Very Large Format Stereoscopic Head-Up Display for the Airport Tower

Stephen Peterson
EUROCONTROL Experimental Centre
Brétigny-sur-Orge, France
stephen.peterson@eurocontrol.int

Magnus Axholt
EUROCONTROL Experimental Centre
Brétigny-sur-Orge, France
magnus.axholt@eurocontrol.int

Stephen R. Ellis
NASA Ames Research Center
Moffett Field, California, United States
sellis@mail.arc.nasa.gov

***Abstract*— Head-up displays typically used in aircraft and automobiles have limited fields of view, generally less than 30°. Such systems can provide large, essentially unlimited fields of regard if they are head-mounted. A new alternative for provision of large fields of regard without requiring a head mount is a large format holographic optical element that can serve as a see-though polarization preserving diffusion screen. It may be adaptable to very large formats (> 2 m). The see-through displays it can support will not provide accommodative relief but will avoid cumbersome head-mounted optics. Some of the psychophysical aspects of this display technology, e.g., luminance, distortion, visual resolution loss, and depth rendering biases have been investigated and are reported below as part of a project to design a practical see-though display that may be used in an airport tower.**

I. Introduction

Head-up displays (HUDs) may prove a valuable addition to or substitute for current information systems in the control tower. The typical HUD as found in cockpits is a see-through display that superimposes information about the aircraft and flight status onto the forward out-of-the-window view, minimizing the number of head-down movements needed to obtain information. Some displayed elements of aircraft HUDs are conformal imagery, meaning that the superimposed symbols appear attached to real objects in the window view. See-through display systems have been proposed for a range of other application domains, emerging into the field of Augmented Reality (AR). An AR display and an aircraft HUD are similar in many ways, but usually differ in the fact that the information presented to the user of an AR display system is adjusted to the user's personal perspective rather than that of the aircraft. User motion tracking systems feed the AR display system with data that allows for spatially registered (i.e. conformal) graphics regardless of user position and orientation. Most AR displays are head-mounted (HMD), just like a helmet-mounted aircraft HUD.

We propose the use of another AR display format, which we refer to as a spatial AR display. It uses a polarization-preserving transparent projection screen. This format is more similar to the panel-mounted aircraft HUD in that the image generating source is not placed on the user's head but rigidly mounted at a distance in front of the user. Because of its potentially large form factor, it can provide, however, a considerably larger field of view as compared to a typical aircraft HUD.

The idea of HUDs in the airport tower is not novel in itself. The first known reference to the concept was made by Hitchcock et al. in the late 1980's at FAA [1]. Potential benefits of such a display system include better display integration/placement, improved low visibility operations, reduced controller memory load, and a kind of x-ray vision in which controllers can see through occluding structures [2], [3]. Human performