timed, and evaluated in terms of accuracy of recall (number of aircraft, and horizontal and vertical placement accuracy) and time taken to reconstruct the scene. In the full experiment, it is intended that reconstruction will be carried out using an computer program so that the accuracy and time can be determined automatically.

3.4 Conflict Detection Experiment

3.4.1 Introduction

As stated in chapter 2, the prime task in air traffic control is to "maintain the safe, orderly and expeditious flow of traffic", and this is done by long-term (strategic) planning using flight strips, and short-term (tactical) planning and monitoring using the radar.

Because a three-dimensional display presents an integrated picture of the spatial disposition of the air traffic, with reduced cognitive workload to interpret it, there may be advantages in using it for long-term planning in conjunction with the flight strips. Although in long-term planning, it may not be necessary to visualise the traffic in spatial terms, nevertheless, an integrated presentation of the traffic situation may help. However, the rôle of a 3D display in long-term planning is difficult to assess within the limits of this research programme, as it would involve long interactive scenarios.

This research therefore concentrates on whether or not the 3D presentation is more or less effective in the purely monitoring and short-term planning rôles — i.e. using the radar-derived positional information to detect and resolve potential conflicts. Detection of conflicts may be taken as a measure of controller awareness — if a controller is more aware of the current state of the traffic (s)he will detect more conflicts. It is conjectured that because of its integrated presentation of the three spatial dimensions and reduced mental integration workload, there will be a situational awareness of the traffic state with the 3D display than with the current PPI.

To maintain safety, aircraft must be greater than specified minimum horizontal and vertical distances apart. If two aircraft are at least the specified distances apart, they are referred to as *separated.* A *conflict* is defined here as a situation which, left unmodified, will lead to a loss of separation, or where separation has already been lost. Each aircraft can be thought of as having a protective "envelope" around it which no other aircraft is permitted to violate. Minimum separation in ATC varies with situation, but typical minima are at 1 000ft vertically, or 5nm laterally—i.e. if aircraft are closer than 1 000ft vertically, then they must be 5nm or more apart horizontally.

If these separations can be assessed only with difficulty in a 3D display, then it places a question mark over the effectiveness of a 3D display for these applications. A conflict detection experiment may help to resolve this issue.

Having detected a conflict, the way in which it is *resolved* is also important. This may require split-second decisions to be made and these are critical, since aircraft must be instructed to manœuvre to avoid conflict with each other and also so that they do not come into conflict with other traffic in the vicinity as a result. A study by McGreevy and Ellis [EMH84] into cockpit air traffic displays showed that pilots made more traffic avoidance manœuvres in the vertical plane with a 3D display than with a plan-view presentation augmented with digital height readouts. If similar behaviour were demonstrated in the ATC context, this may give controllers an "extra dimension" in which to work when resolving conflicts.

A 3D presentation may also help in appraising a developing situation and issuing a resolution instructions accordingly. This could make it useful in conjunction with automated predictive systems such as the Short Term Conflict Alert (STCA) system currently being installed in the UK.

3.4.2 Null Hypotheses

Conflict detection performance, in terms of speed of identification of potential conflicts, is independent of display type.

No more vertical conflict resolution instructions will be given using the 3D display types than in the 2D display. The effectiveness of the manœuvre issued (in terms to separation achieved between aircraft and hazards to other aircraft in the immediate vicinity) will not depend on display type.

3.4.3 Methodology

This experiment conducts a preliminary investigation of conflict detection. Subjects were shown dynamic scenarios in which one or more conflicts exist and were asked to identify conflicts and to issue resolution instructions. It was intended to assess performance by the number of conflicts detected and the time taken to detect the conflicts.

In practice, flow management would be done in conjunction with flight progress strips. This experiment, however, tries to concentrate on use of the display alone in making spatial judgments. All necessary information therefore had to be contained in the display; extra information incorporated included route code (the intended destination of the aircraft — this is shown in the datablock of air traffic control displays) and, for aircraft changing heights, the cleared level.

There are two nuisance variables; the type of conflict and the relationship between the plane of the conflicting aircraft and the line of sight into a 3D display.

Conflicts may be divided into three categories, *viz*:

1. Aircraft approaching head-on, either co-altitude, or assigned the same altitude but climbing/descending, or passing through the same altitude. (Figure 3.4(a)).
2. Aircraft crossing the same geographical point (Figure 3.4(b)).
3. Aircraft overtaking on the same airway (Figure 3.4(c)).

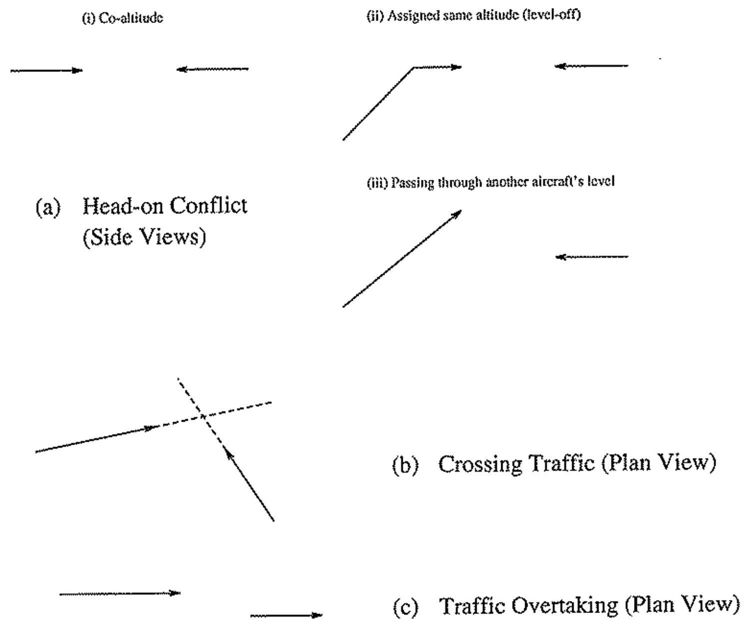

Figure 3.4: Conflict Types

A conflict detection experiment should include representatives of each of these types.

The relationship between the plane of conflicting aircraft and the line of sight into a 3D display also needs to be considered. If two aircraft are converging head-on, the situation will be clearer when viewed from the side (orthogonal to the plane of the conflict) than when viewed head-on (parallel to the plane of the conflict). Scenarios should therefore be devised to include examples of each type of conflict with different relationships to the virtual camera in each case.

Because of the inherent ambiguity of 3D displays, it is also desired to know whether this present a problem in detecting conflicts. Marginal situations should therefore be included.

Finally, the traffic density and complexity may be factors. These may be varied in the main experiment.

The purpose of the pilot experiment was to validate the experimental procedures. Therefore, only two scenarios were presented, and not all combinations of conflict type, plane of conflict, display type and traffic density were evaluated.

4 Displays

4.1 Introduction

This chapter describes the displays of air traffic position data used in the experiments.

CAA NATS provided processed radar data taken from the Heathrow radar for the purpose of these experiments (15 minutes of data taken from ~09:00–09:15 on a morning in April 1994). The origin of the world coordinate space was therefore chosen to be at the radar head at London Heathrow airport, with the world positive x-axis running East from this point, the positive y-axis running North and the positive z-axis running upwards.

Some figures illustrating the display formats used are included in this chapter. It should be borne in mind that some of these are only schematic and for those figures which are taken from actual "screen dumps", the quality of the small greyscale print does not compare with the colour high-resolution display. The reader is therefore advised not to take these as being entirely representative.

4.2 Display Content

4.2.1 General

There is a wide variation amongst existing PPI radar displays as regards to, for example, the refresh method of the CRT (calligraphic with long-persistence phosphor or raster with short-persistence phosphor), the use of colour, and symbology (anything from display of raw MTI-processed radar returns, with or without computer-generated symbology, to a fully synthetic display). The pilot experiment subjects and the main experimental subjects therefore have experience of a wide variety of radar displays. Rather than try to replicate the look of one of the existing displays, it was decided to implement a mixture of emerging standards and "home-brew" symbology, with the pilot subjects giving feedback as to the success or otherwise of the chosen format. Given the variation amongst displays, ATCOs seem to be flexible as regards the format used so long as the necessary information is clear.

The colours were based on the latest interim NATS colour standard [RM92]. As specified in the standard, aircraft-related symbology were displayed on top of all other symbols (map and datablocks) so that they were always visible.

All displays showed air traffic in a 100×100 nautical mile area around the radar head at Heathrow. A "video map" was provided which showed the London Heathrow and Gatwick CTAs and CTRs, the boundary of London Terminal Manœuvering Area (LTMA) and parts of various airways, as well as a section of coastline. Map data were obtained from CAA NATS. As an aid to assessing distance, range rings were superimposed on the map at 10nm intervals centered on the radar head. (Range rings were shown instead of a square grid since these are more familiar to ATCOs, being provided as an option on a PPI radar display.) A schematic of the video map area and range rings are shown in Figure 4.2.1.

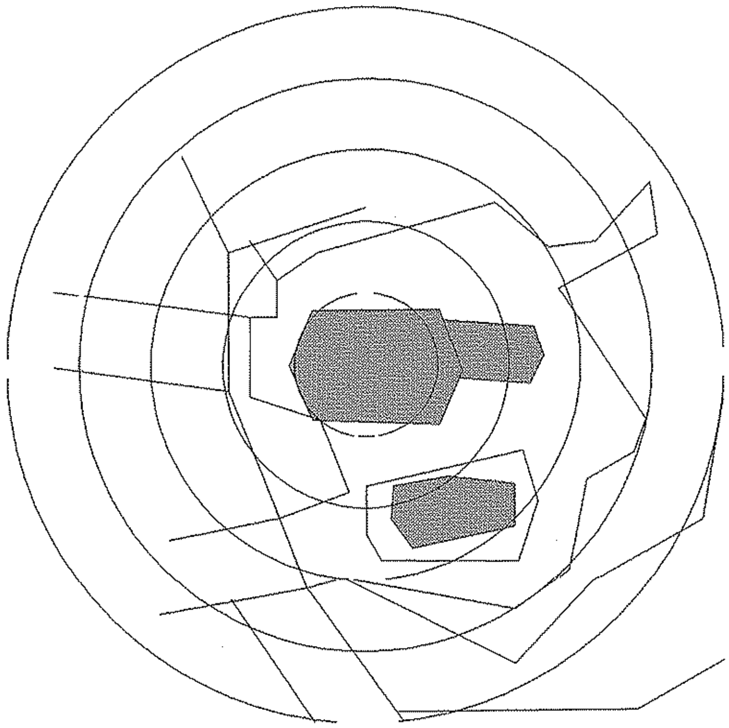

Figure 4.1: Display Area: Video Map and Range Rings

4.2.2 Datablocks

Alphanumeric information for each aircraft was presented in a datablock, tied to the respective aircraft by a *leader line* of single pixel width.

Datablock content is more standardised between displays in the UK, generally including flight number or mode-A transponder code, mode-C code (for aircraft with altitude-encoding transponders) and destination code (for aircraft with flight plans logged in the FDP computer). In these experiments, datablock content varied with the task. The most information shown included:

- Mode-A code.
- Mode-C code (height/altitude in hundreds of feet, or flight level, as reported by the air data computer aboard the aircraft).
- Climb/descend status symbol. This was derived from the current and previous mode-C code for each aircraft. A caret was used to indicate a climbing aircraft, a lower-case vee to indicate a descending aircraft.
- Cleared altitude/flight level.
- Destination code. This either refers to a route or to an aerodrome; in the latter case, the last two letters of the aerodrome's ICAO code are used.

A schematic of an aircraft symbol and its associated trail and datablock is shown in Figure 4.2. The figure shows aircraft with mode-A ident 1245, currently at FL220 and descending to its cleared level of FL110, destination London Gatwick airport (airport ICAO code EGKK, abbreviated to KK in the destination code field).

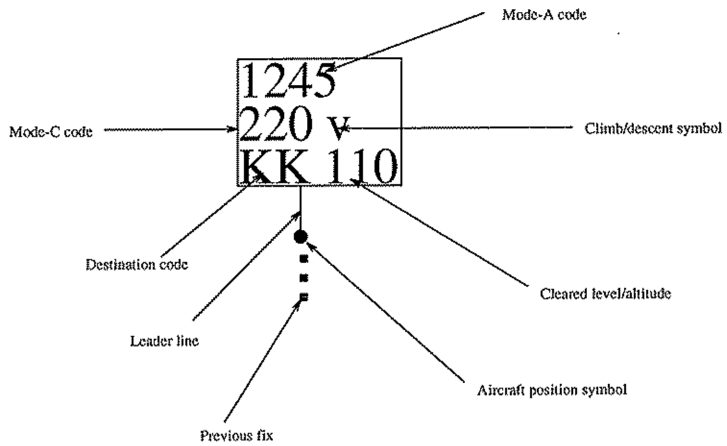

Figure 4.2: Aircraft Position Symbol and Datablock

A simple algorithm was used to help to avoid datablocks overlapping. In this algorithm, the rectangular bounding box of the datablock is computed and first tried in the "North" position relative to the corresponding aircraft symbol. If this overlaps with another datablock already displayed, then it is tried in the "Northeast" position, and so on for the six positions (E, SE, S, SW, W, NW). It is displayed in the first of these positions in which it does not overlap another datab;ock — otherwise, if it overlaps in all eight positions, it is displayed in the "North" position by default. The nearest corner or edge of the datablock's bounding box to the aircraft position was set at 6 pixels away from the centre of the aircraft position symbol.

For decluttering in the 3D TTW displays, datablocks showed only the mode-A code by default. Pressing and holding the SPACE bar then caused the whole of the datablock to be displayed.

4.2.3 Display Types

This section describes the format of the four different display types. Figures 4.3, 4.4 and 4.5 illustrate the same traffic pattern in the 2D, 3D TTW perspective and 3D TTW parallel projection formats respectively.

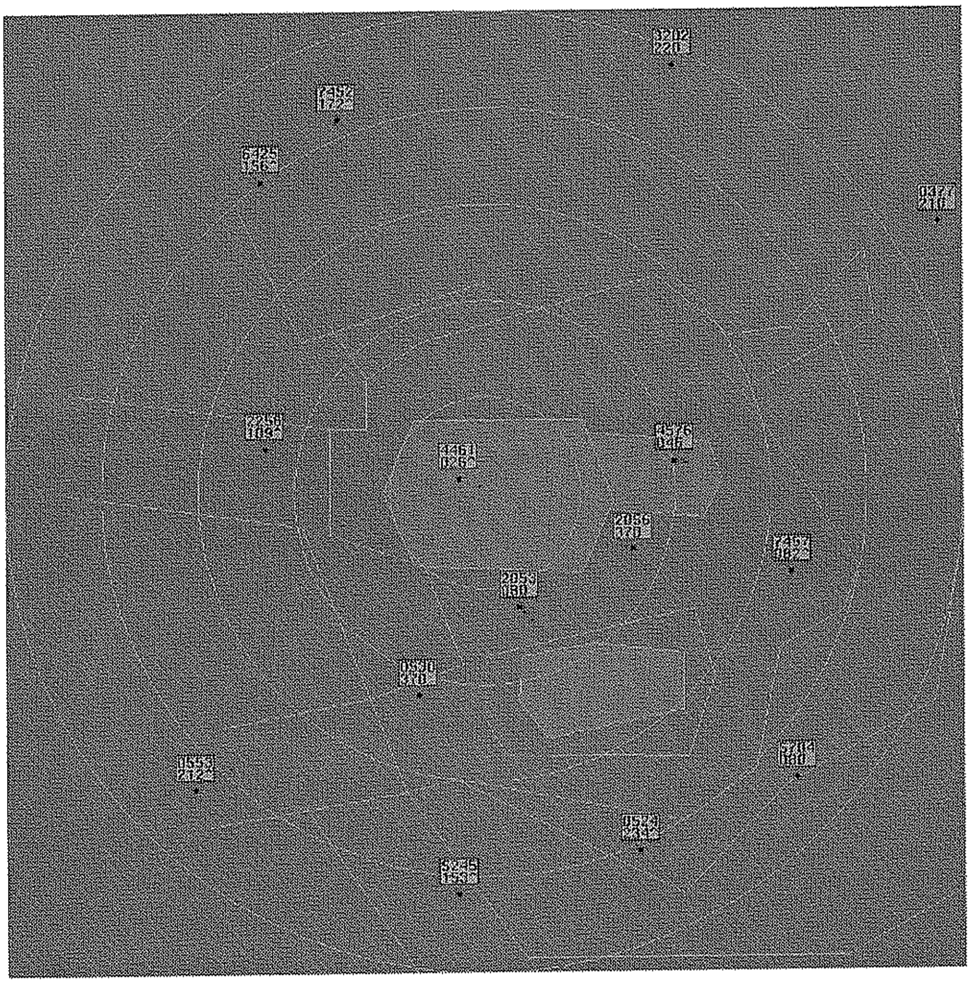

Figure 4.3: 2D aircraft position display

The 2D and TTW displays were implemented on a Silicon Graphics (SGI) Indy workstation, using the C programming language and the SGI GL graphics library. The VR display was implemented in C on a Division ProVision 100-VTX machine running Division's dVS virtual environment operating system.

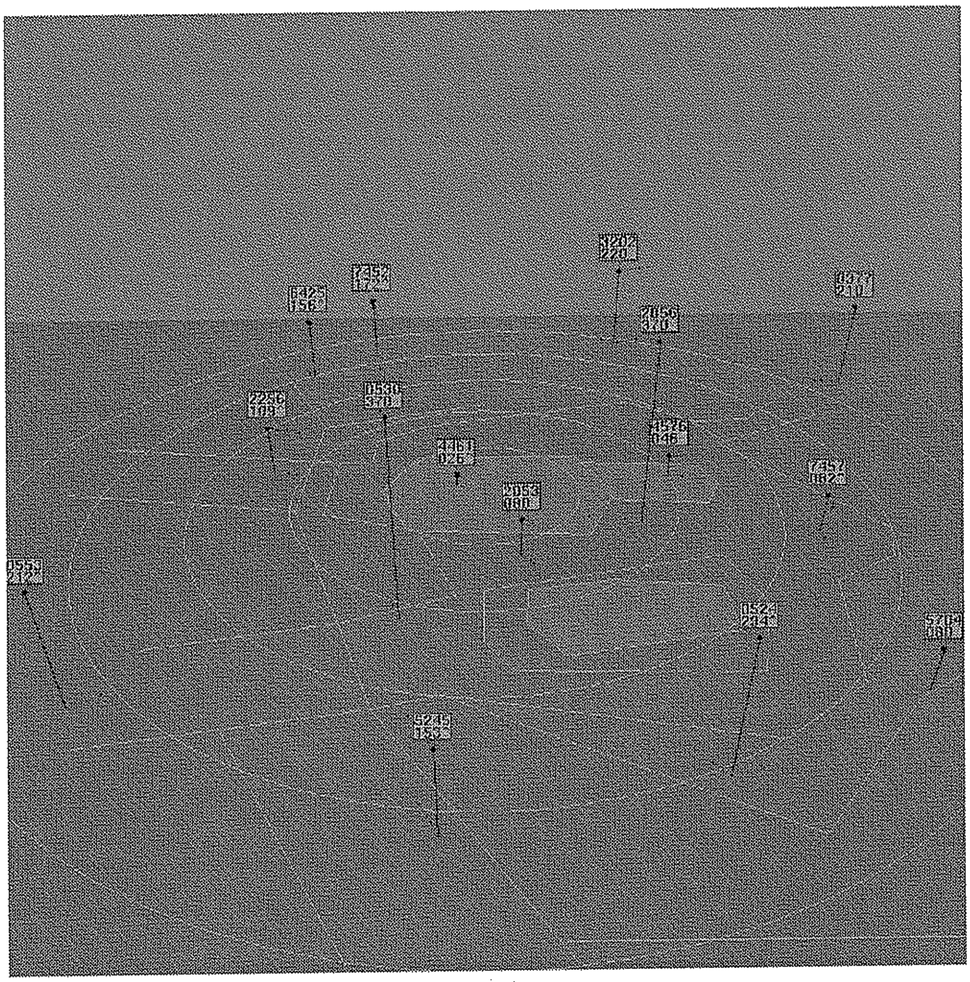

Figure 4.4: 3D TTW aircraft position display: perspective projection

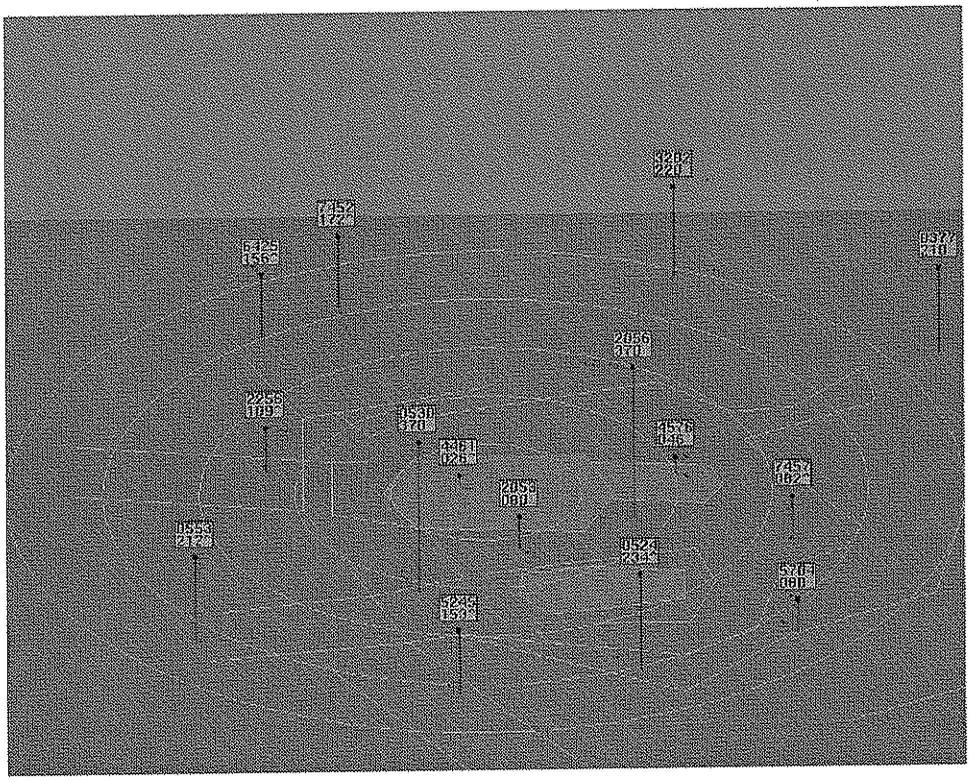

Figure 4.5: 3D TTW aircraft position display: parallel projection

2D Display

In the 2D display, the aircraft current position was represented by a black circle of 2 pixels radius centered on the last fix. Aircraft trailing histories (up to 8) were displayed as single black pixels. The positions of aircraft were updated every 6 seconds (corresponding to the radar sweep time).

When aircraft positions were updated, the symbols moved in single pixel steps rather than in large jumps; this was to enable smooth updating, which is cited by Strutt [Str91] as being less distracting and therefore preferable. However, the datablocks tended to 'jump' whenever the datablock avoidance algorithm dictated a change in their position relative to the aircraft symbol, since the datablock overlap avoidance was re-computed for every frame during the smooth movement of aircraft symbols.

3D TTW Display

The same symbology was used as in the 2D display, except that the black air position symbols also had corresponding white ground symbols representing their "shadows" from a light source at infinity above the centre of the display. The air and ground current position symbols were joined by a vertical drop line. Vertical exaggeration was $\times 4$. Note that symbol size was constant and did *not* vary with z-depth in either the parallel or perspective projections.

For all experiments except the parameter reading experiment, a perspective projection was used. In the parameter reading experiment, parallel projection was also used. For the perspective display, the viewing parameters were:

- Geometric field of view = 40°
- Azimuth = 180° (i.e. looking from the south)
- Elevation = -30° (i.e. looking down at an angle of 30° from the horizontal)

The same azimuth and elevation were used in the parallel display. In both cases, the camera position was chosen so that the view covered 100 nautical miles east-west at the centre of the displayed area (i.e. so that all the range rings were fully visible).

3D Stereoscopic Display

The 3D TTW stereoscopic display was implemented using time-multiplexing of left eye and right eye images with LCD shutter glasses (StereoGraphics CrystalEyes, loaned from Division Ltd.). The model of Indy workstation used in these experiments did not contain a hardware stereo buffer, so the XSGIStereo extensions to X11 windows system (which works with GL) were therefore used to implement the stereo buffering. Under this scheme, the display memory is divided into two sections: the memory for the top half of the screen is allocated to the image for one eye and the memory for the bottom half is allocated to the image for the other eye. The monitor is then switched to display only the top 50% of scanlines, and the system then displays the top and bottom halves of the screen memory (i.e. left and right eye images) within the top 50% of scanlines alternately at 60Hz, synchronised with the switching of the LCD glasses. This results in a stereo-capable display, but with only half the vertical resolution of which the system is capable.

Another problem is that the model of Indy workstation used does not support hardware stereo double buffering for animation. Animation was therefore achieved by manipulating the colour lookup table in the colour-mapped display.

The stereo display was only shown to a few subjects and was not evaluated quantitatively. The reason for this was that in the display as developed for the pilot study, the datablocks were only drawn in the left eye view and so were not displayed at any "depth". (This was due to the fact that the datablock overlap avoidance algorithm would otherwise need to be modified, since it must operate in screen space but the left and right eye datablock images must be shown at their projected positions from world space to achieve stereopsis.) It was originally thought that this would present little problem, but the display could not be evaluated even for development until just before the pilot study because the CrystalEyes were not available until then. When the display was evaluated just before the pilot study, it was found that the datablocks being displayed with no depth was conflicting with their corresponding aircraft being displayed with depth, at least partially destroying the depth perception unless a conscious effort was made to ignore the datablocks. It was felt that this would not allow the displayed to be evaluated effectively.

VR Display

In the VR display, the scene is shown on a binocular head-mounted display (HMD), which comprises a helmet containing two colour LCD display screens (resolution 360×240 pixels) and wide-angle optics to show the separate images to each eye. The display fills the field of view and none of the outside world is visible when wearing the display. The display is visually coupled (i.e. the scene changes in response to the viewer's head movements) using a Polhemus FastTrak AC electromagnetic 3D position sensor mounted on the HMD.

The appearance of the display on the VR machine was necessarily different to the other displays due to the rather more limited resolution and the fact that the ProVision cannot draw line segments (only polygons) and cannot display primitives in screen space (only world coordinate primitives are allowed).

The camera parameters were as provided by the dVS virtual environment. A perspective projection was used. The ground detail (video map and range rings) was generated using a texture map (i.e. mapping an image onto a polygon mesh). Aircraft drop lines were represented by equilateral triangle cross-section tubes. Symbology associated with the aircraft were tetrahedra — white for the aircraft air position symbol, with smaller tetrahedra for the trailing histories: cyan for air plots and yellow for their ground shadows (no ground shadow was provided for the aircraft position symbols). Since these objects are defined in world coordinate space on the ProVision, they are subject to perspective and are rendered smaller with increasing z-depth.

The world was scaled to $6'$ per 100nm horizontally (so the display area was $6' \times 6'$) and $2'$ per $10\,000'$ vertically.

Textual information (datablocks) were omitted. It was found that in earlier prototype displays, because the text characters had to be defined as polygons in world coordinate space, they needed to be re-scaled and re-orientated every time there was a change in viewer head position or orientation in order to be displayed at a fixed size and always facing the viewer. This impacted severely on system performance. Because the text characters were defined in world coordinate space, the datablock overlap avoidance algorithm could not be implemented (this must work in screen coordinate space) so labels frequently overlapped. This was compounded by the fact that because of the poor resolution of the display, the text had to be scaled to quite a large size to be legible at all. Scaling the text to fixed apparent size also had the rather incongruous effect that when the whole scene was viewed from far away, the aircraft symbology was very small compared to the text, which dominated the image.

It was therefore intended to adopt a "selective datablock" strategy for the pilot and subsequent displays. The idea was that normally, no text would be shown at all, and the viewer would have to pick out an aircraft explicitly to get more information on it; for example, by having a "beam" coming from the cursor (representing the hand position) at the press of a button which would then be intersected with the desired aircraft. However, this was not ready in time for the pilot experiment, largely because the prototype display described above was implemented on an earlier model of ProVision running a different and incompatible version of dVS and so the code had to be ported to the new platform. (This was done because the other machine lacked the texture mapping facilities that allowed a video map to be displayed.)

This precluded the VR display from being included in any quantitative experiments. Subjects were shown the VR display without datablocks, however, and allowed to use it freely, to gain subjective feedback.

5 Experiment Design and Procedures

5.1 Introduction

This chapter describes the design of the experimental tasks, and the procedures it was intended to carry out. How the tasks were conducted in practice is detailed in §6!

Because this was only a pilot experiment, it was rather reduced in scale and scope compared to a full study; it was used to validate experimental techniques rather than to gain any useful data. Therefore, elements of the investigation which will be included in the main experiment, such as evaluation across all four display types, perhaps with multiple levels of traffic density and complexity for some of the tasks, and comparative experiments with a student population, were not carried out here.

Quantitative tasks were conducted using a Silicon Graphics Indy workstation with only the 2D PPI and pseudo-3D TTW displays (the stereoscopic 3D TTW and VR displays being incomplete at the time, as explained in chapter 4). After the quantitative tasks were complete, subjects were invited to comment on the experiments and were offered the opportunity to use the VR display.

It is intended to have the experimental programme automated for main study, to the extent that the entire session would be run by computer with minimal intervention by the supervisor; however, there was insufficient time to automate the task presentation to the required degree, so some degree of manual setup was required by the supervisor.

5.2 Subjects

Nine subjects participated in the pilot study over the course of a week. These were ex-ATCOs currently working in various non-operational capacities within CAA NATS (although one hoped to return to operational duty in the near future). Subjects had a variety of operational backgrounds, ranging from civil ATC without radar to military air traffic control. All were enthusiastic and interested in the research, and were commendably objective, despite having reservations about the applicability of 3D displays to current operational practices.

5.3 Pre-Experiment Briefing

On arrival, subjects were asked to read a background information sheet (see §A.1) and the instruction sheet for the first task (see §A.2) whilst the first task (parameter reading experiment) was prepared. No questionnaire was given to subjects, the type of information required having yet to be determined (apart from demographic information, which would have been useless for the pilot study). It was explained that the experimental programme was a pilot study, and that open, constructive criticism was welcome. It was intended to have subjects write comments at the end of the whole experiment.

5.4 Task 1: Parameter Reading

This task was to read the azimuth and distance between 20 pairs of aircraft on a 3D display. Subjects were allocated into one of two groups: parallel projection or perspective projection.

Having read the instruction sheet (§A.2), subjects were shown the display. The supervisor repeated the instructions and answered any questions. Subjects were then given a response sheet on which to record their answers, and wrote down the first response with the supervisor watching. That having been done, the supervisor left them to complete the task on their own (although being nearby in case of any problems). After this, they were debriefed before the next task.

5.5 Task 2: Recall

The purpose of this task was to examine the recall of a static air traffic scenario. Subjects were allocated into one of two groups: 2D PPI presentation or 3D TTW presentation (perspective projection). The subjects were requested to read the instruction sheet (§A.3) whilst the task was being set up.

It was intended to present subjects with two static scenarios each containing 15 aircraft for 60s each. After this time, the scenario would be removed and subjects would re-create the scenario on a piece of paper showing the airspace area in plan view, with range rings and video map, but without aircraft, attempting to recall position and height of all the aircraft. Subjects were asked to try to mark the positions of the aircraft on the piece of paper and to indicate their heights where possible, either by writing a number or by drawing the length of the drop line for the aircraft. The first scenario was intended as a training exercise, the second as a genuine 'run'. After the second run, subjects were to be debriefed.

The reconstruction was to be timed, and evaluated in terms of accuracy of recall (number of aircraft, and horizontal and vertical placement accuracy) and time taken to reconstruct the scene.

5.6 Task 3: Conflict Detection

Following debriefing from the second task, subjects were asked to read the next instruction sheet (§A.4) whilst the supervisor set up the conflict detection task. Subjects were allocated into one of two groups: 2D PPI or 3D TTW display (perspective projection).

When subjects had read the instruction sheet, a real traffic sample was shown in order to familiarise them with the dynamic display, whilst the instruction sheet was read by the supervisor. In the 3D display, the use of the SPACE key in presenting a full datablock was demonstrated. Subjects were then left to examine the display for unlimited time until they reported ready.

Two dynamic scenarios were then presented, each of 3 minutes duration. Subjects were asked to identify conflicts and give resolution instructions. There was one actual conflict per scenario, but this number was not revealed to the subjects. It was intended to record the time to detect the conflict to assess situational awareness, and to ask for a resolution instruction in case of a conflict in order to see whether a 3D display encouraged more use of the vertical plane rather than the horizontal in issuing avoidance instructions. The first scenario was intended for training; the second was recorded.

Following the two scenarios, subjects were then debriefed on the task.

The two scenarios contained one conflict each, and were based on real traffic samples, with one computer-simulated aircraft to cause the conflict. The first scenario had two aircraft on the same airway travelling in the same direction, but slowly converging, with the lower climbing through the level of the upper. The second scenario contained a crossing conflict: a aircraft level at high altitude against a climbing aircraft crossing its path. This second conflict actually resulted in a merged plot on the display, and a minimum separation of less than 0.1nm (i.e. a probable collision).

5.7 VR Display

Having completed the quantitative tasks, subjects were given the opportunity to try the VR display prototype. Subjects were asked to read the instruction sheet (§A.5), which described the VR display and its operation, before trying the display.

Before donning the HMD, subjects were again briefed verbally, and the HMD components pointed out. When subjects put the HMD on, they were given a hand-held "3D mouse" control device which was displayed as a cursor in the virtual world corresponding to their hand position. Navigation in the virtual environment was then described and subjects were asked to perform simple navigation tasks until they were used to it. They were then allowed to explore the environment at their leisure for unlimited time. The environment showed a culled set from the real traffic sample.

6 Results

6.1 Introduction

This chapter presents the results of the experiments, both subjective and quantitative, where such were obtained.

6.2 Subjective Data Gathering

As explained above, it was originally intended to have subjects fill out a comment sheet having completed all quantitative experiments. However, this proved to be impractical, since the subjects were commenting whilst the tasks were in progress, and if comments had been deferred until the end of the experiment, some useful feedback may have been lost.

6.3 Parameter Reading Task

6.3.1 General

The pilot study was carried out with ex-ATCOs of varying experience. Several commented on the fact that they were "rusty" and that current ATCOs would be able to estimate angles and distances with greater proficiency, this being done routinely in radar control work.

None of the subjects had used a 3D display previously. Most reported that subjectively, the non-uniformity of ground distance with position on the display (and allied to this, the distortion apparent in the range rings) presented a difficulty. Regarding the perspective display, one subject commented that the fact that the drop lines of the same *actual* height would appear to be *different* lengths depending on the z-depth due to perspective, and that this would negate any usefulness which they might otherwise have.

One subject commented on the lack of a north index in the display for a reference, but was able to perform the task when it was indicated that the camera azimuth was 180°.

Subject 7 reported that estimation of distance presented more of a problem than heading.

One subject who had had operational experience within LATCC reported that he used his local knowledge (knowledge of relative bearings and distances of features on the ground map) as an aid.

A couple of subjects used a pen up against the display screen to measure distances instead of "eyeballing". (The instructions did not prohibit this—in fact, they did not mention it.) This technique is not uncommon amongst operators of plan view radar displays.

6.3.2 Instruction Sheet

The instruction sheet should be modified for subjects to make it less specific to ATCOs. It used an incorrect technical term (it talked about *heading* rather than the correct term which was *bearing*). Non-ATCOs would not necessarily be aware of a distinction, so a neutral term (e.g. azimuth angle) should be used, or clarification should be made, perhaps with a diagram.

The instructions also omitted to clarify whether to use the air plot or the ground plot of aircraft position for the task (in fact, in this task it made no difference since the aircraft were co-altitude, but some subjects requested clarification).

Even though the instruction sheet pointed out that the bearing was to be taken from aircraft 1 to aircraft 2, this still caused some initial confusion with some subjects providing reciprocal bearings for one or two responses until they recognised their mistakes and queried the supervisor. (Although the bearing of target 2 from target 1 is just the reciprocal bearing of that from target 1 to target 2, and subjects corrected their mistakes by adding (or subtracting) 180° from the first bearing they gave, requesting the bearing or its reciprocal may influence perception.)

All of the above would probably be rectified by clearer instructions, perhaps illustrated with examples (either on the sheet itself, or on the computer as a "training" exercise).

6.3.3 Quantitative Results

Out of the nine subjects who took part, four were shown the parallel projection, five the perspective projection. Raw and statistical data are tabulated in Appendix B, and figures 5.1–5.4 show scatter plots of the observed versus actual angles and distances for parallel and perspective projections.

Angle

For the parallel projection, linear regression analysis gave:

$$\lambda^1_{ra} = 6.43 + 0.14x + 0.96\theta_{ra} \tag{6.1}$$

with $R^2 = 0.99$ and a significance $P = 8 \times 10^{-75}$. (This means that the model accounts for 99% of the variation in the data, and the chance of the same results being produced from random data was 8×10^{-75}.) For the perspective projection, linear regression gave:

$$\lambda^2_{ra} = 26.01 + 0.82\theta_{ra} \tag{6.2}$$

with $R^2 = 0.72$ and a significance $P = 4 \times 10^{-21}$.

For the angle error term ϵ_{ra}, in the parallel case:

$$\epsilon^1_{ra} = 6.43 + 0.14x - 0.04\theta_{ra} \tag{6.3}$$

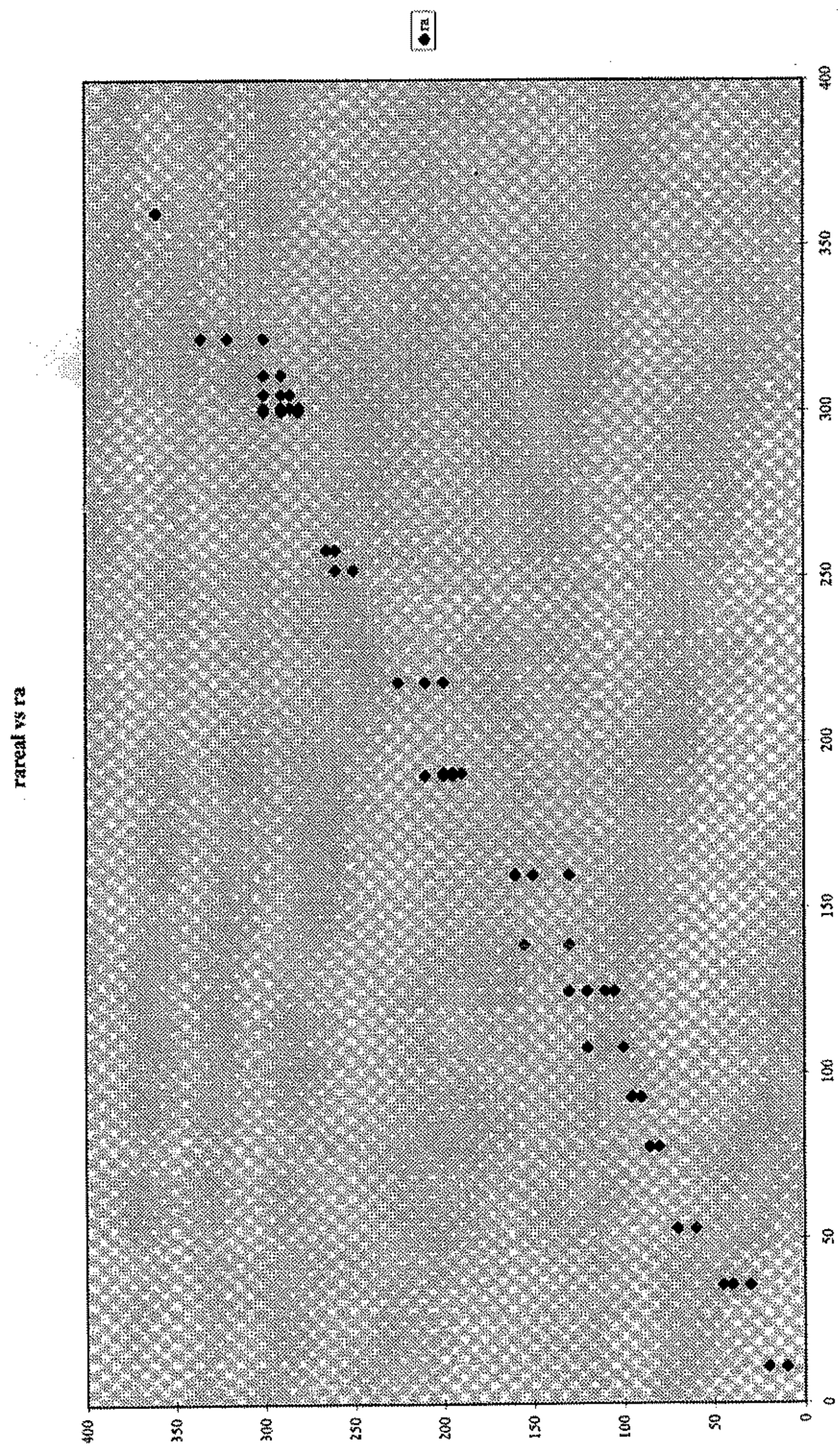

Figure 6.1: Parallel Projection: λ_{ra} vs. θ_{ra}

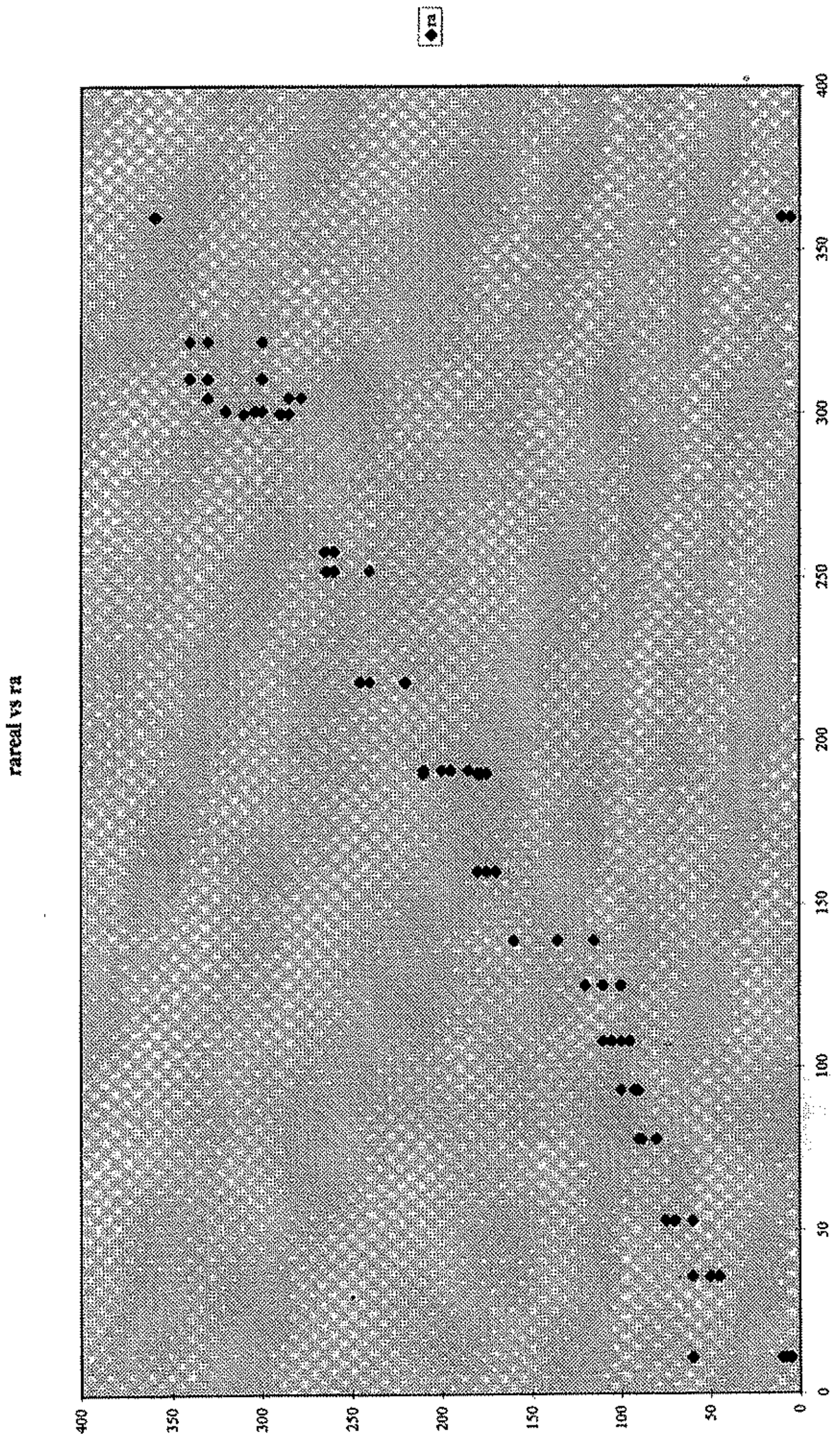

Figure 6.2: Perspective Projection: λ_{ra} vs. θ_{ra}

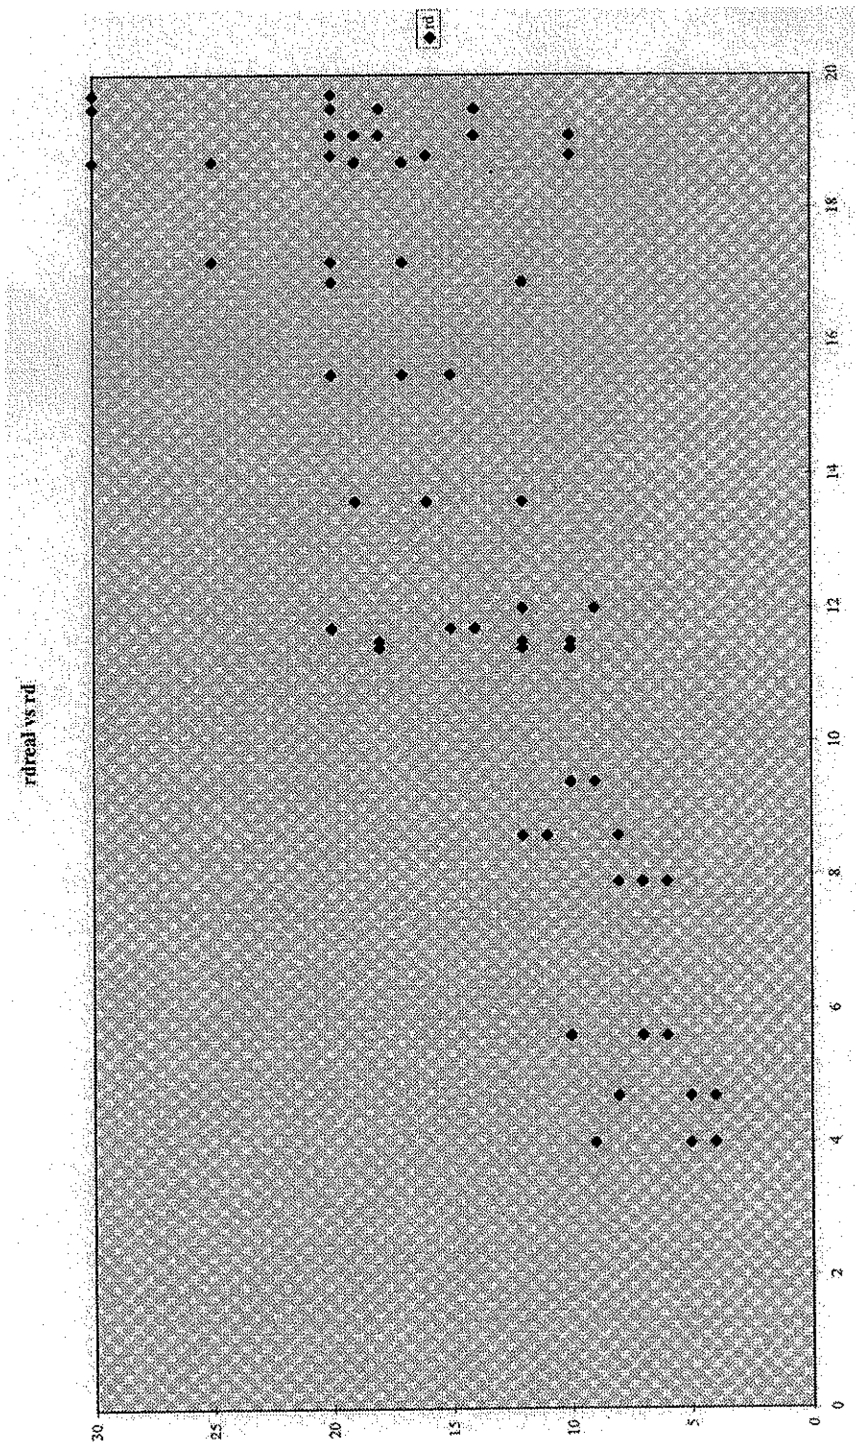

Figure 6.3: Parallel Projection: λ_{rd} vs. θ_{rd}

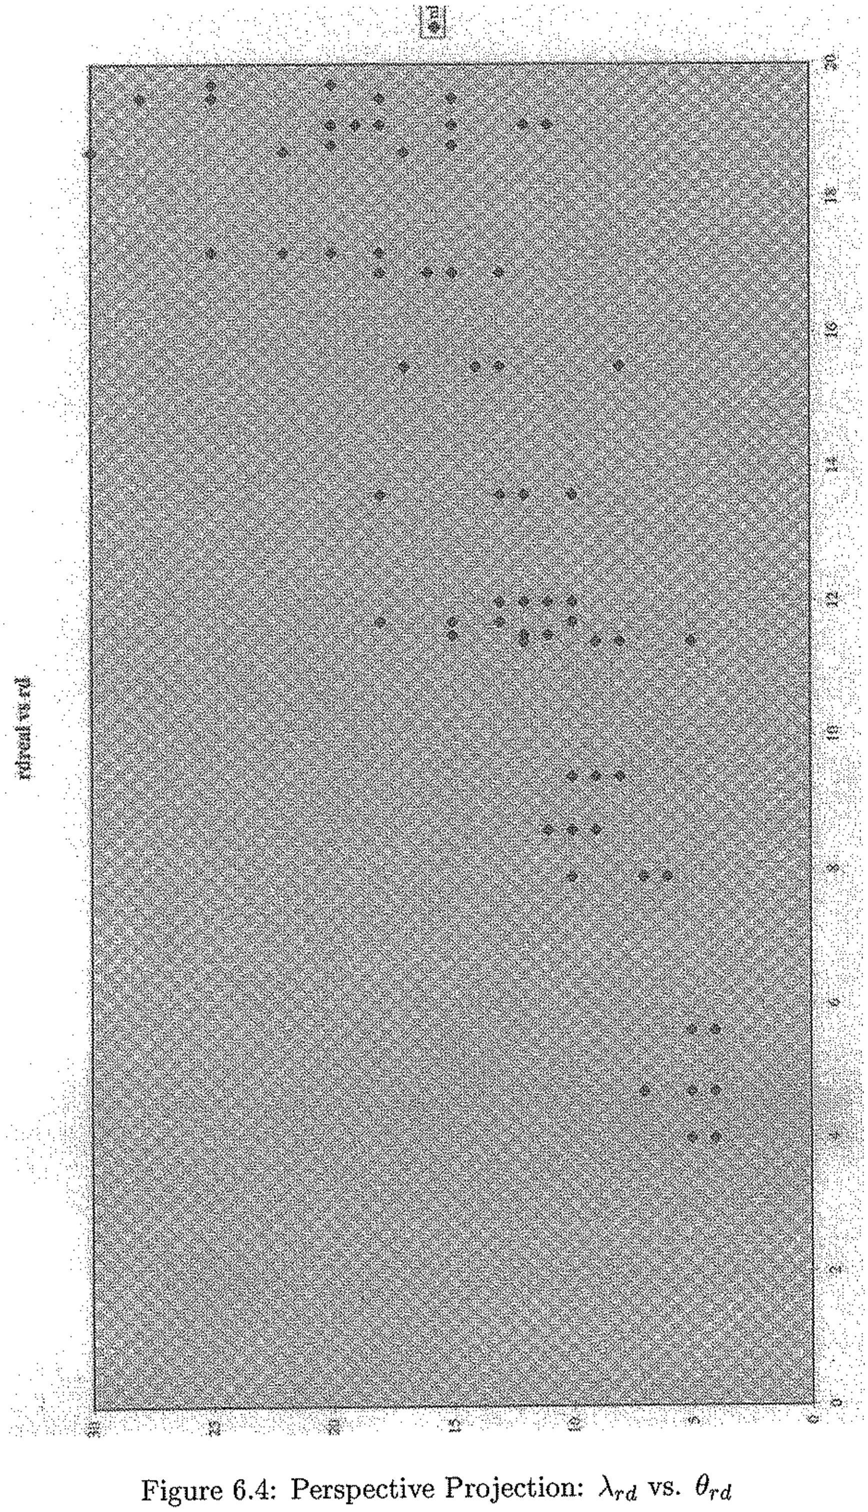

Figure 6.4: Perspective Projection: λ_{rd} vs. θ_{rd}

with $R^2 = 0.21$ and a significance $P = 4 \times 10^{-4}$. In the perspective case, the regression analysis was not significant.

In the case of the absolute error, $|\epsilon_{ra}| = f(x, y, \theta_{ra})$, linear regression analysis showed no correlation for either projection.

These equations suggest that observed azimuth angle and angle error are dependent on the variable x in the parallel projection, but *not* in the perspective projection. In the parallel projection, the contribution to the observed angle appears to grow with increasing x—i.e. when targets are in the left half of the display, there is a tendency to under-estimate the true azimuth, and when the targets are in the right half of the display, there is a tendency to over-estimate the angle. The high value of R^2 in equation (5.1) makes these results particularly believable.

Further, in both parallel and perspective projections, the results suggest that *there is no dependency on y*. Since y corresponds directly to the z-depth of the targets, this was unexpected since it was hypothesised that linear perspective in the perspective projection might influence the reading of angle.

Looking at the error term for parallel projection in eqn. (5.3) there appears to be a positive dependency on x and a slight negative dependency on θ_{ra}.

Distance

Linear regression was not valid for distance analysis since as can be seen in Figures 5.3 and 5.4, variance amongst the λ_{rd} is not constant with increasing θ_{rd}. Therefore, a logarithmic fit was tried. Using the model

$$\ln \lambda_{rd} = f(x, y, \theta_{rd})$$

regression analysis gave:

$$\ln \lambda_{rd}^{1,2} = 1.4 + 0.083\theta_{rd} \qquad (6.4)$$

$$\Rightarrow \quad \lambda_{rd}^{1,2} = 4.06e^{0.083\theta_{rd}} \qquad (6.5)$$

for both parallel and perspective projection, with $R^2 = 0.73$.

The observed distance therefore appears to be independent of projection type, x and y.

6.4 Recall Experiment

6.4.1 Introduction

This was the least effective experiment in the study, but still yielded interesting information. Subjects generally found the task to be very difficult, due to the sheer amount of information to be committed to memory and the short time of presentation. It was felt that the examining the results obtained would not be worthwhile, and so they are not analysed here.

6.4.2 Problems with Procedures

It was intended to use the first scenario for training, and the second as a recorded run, and to present both identically. In practice, however, because of the unfamiliarity with the 3D display format in particular, the first training scenario was presented whilst the supervisor re-read the instruction sheet to the subject and explained the format of the display, and subjects were then allowed to examine the display for unlimited time until they reported ready to continue. Subjects reported using this time to familiarise themselves with the format of the display (particularly in the case of the 3D presentation), and to work out a method of memorising the scenario. When they reported ready, the second scenario was presented for 60s and subjects were asked to reconstruct it.

It was intended that the reconstruction be timed. However, the presence of the supervisor meant that the subjects tended to make comments whilst performing the reconstruction and so the time of finishing was difficult to determine. This should be solved in the main experiments by removing the supervisor for the reconstruction and carrying out the timing by automating the procedures.

Queries raised during the reconstruction soon highlighted the fact that there should have been a full training run prior to the recorded run, with the same procedures as the recorded run, in addition to an unlimited familiarisation period. This would have allowed any questions to be raised during the training run (during which time the supervisor would be on hand) instead of in the recorded run.

The reconstruction task given the 3D stimulus presented problems in that subjects were asked to draw the approximate lengths of the height poles if they could not remember the actual heights, but only one subject did so. This was not made explicit in the instructions, and should perhaps have been clarified by an illustration in the instruction sheet. One subject commented on the potential difficulty of reproducing a display presented in 3D in a plan view. This may have been a factor for more than one subject. In both display formats, subjects were reluctant to guess heights, often preferring to omit these, although whether or not this was because they genuinely couldn't remember the heights even approximately, or was due to some other factor wasn't clear.

One subject did report that the task had a demoralising effect. This may have had an impact on the following task (conflict detection). As related by one subject, ATCOs as a group generally find their job stimulating and take pride in being able to perform a challenging task well. The demoralising effects of this task may therefore have had a more severe psychological impact than may be expected with other groups.

One subject was accidentally shown the stimulus for 2 minutes, and seemed to fare better at the recall. This indicates that the difficulty found in the task was due to the sheer amount of information and/or lack of time, rather than being inherent in the memorisation and recall of a scenario *per se*.

6.4.3 Task Performance

The most interesting aspects of this experiment came as subjects related the ways in which they were performing the task.

It was expected that subjects would remember the traffic position in terms of the overall spatial pattern, and try to recall the heights in the 2D display by reading the datablocks and in 3D display by a combination of memory of the length of the height poles as well as the datablock. In fact, the subjects did not discernibly conform to this behaviour.

In the training scenario, some subjects attempted to devise a strategy for memorising the scenario. One said that she started at the centre and of the display and worked outwards, another that he tried to quarter the display area and remember the traffic in each quarter. In general, however, the method of memorising the traffic did not appear to be visually based in terms of chunks of spatial patterns, as expected, but in terms of air traffic control patterns and approximate levels. Some remembered traffic in terms of it being at high, medium or low level, and tried to classify the aircraft in terms of a familiar traffic behaviour pattern: for example, Heathrow inbounds, overflights, etc., instead of chunking into spatial groups as expected.

6.5 Conflict Detection

6.5.1 Familiarisation

Before presenting the task, subjects were shown an dynamic scenario for familiarisation. This comprised all traffic movements within the displayed area contained in the radar data file, minus general aviation traffic (i.e. all aircraft with squawk code 7000) and non-mode-C transponder-equipped aircraft (for which no height information was available). Additionally, traffic below an altitude of 3 000′ (i.e. a mode-C code of less than 30) was filtered out. Because of the large area displayed (compared to the average sector size) and the busy time of day (around 9:00 a.m.), there was a lot of traffic and some subjects were initially overwhelmed. This may have had a slight negative influence on the subjects (however, it was pointed out that there would be far fewer aircraft in the actual task).

Most subjects examining the 3D display tried to work out a strategy for using this display at this point, and one spent over 10 minutes on the familiarisation.

6.5.2 Task Performance

The supervisor was on hand the whole time, and subjects were encouraged to talk about what they were doing and how they were doing it. Whilst this was very useful, it made it difficult precisely to time when each conflict was identified.

The use of the datablock to present the cleared level and destination code was unfamiliar to subjects, and a lot of conflicts were reported which closer examination of the cleared levels would have revealed were not conflicts at all (mostly aircraft being cleared to a level above that of lower traffic). One controller was confused when the aircraft did not obey the instructions—it was not pointed out sufficiently well that the scenarios were fixed and not interactive. When this was explained to later subjects, this had the side effect that once aircraft were classified as conflicting (or potentially conflicting), they were ignored.

One subject used the SPACE key as a *de*clutter button in the 3D display, in the opposite sense to that intended—i.e. he held it down most of the time to display the full datablocks, and released it where he wanted the display decluttered.

All subjects identified the slowly converging overtaking traffic conflict, but weren't sure about the precise distances involved. They tended to be conservative; e.g. the tracks were converging, so controllers tended to stop the lower aircraft's climb, sort out the lateral separation problem with a course change, then resume climb when lateral separation was seen to be clear.

The crossing conflict was actually missed by some subjects, or identified too late in the case of one subject, especially on the 3D display. Several sources of confusion were cited as reasons:

1. The airspace sector shown was far too large compared to real sectors, and early subjects were tending to "assign" themselves to a particular sector and ignore other traffic! Later subjects were told explicitly to consider all traffic.

2. It was expected that the drop lines would be used as a cue to height and that the 3D display would be sufficiently clear. However, the display tended to appear very cluttered, with inadequate depth cues meaning that separation of a cluster of aircraft along the same line of sight was difficult. Two air plots could appear next together even though their ground plots were well separated due to height differences (i.e. this

was interfering with the "height in the visual field" depth cue), and there were insufficient depth cues to separate the two easily.

- The perspective projection exacerbated this since aircraft near to the horizon tended to be clustered together.
- Perspective projection also caused problems in that fixed-size drop lines vary in displayed length according to its z-depth, negating their value in the opinion of two subjects.
- With two or more aircraft along similar line of sight, it was difficult to determine heights and which air plot corresponded to which ground plot.

3. Correct identification of aircraft was made more difficult by a datablock crossover problem where several air plots were shown in the same region of the display. (See §6.7.)

ATCOs tended to look at lateral and vertical separations separately. From the comments as the task was conducted, some tended to look at the lateral separations of the ground plots, then try to follow the height poles up to the air plots to find the associated labels to read the heights. Clutter and insufficient depth cues made this difficult if several aircraft were along the same line of sight.

6.6 VR Subjective Trial

All bar one subject tried the VR display (and the one who didn't was constrained by time rather than by lack of interest). All enjoyed the experience and found it to be novel and fun. Once navigation was mastered, none reported any problems with using the display or after use, or ill-effects resulting from its use (such as motion sickness or disorientation).

Although the technology was limited (in terms of resolution, display update rate and the cumbersome apparatus) subjects were objective and tended not to be put off by the present limitations, preferring to view the technology as developing and to see whether or not it could be made use of in an ATC context.

Because of the sense of presence an immersive display generates, it was felt that subjects might somehow view the data in a new way, and this was indeed observed in the subjects. Some subjects were also pilots and they commented that the viewpoint was more that of a pilot than an ATCO: such an egocentric display might be useful in cockpits, they suggested, (presumably with a see-through HMD, or a TTW display) for gaining an awareness of the disposition of other traffic (which at the moment has to be done by listening to R/T and is therefore vague), but would not be useful from a controller's point of view with current working practices.

Additional evidence of subjects viewing the data from a new perspective was that some subjects were fascinated watching the traffic patterns: watching inbound aircraft enter a stack and then leave it to head for the extended centreline, turning onto it and landing, or watching outbounds climbing out from an aerodrome, then turning and climbing over a stack or going under it. Two subjects used the display for over 15 minutes and one for over 30 minutes (even though the radar data was only for 15 minutes and so had to be replayed). One commented that he could see the climb angle by the angle of the air trailing histories particularly clearly and so see whether or not a particular aircraft would be able to climb over the stack at present rate of ascent. (The symbols were not sufficiently clear or large to enable him to do this in the TTW display.)

Subjects also tried to place themselves inside aircraft to follow them or to position themselves at an airport to gain a "tower" view of the traffic.

Regarding potential applications, one subject thought that it might be useful in airspace planning. Here, airways are planned and tested by using simulated aircraft with representative performances to fly the routes and procedures. These are displayed in 2D plan and profile views; however, the subject felt that the pilot's perspective as afforded by the VR display might give a better feeling for how well the airway is suited to the simulated flight profiles.

A couple of subjects thought of potential applications in training, for debriefing trainees after exercises.

Two applications were proposed to subjects, who were asked to comment. One was a low cost control tower visual control room simulator. Full-size simulators are constructed for training, with large visual displays providing a panoramic view, but some training could be done using a much cheaper immersive display. The other was a see-through HMD for control tower applications on which taxiways and aircraft could be shown under conditions of restricted visibility superimposed on the outside world. At present, under restricted visibility, ground movement control is done using a surface movements radar, which gives a plan view. A see-through HMD would give an overlay picture instead, and would not be restricted to aircraft on the ground (aircraft could appear as a radar target box, perhaps with mode-A and mode-C information appended).

Most subjects thought both applications feasible and potentially useful.

6.7 Display Evaluation

Because of the unreadiness of the VR display, the comments below apply only to the 2D PPI and 3D TTW displays on the SGI Indy.

6.7.1 Colours

Colours were generally acceptable. Some subjects were familiar with the colour standard; only one subject said that he did not like it, but that was a criticism of the NATS standard rather than the implementation.

6.7.2 Area

The displayed area was generally considered to be far too large compared to the size of sectors which ATCOs usually control. This was no problem in the parameter reading task (except perhaps that the separations were viewed against a much larger area than normal) but interfered with the conflict detection experiment, to the extent that some subjects were ignoring traffic in some parts of the display as they had "assigned" themselves to one particular sector (e.g. Heathrow approach) and were ignoring traffic outside that sector (e.g. overflights or Gatwick traffic) as non-pertinent.

6.7.3 Datablock Format

There was also some unfamiliarity with the datablock in the conflict detection experiment. Since the rationale was to get subjects to extract information purely from the display and to dispense with flight progress strips, the cleared altitude was included in the datablock next to the destination code. The full datablock format used is shown in Figure 4.2. Because the presence of the cleared flight level was unfamiliar to the subjects, it tended to be ignored. The general consensus was summed up by one subject:

> Datablock layout is unfamiliar—and needs time to get used to.

Another subject liked the fact that the cleared level and actual level are not vertically adjacent but offset by the route code, since this helps to distinguish the two and reduce the chance of one being read for the other.

The familiarity problem could be solved by greater exposure and training before the main experimental run.

6.7.4 Legibility of Symbology and Datablocks

Generally, subjects found that symbology viewed in isolation to be legible; however, some commented that the trail dots were rather small in the 2D and 3D TTW versions (being single pixels on a high-resolution display).

Problems were found with the size and legibility of some of the aircraft-related symbology, particularly when viewed against a datablock. One subject commented:

> Trail data are difficult to read when garbled with a datablock.

This was generally supported by the other subjects.

Subjects found the aircraft trailing dots to be generally too small (the single pixel is quite small given the high display resolution) and information generally difficult to see against a datablock, even though aircraft-related symbology is guaranteed not be obscured by datablocks. Geoff Strutt suggested the use of XOR plotting to alleviate this difficulty.

6.7.5 Datablock "Cross-Over"

There was a problem in the datablock overlap avoidance algorithm in that datablocks related to targets drawn close to each other could "cross", so that it could be difficult to associate a datablock with a given target (see Figure 5.5). Again, there were problems when the datablocks actually did

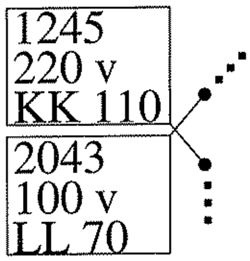

Figure 6.5: Datablock Crossing Problem

overlap, too. Part of the problem is related to the thinness of the leader line (only one pixel) and the subsequent difficulty in seeing it.

7 Conclusions and Recommendations

7.1 Introduction

This pilot study has been extremely revealing in terms of showing the shortcomings of the experimental approached adopted, highlighting areas which should be investigated, and showing how the air traffic controller carries out his/her tasks. The data gathered in the experiments will help to shape the main experimental programme. The conclusions and recommendations for the main study are presented in this chapter.

7.2 Subjective Opinion and Questionnaires

As was observed in the previous chapter, it was originally intended to wait until the end of the whole experiment to debrief subjects to gain subjective opinion data: however, this turned out not to be feasible, so subjects tended to be interviewed whilst performing the tasks (in the case of the conflict detection task and recall tasks) or shortly afterwards.

As this was a pilot study, this was no bad thing, since the aim was to gain as much feedback as possible; however, in the main experiment, a better approach might be to present questionnaires at the *end* of each task to avoid supervisor interference effects. Greater automation of the experimental procedures would also help considerably. For verbal de-briefings (and perhaps whilst the tasks are being carried out), subjects should also be tape-recorded, due to the difficulties of *verbatim* transcription.

In additional to questionnaires or interviews following each task, there should be an entry questionnaire before all tasks and an exit questionnaire/interview.

The entry questionnaire should include any relevant personal data on the subject (for example, age and sex) and should also include operational background (for example, experience with what displays (e.g. no radar, raw MTI plus symbology, fully synthetic colour, any stereo training), what types of control (e.g. GMP/GMC, approach, en-route, military)). Non-operational experience, such as participating in other trials or reviewing new technologies, and whether or not the subjects have piloting experience may also be relevant.

The exit questionnaire should include things like overall impressions of the the display format used.

Post-task questionnaires and interviews should probe the method used to perform the task, as well as overall opinion about the task itself. Self-rating of performance should also be included — the supervisor was interested to note that some subjects thought that they had performed a task badly where they had in fact performed it well, and *vice versa*. Such opinions might be useful in determining whether or not a particular display format gives a *perceived* benefit irrespective of quantitative performance.

7.3 Procedures

Many problems encountered with 3D displays appeared to be either training related (i.e. the training received by ATCOs is based on the use of the 2D PPI radar display, and may not be suitable for working with a three-dimensional presentation) or due to insufficient experience with the 3D display formats before the tasks were carried out.

Further, the pilot subjects were exposed to both 2D and 3D display types during the experiment. In the main experiment, subjects should perhaps be restricted to a single type of display. Subjects may also have to be give more training sessions in order to get used to the 3D presentation and factors like unfamiliar datablock formats.

Non-expert subjects may also present a problem in that they may either require training (which may bias their performance) or perhaps should be given sufficient familiarisation time to develop their own strategies. (This may not be relevant for the parameter reading task, but will be particularly relevant to the conflict detection task, where concepts such as "cleared levels" and destination codes will have to be explained.) The task description sheets will also have to be designed in such a way that they will be equally applicable to expert and novice subjects alike.

7.4 Parameter Reading Task

7.4.1 Ground or Air Plot

Queries from subjects about whether or not to use air or ground plots when estimating azimuth angles and distances raises questions about which plots should be used when estimating distance and angle in general.

If aircraft are not at the same height, then this is likely to make the reading the distance and angle between them using the air plots significantly more difficult and encourage misreadings, since the angle between the air plots on the display will comprise both azimuth and elevation components. In general, subjects perhaps ought to be told to estimate the angle and distance between the ground plots.

However, the perception of elevation and azimuth may be improved by symbolic enhancements or training and is dependent on display parameters as well.

7.4.2 Angle

From the analyses, equations (5.1) and (5.2) in §6.3 show that for both projections, there is a strong correlation between the observed and real azimuth angle. The error was surprisingly low, considering that some subjects reported themselves to be "rusty" and that they were all using an unfamiliar display format. From Table B.3, mean and standard deviations of the angle error term ϵ_{ra} and its absolute value $|\epsilon_{ra}|$ are lower for the parallel than for the perspective projections, and the coefficient of θ_{ra} in eqns. (5.1) and (5.2) is also higher in the parallel projection than in the perspective projection. This suggests that angle accuracy is greater for the parallel projection than for the perspective projection.

The fact that there was no dependency on the y parameter (i.e. z-depth) in either projection was surprising, since it was hypothesised that the effects of linear perspective might influence the reading of azimuth.

7.4.3 Distance

For distance, the same results (equations (5.5) and (5.6)) resulted from both projections. There was *no dependency on the y variable*, and there was *no dependency on projection type*. These were again contrary to expectation, especially given subjective reports.

7.4.4 Implications

One of the purposes of this experiment was to determine whether perspective or parallel projections are more suitable for 3D displays for air traffic control. From this preliminary investigation, it can be seen that contrary to subjective opinion and hypothetical expectation, there is no difference of projection type on the reading of distance, and regarding angle, the projection type does have an influence but the z-depth is not a factor in either case. Overall, parallel projection seems to give better performance in estimation of angle than perspective projection. It would probably be possible to train operators to correct any perceptual errors (a large part of which seems to be a constant bias term in both cases), or to employ symbolic enhancements to reduce errors in both cases (at the expense of greater clutter).

In order to make any recommendations, the experiment should be continued using a larger sample size, and also looking at the effects of stereopsis. In addition, comparison should be made with a 2D display as well, although whether this should be a between-subjects study (i.e. each subject does the task with a 2D PPI and one other type of display) or a within-subjects study (i.e. subjects are allocated to one type of display *only* for this task, one of the display types being the 2D PPI) is not clear at present. Novice subjects should also be used for comparison.

7.5 Recall

The results of this experiment were interesting because, although no good quantitative data were obtained, the experiment showed a lot about the ways in which ATCOs think.

The main problems which subjects found were twofold:

- lack of presentation time
- too much information to recall

With subjects being overloaded (recall that the number of items which may be retained in working memory is 7±2) it was thought that subjects would chunk position in terms of spatial pattern, but from interviewing the subjects, this is not the way in which they memorised the scenario. It would be interesting to compare this behaviour with that of non-expert subjects.

There is evidence that subjects memorise the traffic at least partly in terms of familiar traffic patterns. It is speculated that this behaviour reflects the methods by which ATCOs work; they tend to memorise the traffic as part of the overall mental "picture". They do not try to keep track of the precise location of each aircraft; rather, they know the disposition of each aircraft in terms of its approximate relationship to significant points in the traffic pattern (perhaps temporal as well as spatial) and to other traffic. Where aircraft are separated vertically, the precise lateral position is not important; similarly, where aircraft are separated laterally and are not on converging courses or overtaking, height is not so important. So thinking purely in terms of where an aircraft is at a particular point is may be artificial for ATCOs.

Another possibility for the poor performance of subjects is due to social factors associated with the task of air traffic control; ATCOs take pride in doing what is an important and difficult job well, and so perhaps like to be as precise as possible in these tasks. The instructions were to "make a best effort"—it is possible that this was interpreted as "precision is highly important", and in their attempts to go for precision, chunking was rejected to group items. Consequently, in attempting to remember all items to high precision, a lot of information was lost or remembered incorrectly.

The experiment as a whole was felt by the subjects to be rather artificial *vis-à-vis* the task of air traffic control, since ATCOs do not try to remember precise spatial position, so this raises questions as to the validity of the experiment from this point of view:

- Is the experiment effective—does it actually measure what we want it to measure?
- Is what we're trying to measure really relevant?

These issues require further clarification.

Assuming that the overall objective of this experiment is still valid, there is the question of how to approach it. There are a number of possibilities:

- Retain current experiment, but modify to reduce the number of aircraft or to increase the time.
- Place aircraft in traffic patterns (as they would be found in real life), so that controllers can relate the image to a pattern.
- Show a static image, then show several alternatives and have the subject pick out the correct one.
- Show a dynamic scenario, containing fewer aircraft, for longer period of time (more akin to Burnett & Barfield's experiment) but still no flight progress strips.
- Replicate Burnett & Barfield's experiment, with flight progress strips.

It will be most interesting to compare with non-expert subjects, since these do not have training biases and are perhaps more likely to exhibit the expected behaviour. More general conclusions regarding the effects of display type on memory and recall may be drawn by comparing with non-expert subjects. This is likely to prohibit the use of flight progress strips, however.

7.6 Conflict Detection

The main problems here were lack of familiarity with the 3D display formats, and the fact that

> "The airspace being studied is very large and complex in ATC terms."

Because of the large airspace, some subjects were ignoring some traffic as non-pertinent to the area which they had assigned themselves to control. Also, the animated real traffic sample presented for familiarisation was far too cluttered; a smaller area with its commensurately smaller traffic sample would probably be better for familiarisation in future.

As far as the timing of conflicts is concerned, perhaps a key should be pressed in future when each conflict is detected.

Again, there was evidence that subject behaviour was being influenced by their training, and that subjects were viewing the problem in horizontal and vertical dimensions separately instead of making full use of the three-dimensional integrated visualisation. As regards how the subjects were performing the task using the 3D display, some were trying to work out the lateral geometry using the ground plots, then following the drop lines up to find the associated datablocks to extract height information. If Burnett & Barfields' approach of attaching the datablock to the ground plot were adopted, it is doubtful that the ATCOs would have used the air plots at all, given the high level of clutter in the present display making them difficult to read. Again, comparison with non-expert subjects and interviews of ATCOs to try to elicit the mental representation would be useful.

Some subjects may also have been ignoring the drop lines due to the excessive clutter and insufficient depth cues in order to separate foreground and background targets along the same line of sight. Further, linear perspective rendered the drop lines useless in at least one subject's view.

One subject who had been an RAF photographic interpreter and therefore trained in stereoscopic viewing was afterwards shown the same scene in stereo. He commented that the scene "jumped straight in". In the VR display, subjects also commented that they had fewer problems with sorting foreground from background traffic. This may have been related to the stereopsis and also to the different camera parameters in the VR display. The chosen GFOV for the TTW displays had a slightly 'wide angle' effect and tended to enhance perspective. Perhaps a slight 'telephoto' should be adopted instead. Viewpoint elevation angle may also be a factor—this demonstrates how many design problems accrue to 3D displays!

7.7 VR

Although the VR display could not be evaluated quantitatively, subjective responses allowed the following preliminary conclusions to be drawn.

Regarding air traffic control, the main problem with the VR display appears to be that it presents an egocentric as opposed to exocentric perspective; a pilot's perspective as opposed to a controller's perspective.

With an immersive display, there is also the problem that if the viewer is allowed to place the viewpoint within the dataset, traffic outside the field of view will not be shown. Perhaps the "outside looking in" perspective should be enforced to avoid this problem by constraining movement such that the viewer is not allowed into the airspace, but can moving and rotate the controlled volume instead.

From the point of view of rapid navigation, two features may be desirable:

- A "take me home" facility to return to a predefined viewpoint covering the whole sector, in case of disorientation or the need to see all the traffic quickly.

- Limitations to the viewer's travel to stop the viewpoint from being too far away or "underground".

7.8 Display

The main problems experienced with the display were the datablock crossover problem and legibility of symbology, particularly the trailing histories. This will have to be investigated further before the main study.

Concern was raised over the effects of variation in ambient and display light intensity. The latter could be standardised by measuring from the screen with a light meter and adjusting display brightness until this was uniform. There may be fluctuations in ambient light levels, although natural light appears to be swamped by the artificial lighting. A controlled ambient light environment may be difficult to achieve, however.

7.9 ATC as a Three-Dimensional Problem

From this pilot study into 3D displays for ATC, it was felt that the most important result was the evidence that the way in which ATCOs perform their task may be unsuitable for the efficient use of a 3D display.

The initial thought regarding 3D displays for ATC is that since air traffic control is a problem in three-dimensional space, then there must be an advantage in using a 3D spatial representation. However, this is not necessarily the case. Evidence from the pilot study shows that air traffic controllers think in three dimensions, but not necessarily in the way that is commonly thought — i.e. the mental "picture" is not necessarily a spatial one.

As already discussed, ATC is a problem of flow control. Air traffic controllers appear to think about the relationship of aircraft with respect to position in a traffic pattern, about temporal relationships (e.g. estimated time to waypoint/beacon) and about separation in terms of lateral and vertical dimensions, not regarded as an integrated whole but as separate dimensions. As one controller related, when accepting an aircraft into the sector, one looks at route (is it likely to cross anything?), lateral separation (is it near anything horizontally?) and vertical separation almost separately. If its route is clear of other traffic and nothing else is in the vicinity, then it can be accepted and placed into the "background" with minimum monitoring. If it is separated laterally and lateral separation is not likely to be violated by crossing or overtaking traffic, then there is no need to consider vertical separation. In a crowded airway, where things are likely to be busy, one may opt for adding vertical separation as an extra safeguard. Similarly, if traffic is separated vertically then there is no need to consider lateral separation except where aircraft may be climbing or descending through each others' levels.

This method of thinking appears to be partly based around the use of the PPI radar display, and seemed to occur even with the 3D displays. Some subjects commented that they could extract height information perfectly well from the existing datablocks, and any speed advantage resulting from a direct analogue visualisation would be very minor. In the recall task, a few subjects sorted traffic into height bands on looking at the display. Is it speculated that this is common working practice with a PPI display, since the ATCO is only interested in traffic within his/her sector; any traffic that is too high or too low is simply ignored. As one subject related:

> Difference between 3D air plot and ground plot can cause confusion when judging separation. Most controllers think in terms of level separation for sectorisation and tend to ignore traffic under or over their normal work area. This makes high-level conflict detection difficult for Approach Control and low level conflict detection difficult for Upper Air Controllers.

This is not to say that 3D displays are entirely useless; just that current operational practices are not matched to their use. There may be applications in other areas of ATC. As one subject said:

> An *approach* radar is quite different from en-route radar. It may be worthwhile examining these proposals in relation with the different ATC tasks.

There may be a large element of training. As related by other subject:

> I would need re-training to be at all happy with the 3D presentation. I experienced great difficulty in adapting to the elliptical presentation of range rings. No problems felt with the conflict resolution.

7.10 Parallel or Perspective Projection?

Based on quantitative evidence from the parameter reading task and from qualitative evidence from subjects, it seems that parallel projection is superior to perspective projection for 3D air traffic control displays, for the following reasons:

1. Azimuth angle reading error is lower for the parallel than for the perspective projection.
2. Distance reading is the same for the parallel as the perspective projection.
3. Subjects subjectively preferred the parallel display.

4. Drop lines representing the same height will be displayed as the same length in the parallel projection. (Although the reading of drop line length in perspective and parallel projections was not investigated, but drop lines were ignored by at least one subject in the perspective display for this stated reason.)

5. No "clustering" of aircraft near to the horizon in the perspective display, so less clutter.

The parallel projection should therefore be adopted as the "standard" 3D projection for the main experiments.

8 Glossary

ATC Air Traffic Control A system by which aircraft are controlled by ground-based controllers to ensure their safe and timely transit from point of departure to destination.

ATCO Air Traffic Control Officer A trained and licensed individual who carries out the task of air traffic control.

CAA Civil Aviation Authority The authority responsible for civil aviation in the United Kingdom.

COP Centre of Projection The position of the 'virtual camera' used to project a 3D scene onto a 2D surface.

CSCW Computer-Supported Co-operative Working A system whereby co-operative work of individuals is supported in some way by a computer system, by sharing information between parties or more sophisticated support. ATC is an example of CSCW.

FDP Flight Data Processing The ATC computer system which files flight plans and tracks aircraft through the air traffic control system, providing information to ATCOs in the form of flight progress strips and callsigns in place of mode-A codes on the radar display.

FPS Flight Progress Strip A paper strip used by air traffic controllers on which all pertinent details relating to an aircraft are printed. This is annotated by the air traffic controller in the course of his/her job, so that it always reflects the state of the aircraft.

LCD Liquid Crystal Display These are referred to in two usages here. LCD TVs are used in the virtual reality head-mounted displays as they offer a safe, compact and comparatively low-cost visual display compared to the relatively large, high-voltage, expensive cathode ray tubes normally used in televisions. LCD shutter glasses used to implement a stereoscopic display place a single liquid crystal panels in front of each eye. These are normally transparent; however, when a signal is applied, they become opaque. They can therefore be used to present an alternative image to each eye if synchronised to an image source, to give a stereoscopic display.

MTI Moving Target Indicator A radar which only shows moving targets. Processing is applied to raw radar returns to filter out targets which are not doppler shifted; i.e. that do not have a velocity component in a radius from the radar head. Thus, clouds and terrain may be removed from a radat display.

NATS National Air Traffic Services Part of the UK Civil Aviation Authority, NATS is responsible for airspace planning and overall air traffic control within the UK.

PPI Plan-Position Indicator A two-dimensional display giving a plan-view of an area. Generally, this term refers to a radar-derived plan-view display of position of objects in azimuth and distance from a radar receiver.

RDP Radar Data Processing Signal processing associated with the radar system.

R/T Radio Telephone A half-duplex voice radio link between two or more stations.

SSR Secondary Surveillance Radar A radar, often located beside the primary radar, sends out an interrogation signal. This triggers a transponder on board an aircraft to send back its mode-A identification code and, if an altitude-encoding transponder, its mode-C code which gives its height. Thus, the identification and height of a target may be determined.

TTW Through the Window A display paradigm referring to the viewing of a virtual scene "through the window" of a computer screen. Thw world of the viewer and the virtual environment thus remain separate.

VPN View Plan Normal In a 3D projection, this is vector from the centre of projection normal to the view plane. In the case of the virtual camera not being 'off-axis', it is the line of sight from the virtual camera position into the scene.

VR Virtual Reality A computer-generated environment, perhaps incorporating multi-sensory elements. The term is primarily used to describe environments where the visual element is head-coupled and viewed through a head-mounted display: the effect is to "immerse" the viewer so that (s)he becomes part of (and has a compelling sense of "presence" in) the virtual environment.

A Instruction Sheets

A.1 Introduction Sheet

Air Traffic Control Displays Pilot Study

Thank you for taking the time to participate in this study. Your help is greatly appreciated. The purpose of my research is to evaluate different display technologies for air traffic control. To this end, you will be given three tasks concerned with various perceptual aspects of the job of air traffic control.

- Task 1 will be to read the relative heading and distance between pairs of aircraft on a three-dimensional display.
- Task 2 will be to memorise a static traffic scenario, and then try to recall it.
- Task 3 will be a conflict detection task. A couple of short scenarios will be presented on a display. You will be required simply to identify any potential conflicts (loss of separation) that exist.

After the tasks, you will then be shown some of the different display types and be invited to comment.

Please do not hesitate to ask the supervisor at any time if there is anything that you do not understand, or which is not clear.
Since this pilot study is running throughout this week, in the interests of not prejudicing the research, please do not discuss this with others until after Friday 30 September.

A.2 Parameter Reading Task

Task 1
Heading and Distance Estimation

In this task, you will be asked to estimate the heading and distance between 20 different pairs of aircraft on a three-dimensional display.
The display will show a 100×100 nautical mile region around the Heathrow radar, looking to the north. A video map is displayed on the "ground" which shows the Heathrow, London City and Gatwick CTAs in a light tan colour and the various airspace boundaries and airways. Range rings are also shown at 10nm intervals centered on the radar head at Heathrow.
Aircraft are represented as black "dots" in the air, with a 'drop' line joining their positions with the ground, which is represented by a white "dot".
Each aircraft has a label (datablock) in a box attached to it. The top line of the datablock contains the mode-A identification number of the aircraft. The bottom line shows its mode-C altitude (00s of feet or flight level). *The altitude will be ignored in this task.*
In the experiment, pairs of aircraft will be shown on the screen one after the other. For each pair, please write on the response sheet provided the estimated heading from aircraft 1 to aircraft 2 and the distance between them in nautical miles. Please take as much time as you like, and try to be as accurate as you can. When you have finished with one pair of aircraft, press the SPACE BAR on the keyboard to move to the next pair of aircraft. When you have completed all 20 pairs, please tell the supervisor.

Please tell the supervisor when you are ready to proceed.

A.3 Recall Task

Task 2
Scenario Recall Task

For this task, you will be shown two static displays showing an air traffic scenario. For each display, you will be asked to memorise the air traffic pattern, and then to draw it on a piece of paper. The first display is just for familiarisation. The second will be recorded.

Each display shows a static air traffic scenario over the same area as the previous task. This will be shown for one minute; after this, it will removed from the screen and you will be given a piece of paper showing the radar map of the area. You will then be asked to mark on this piece of paper the positions of the targets and their altitudes from memory. Please try to make a best effort (for example, fill in a best guess for the altitude if you don't remember it precisely), working as quickly and as accurately as possible. Please tell the supervisor as soon as you have finished.

In these displays, some additional information will be given. "History" trails will be shown attached to each aircraft, (if you are shown a 3D display, these will be black for trails in the air, and their "shadows" on the ground will be white) giving its previous positions over the last 8 radar sweeps. A datablock (shown as text characters in a box) attached to each aircraft by a thin leader line shows the mode-A transponder code and mode-C height (00s of feet or flight level), possibly with a small character after it: an up arrow indicating that the aircraft is ascending, a down arrow indicating that the aircraft is descending, or no character to indicate that the aircraft is in level flight.

For this task, I am interested purely in position and height, and the previous positions, identification and vertical trend (climb/descend) can be ignored.

A.4 Conflict Detection Task

Task 3
Conflict Detection Task

The purpose of this task is for you to detect conflicts in an animated traffic scenario.

Familiarisation Display

Prior to the task, you will be shown a short animation of radar data to familiarise you with the display. The radar data are taken from Heathrow radar, at around 09:00 on a morning in April. The aircraft positions are updated every 6 seconds. Don't worry that the display shows a lot of aircraft; you will be shown far fewer in the task!
The display contains a rudimentary datablock overlap avoidance algorithm — that is, the computer will try to move the datablocks around to avoid them overlapping. They may therefore "jump" as the aircraft move, but this is nothing to worry about. The algorithm used is not perfect, however, and so sometimes the datablocks may overlap for short periods.
If you are shown a 2D display, the datablocks will show full information (mode-A code, mode-C height information and climb/descent, as for the previous task). If you are shown a 3D display, since the height is represented graphically, only the mode-A codes will be shown unless the SPACE BAR on the keyboard is pressed — pressing and holding this key for more than 0.5s will cause the full datablock to be shown. Releasing the key will cause only the mode-A code to be shown again.

Conflict Detection Task

In the task itself, you will be shown two short (less than 4 minute) animated traffic scenarios containing several aircraft. You will be asked simply to watch the scenario unfolding and to tell the supervisor when you think that any aircraft may lose separation.
Here, separation is defined as 3nm laterally, or 1 000 feet vertically (i.e. aircraft must be 3nm or greater apart horizontally if they are within 1 000ft of each other vertically, or they must be 1 000ft or greater apart vertically if within 3nm of each other horizontally). When you think two or more aircraft are in danger of coming into conflict with each other, either imminently or at some time in the next few minutes, please tell the supervisor immediately, with the following information:

Conflict Detection Task contd.

- The mode-A transponder codes of the aircraft involved.
- A conflict resolution manœuvre (e.g. turn aircraft 1053 left 20 degrees; descend aircraft 2047 immediately to FL330).

Please note that there may be more than one conflict in the scenario.
The datablocks also contain additional information — an extra line in the datablock shows the cleared altitude (if any) and route code. For example, the datablock:

```
1023
220v
210 LL
```

refers to aircraft mode-A code 1023, at FL220 and descending to its cleared level of FL210, its destination London Heathrow (ICAO code EGLL).

A.5 VR Display Instructions

Virtual Reality Display Instructions for Use

The virtual reality (VR) machine uses a head-mounted display (HMD) to show a three-dimensional scene filling your view which will change as you move your head. The scene is similar to the 3D air traffic displays on the computer, but this time you will be "inside" it as opposed to looking at it "through the window" of a computer screen. Before the task, you will be given a short familiarisation session to get you used to it.

The HMD comprises a helmet containing two small television screens, one for each eye, with some wide-angle optics. Spectacles may be worn with this display. The helmet is connected to the VR machine by a cable at the back. Next to the cable is an power switch, and a nut which tightens or loosens a headband. Before you put the helmet on, please ensure that the nut is unscrewed. Then place the helmet on your head and tighten the nut so that the helmet is comfortable but will not fall off if you lean over.

Take a little time to look around you when you first become "immersed", to get used to the scene changing when you move your head. Try squatting down and tilting your head to the side, and notice the effect. Also, try turning on the spot, and again notice how the scene changes.

You will also be given a hand-held device on which there are several buttons: three on the top (left, centre and right) and two at the front (top and bottom). You will be able to see the position of the device as an arrow if you look at the hand holding it in the virtual world. Try moving your hand about and notice how the arrow changes direction with it.

Moving in the virtual "world" can be accomplished in two ways — you can either step in any direction or you can "fly". Walking anywhere is rather restrictive because of the cable. Flying is therefore the preferred method of moving.

Flying is accomplished with the buttons on the top of the hand-held device. Pressing the left-hand button on the top of the device moves you forward in the direction in which the arrow (i.e. your hand) is pointing. Pressing the right-hand button moves you backwards in this direction. Notice that you can look sideways whilst travelling—just turn your head in any direction whilst keeping the arrow pointing in the desired direction of travel. As an exercise, without taking a step, try to fly to the north of the displayed virtual area.

When you feel that you are familiar with the virtual environment, please tell the supervisor that you are ready to proceed.

B Parameter Reading Experiment Tables

B.1 Stimulus Values

Table B.1 below gives the values of the variables in the 20 stimuli used in the experiment. Symbols are as in §3.2.3.

Stimulus	θ_{RA} (deg)	θ_{RD} (nm)	x (nm)	y (nm)	$T1_x$ (m)	$T1_y$ (m)	$T2_x$ (m)	$T2_y$ (m)
1	139	13.6	3.8	-37.4	-1225	-59799	15310	-78820
2	301	18.7	-22.6	-43.7	-27030	-89909	-56735	-72061
3	258	11.7	-3.8	13.2	3562	26716	-17646	22208
4	53	19.1	-22.2	-20.2	-55275	-48086	-27007	-26784
5	190	19.1	41.5	5.6	79981	27807	73835	-7051
6	160	11.4	-35.7	41.5	-69772	86834	-62546	66982
7	36	4.7	-34.6	4.9	-66681	5557	-61561	12604
8	11	12.0	-3.5	29.2	-8608	43199	-4365	65028
9	305	8.6	19.4	-25.0	42480	-50901	29424	-41759
10	218	4.0	-3.5	-11.0	-4204	-17465	-8768	-23306
11	322	16.9	-1.2	-36.4	7417	-79796	-11865	-55117
12	300	11.5	-19.0	-16.1	-25983	-35164	-44439	-24509
13	78	18.8	13.6	-2.6	8164	-8440	42243	-1196
14	108	19.5	11.1	25.5	3386	52840	37755	41673
15	191	7.9	-35.4	-28.8	-64207	-46187	-67000	-60558
16	360	15.5	-1.3	-9.0	-2409	-31041	-2409	-2317
17	93	17.2	3.8	15.3	-8874	29188	22958	27520
18	125	9.4	-17.9	-7.7	-40307	-9274	-26037	-19266
19	311	5.6	-42.4	-22.9	-74660	-45843	-82492	-39034
20	252	19.7	5.6	31.2	27739	63461	-6983	52179

Table B.1: Values of Parameters used in the Experiment

B.2 Result Data

These tables show results of the parameter reading experiments. Tables B.2–B.21 give subject readings and errors according to stimulus. In these tables, ϵ_{ra} is taken modulo 360° to keep it in the range $-180° \leq \epsilon_{ra} \leq 180°$

(i.e. 359°≡ −1°). Figures 5.1–5.4 show scatter plots of the observed versus actual angles and distances for parallel and perspective projections.

In these tables, projection 1 is parallel, projection 2 is perspective.

Stim	Subject	Projection	λ_{ra}	ϵ_{ra}	λ_{rd}	ϵ_{rd}
1	1	2	160	21	12	-1.6
1	2	1	155	16	12	-1.6
1	3	1	130	-9	16	2.4
1	4	2	115	-24	18	4.4
1	5	1	130	-9	19	5.4
1	6	1	130	-9	12	-1.6
1	7	2	160	21	10	-3.6
1	8	2	135	-4	13	-0.6

Table B.2: Stimulus 1: $\theta_{ra} = 139°$, $\theta_{rd} = 13.6$nm

Stim	Subject	Projection	λ_{ra}	ϵ_{ra}	λ_{rd}	ϵ_{rd}
2	1	2	300	-1	22	3.3
2	2	1	300	-1	17	-1.7
2	3	1	280	-21	25	6.3
2	4	2	304	3	22	3.3
2	5	1	290	-11	19	0.3
2	6	1	285	-16	30	11.3
2	7	2	320	19	30	11.3
2	8	2	300	-1	17	-1.7

Table B.3: Stimulus 2: $\theta_{ra} = 301°$, $\theta_{rd} = 18.7$nm

Stim	Subject	Projection	λ_{ra}	ϵ_{ra}	λ_{rd}	ϵ_{rd}
3	1	2	265	7	13	1.3
3	2	1	265	7	14	2.3
3	3	1	265	7	14	2.3
3	4	2	265	7	10	-1.7
3	5	1	260	2	15	3.3
3	6	1	265	7	20	8.3
3	7	2	260	2	15	3.3
3	8	2	265	7	18	6.3

Table B.4: Stimulus 3: $\theta_{ra} = 258°$, $\theta_{rd} = 11.7$nm

Stim	Subject	Projection	λ_{ra}	ϵ_{ra}	λ_{rd}	ϵ_{rd}
4	1	2	070	17	19	-0.1
4	2	1	060	7	20	0.9
4	3	1	070	17	19	-0.1
4	4	2	075	22	19	-0.1
4	5	1	070	17	20	0.9
4	6	1	070	17	19	-0.1
4	7	2	070	17	18	-1.1
4	8	2	060	7	20	0.9

Table B.5: Stimulus 4: $\theta_{ra} = 053°$, $\theta_{rd} = 19.1$nm

Stim	Subject	Projection	λ_{ra}	ϵ_{ra}	λ_{rd}	ϵ_{rd}
5	1	2	175	-15	12	-7.1
5	2	1	195	5	18	-1.1
5	3	1	200	10	18	-1.1
5	4	2	178	-12	11	-8.1
5	5	1	210	20	14	-5.1
5	6	1	210	20	10	-9.1
5	7	2	210	20	15	-4.1
5	8	2	180	-10	11	-8.1

Table B.6: Stimulus 5: $\theta_{ra} = 190°$, $\theta_{rd} = 19.1$nm

Stim	Subject	Projection	λ_{ra}	ϵ_{ra}	λ_{rd}	ϵ_{rd}
6	1	2	170	10	5	-6.4
6	2	1	160	0	10	-1.4
6	3	1	130	-30	12	0.6
6	4	2	175	15	8	-3.4
6	5	1	150	-10	18	6.6
6	6	1	160	0	10	-1.4
6	7	2	180	20	12	0.6
6	8	2	170	10	9	-2.4

Table B.7: Stimulus 6: $\theta_{ra} = 160°$, $\theta_{rd} = 11.4$nm

Stim	Subject	Projection	λ_{ra}	ϵ_{ra}	λ_{rd}	ϵ_{rd}
7	1	2	050	14	5	0.3
7	2	1	030	-6	5	0.3
7	3	1	030	-6	5	0.3
7	4	2	045	9	4	-0.7
7	5	1	045	9	8	3.3
7	6	1	040	4	4	-0.7
7	7	2	045	9	7	2.3
7	8	2	060	24	5	0.3

Table B.8: Stimulus 7: $\theta_{ra} = 036°$, $\theta_{rd} = 4.7$nm

Stim	Subject	Projection	λ_{ra}	ϵ_{ra}	λ_{rd}	ϵ_{rd}
8	1	2	005	-6	12	0.0
8	2	1	010	-1	12	0.0
8	3	1	010	-1	9	-3.0
8	4	2	008	-3	10	-2.0
8	5	1	020	9	12	0.0
8	6	1	020	9	12	0.0
8	7	2	060	49	13	1.0
8	8	2	010	-1	11	-1.0

Table B.9: Stimulus 8: $\theta_{ra} = 011°$, $\theta_{rd} = 12.0$nm

Stim	Subject	Projection	λ_{ra}	ϵ_{ra}	λ_{rd}	ϵ_{rd}
9	1	2	285	-20	11	-2.4
9	2	1	300	-5	8	-0.6
9	3	1	285	-20	12	3.4
9	4	2	278	-27	11	2.4
9	5	1	290	-15	11	2.4
9	6	1	285	-20	12	3.4
9	7	2	330	25	10	1.4
9	8	2	285	-20	9	0.4

Table B.10: Stimulus 9: $\theta_{ra} = 305°$, $\theta_{rd} = 8.6$nm

Stim	Subject	Projection	λ_{ra}	ϵ_{ra}	λ_{rd}	ϵ_{rd}
10	1	2	245	27	4	0.0
10	2	1	210	-8	5	1.0
10	3	1	200	-18	4	0.0
10	4	2	245	27	5	1.0
10	5	1	225	7	9	5.0
10	6	1	225	7	4	0.0
10	7	2	220	2	5	1.0
10	8	2	240	22	5	1.0

Table B.11: Stimulus 10: $\theta_{ra} = 218°$, $\theta_{rd} = 4.0$nm

Stim	Subject	Projection	λ_{ra}	ϵ_{ra}	λ_{rd}	ϵ_{rd}
11	1	2	340	18	13	-3.9
11	2	1	335	13	12	-4.9
11	3	1	300	-22	20	3.1
11	4	2	300	-22	16	-0.9
11	5	1	300	-22	20	3.1
11	6	1	320	-2	20	3.1
11	7	2	330	8	18	1.1
11	8	2	330	8	15	-1.9

Table B.12: Stimulus 11: $\theta_{ra} = 322°$, $\theta_{rd} = 16.9$nm

Stim	Subject	Projection	λ_{ra}	ϵ_{ra}	λ_{rd}	ϵ_{rd}
12	2	1	300	0	12	0.5
12	3	1	290	-10	18	6.5
12	4	2	285	-15	11	-0.5
12	5	1	280	-20	12	0.5
12	6	1	280	-20	10	-1.5
12	7	2	310	10	15	3.5
12	8	2	290	-10	12	0.5

Table B.13: Stimulus 12: $\theta_{ra} = 300°$, $\theta_{rd} = 11.5$nm

Stim	Subject	Projection	λ_{ra}	ϵ_{ra}	λ_{rd}	ϵ_{rd}
13	1	2	080	2	20	1.2
13	2	1	080	2	16	-2.8
13	3	1	080	2	10	-8.8
13	4	2	088	10	15	-3.8
13	5	1	080	2	16	-2.8
13	6	1	085	7	20	1.2
13	7	2	090	12	20	1.2
13	8	2	080	2	15	-3.8

Table B.14: Stimulus 13: $\theta_{ra} = 078°$, $\theta_{rd} = 18.8$nm

Stim	Subject	Projection	λ_{ra}	ϵ_{ra}	λ_{rd}	ϵ_{rd}
14	1	2	095	-13	15	-4.5
14	2	1	120	12	18	-1.5
14	3	1	100	-8	30	10.5
14	4	2	105	-3	18	-1.5
14	5	1	100	-8	14	-5.5
14	6	1	100	-8	20	0.5
14	7	2	110	2	28	8.5
14	8	2	100	-8	25	5.5

Table B.15: Stimulus 14: $\theta_{ra} = 108°$, $\theta_{rd} = 19.5$nm

Stim	Subject	Projection	λ_{ra}	ϵ_{ra}	λ_{rd}	ϵ_{rd}
15	1	2	195	4	6	-1.9
15	2	1	190	-1	8	0.1
15	3	1	195	4	7	-0.9
15	4	2	185	-6	6	-1.9
15	5	1	200	9	7	-0.9
15	6	1	190	-1	6	-1.9
15	7	2	210	19	10	2.1
15	8	2	200	9	7	-0.9

Table B.16: Stimulus 15: $\theta_{ra} = 191°$, $\theta_{rd} = 7.9$nm

Stim	Subject	Projection	λ_{ra}	ϵ_{ra}	λ_{rd}	ϵ_{rd}
16	1	2	005	5	17	1.5
16	2	1	360	0	15	0.5
16	3	1	360	0	17	1.5
16	4	2	359	-1	14	-1.5
16	5	1	360	0	20	4.5
16	6	1	360	0	17	1.5
16	7	2	010	10	13	-2.5
16	8	2	360	0	8	-7.5

Table B.17: Stimulus 16: $\theta_{ra} = 360°$, $\theta_{rd} = 15.5$nm

Stim	Subject	Projection	λ_{ra}	ϵ_{ra}	λ_{rd}	ϵ_{rd}
17	1	2	090	-3	18	0.8
17	2	1	090	-3	20	2.8
17	3	1	095	2	25	7.8
17	4	2	092	-1	18	0.8
17	5	2	090	-3	22	4.8
17	6	1	095	2	17	-0.2
17	7	2	100	7	25	7.8
17	8	2	090	-3	20	2.8

Table B.18: Stimulus 17: $\theta_{ra} = 093°$, $\theta_{rd} = 17.2$nm

Stim	Subject	Projection	λ_{ra}	ϵ_{ra}	λ_{rd}	ϵ_{rd}
18	1	2	120	-5	10	0.6
18	2	1	130	5	10	0.6
18	3	1	105	-20	10	0.6
18	4	2	100	-25	8	-1.4
18	5	1	110	-15	9	-0.4
18	6	1	120	-5	10	0.6
18	7	2	120	-5	8	-1.4
18	8	2	110	-15	9	-0.4

Table B.19: Stimulus 18: $\theta_{ra} = 125°$, $\theta_{rd} = 9.4$nm

Stim	Subject	Projection	λ_{ra}	ϵ_{ra}	λ_{rd}	ϵ_{rd}
19	1	2	300	-11	4	-1.6
19	2	1	300	-11	6	0.4
19	3	1	290	-21	7	1.4
19	4	2	300	-11	4	-1.6
19	5	1	300	-11	10	4.4
19	6	1	290	-21	7	1.4
19	7	2	340	29	5	-0.6
19	8	2	330	19	5	-0.6

Table B.20: Stimulus 19: $\theta_{ra} = 311°$, $\theta_{rd} = 5.6$nm

Stim	Subject	Projection	λ_{ra}	ϵ_{ra}	λ_{rd}	ϵ_{rd}
20	1	2	260	8	20	0.3
20	2	1	260	8	20	0.3
20	3	1	260	8	30	10.3
20	4	2	264	12	20	0.3
20	5	1	250	-2	20	0.3
20	6	1	250	-2	20	0.3
20	7	2	240	-12	25	5.3
20	8	2	260	8	20	0.3

Table B.21: Stimulus 20: $\theta_{ra} = 252°$, $\theta_{rd} = 19.7$nm

B.3 Statistical and Regression Analysis

The following tables give statistical and regression analyses for the above data.

	Parallel	Perspective
max $\|\epsilon_{ra}\|$	30	49
min $\|\epsilon_{ra}\|$	0	0
mean ϵ_{ra}	-2.23	3.99
standard dev ϵ_{ra}	11.48	14.38
mean $\|\epsilon_{ra}\|$	9.14	11.89
standard dev $\|\epsilon_{ra}\|$	7.22	8.93

Table B.22: Angle Statistical Analysis

	Parallel	Perspective
max $\|\epsilon_{rd}\|$	11.3	11.3
min $\|\epsilon_{rd}\|$	0.0	0.0
mean ϵ_{rd}	0.97	-0.02
standard dev ϵ_{rd}	3.63	3.43
mean $\|\epsilon_{rd}\|$	2.52	2.44
standard dev $\|\epsilon_{rd}\|$	2.78	2.40

Table B.23: Distance Statistical Analysis

	Parallel			Perspective		
	Coeff	Std Err	tStat	Coeff	Std Err	tStat
Intercept	6.43	2.61	2.47	26.01	13.81	1.88
x	0.14	0.06	2.52	-0.25	0.30	-0.82
y	0.02	0.05	0.36	-0.19	0.29	-0.67
θ_{ra}	0.96	0.01	77.94	0.82	0.07	12.38
R^2	0.99			0.72		
Significance F	8×10^{-75}			4×10^{-21}		

Table B.24: Angle Linear Regression: $\lambda_{ra} = f(x, y, \theta_{ra})$

	Parallel			Perspective		
	Coeff	Std Err	tStat	Coeff	Std Err	tStat
Intercept	6.43	2.61	2.47	5.72	3.46	1.65
x	0.14	0.06	2.52	-0.18	0.08	-2.38
y	0.02	0.05	0.36	0.02	0.07	0.35
θ_{ra}	-0.04	0.01	-3.18	-0.02	0.02	-0.92
R^2	0.21			0.09		
Significance F	4×10^{-4}			0.08		

Table B.25: Angle Error Linear Regression: $\epsilon_{ra} = f(x, y, \theta_{ra})$

Bibliography

[BB91] Meridyth Svensa Burnett and Woodrow Barfield. Perspective versus plan view air traffic control displays: Survey and empirical results. In *Proc. Human Factors Soc. 35th Annual Meeting*, pages 87–91, 1991.

[BLR90] Woodrow Barfield, Rafael Lim, and Craig Rosenberg. Visual enhancements and geometric field of view as factors in the design of a three-dimensional perspective display. In *Proc. Human Factors Soc. 34th Annual Meeting*, pages 1470–1473, 1990.

[Bro94a] Mark A. Brown. Design specifications for air traffic control displays: 2D, 3D and VR. Computer Science Technical Report 676, QMW College, University of London, May 1994.

[Bro94b] Mark A. Brown. Displays for air traffic control: 2D, 3D and VR—a preliminary investigation. Computer Science Technical Report 675, QMW College, University of London, May 1994.

[Bro94c] Mark A. Brown. On the evaluation of 3D display technologies for air traffic control. Computer Science Technical Report 682, QMW College, University of London, August 1994.

[Duk92] Graham Duke. *Air Traffic Control.* Ian Allan, 4th edition, 1992. ISBN 1-7110-2008-6.

[EMH84] Stephen R. Ellis, Michael W. McGreevy, and Robert J. Hitchcock. Influence of a perspective cockpit traffic display format on pilot avoidance maneuvers. In *AGARD Conference Proceedings*, pages 16–1–16–9, 1984.

[ETGS91] Stephen R. Ellis, Gregory K. Tharp, Arthur J. Grunwald, and Stephen Smith. Exocentric judgements in real environments and stereoscopic displays. In *Proc. Human Factors Soc. 35th Annual Meeting*, pages 1442–1473, 1991.

[Gra89] David Graves. *United Kingdom Air Traffic Control.* Airlife Publishing Ltd., 1989. ISBN 1-85310-107-9.

[HH93] R. H. R. Harper and John A. Hughes. What a f-ing system! send 'em all to the same place and then expect us to stop them hitting: Making technology work in air traffic control. In Graham Button, editor, *Technology in Working Order: Studies of work, interaction, and technology*, chapter 7, pages 127–144. Routlege, 1993.

[HHS91] R. R. Harper, J. A. Hughes, and D. Z. Shapiro. Harmonious working and CSCW: Computer technology and air traffic control. In J. M. Bowers and S. D. Benford, editors, *Studies in Computer Supported Cooperative Working*, pages 225–234. Elsevier Science, 1991.

[Lea94] David Learmount. Reading the way: what are the human-factors issues in air-traffic control by datalink? *Flight International*, 146(4447):26–27, 16-22 November 1994.

[Nol90] Michael S. Nolan. *Fundamentals of Air Traffic Control.* Wadsworth Publishing Company, 1990. ISBN 0-534-12246-9.

[RM92] L. Reynolds and C. Metcalfe. The interim NATS standard for the use of colour on air traffic control displays. CS Report 9213, Issue 1.2, UK CAA, Chief Scientist's Division, National Air Traffic Services, 1992.

[Str91] Geoffrey G. Strutt. The human computer interface in air traffic control: Feasibilty study for a three dimensional radar picture. Master's thesis, Middlesex Polytechnic, London, UK, 1991.

[WF88] Christopher D. Wickens and John M. Flach. Information processing. In Earl L. Wiener and David C. Nagel, editors, *Human Factors in Aviation*, chapter 5, pages 111–155. Academic Press, Inc., 1988.

[WTS89] Christopher D. Wickens, Steven Todd, and Karen Seidler. Three-dimensional displays: Perception, implementation and applications. Technical report, Aviation Research Laboratory, Institute of Aviation, University of Illinois at Urbana-Champaign, 1989. Subcontract for AAMRL/HEL/CSERIAC Wright-Patterson Air Force Base, Contract DLA900-88-D-0393.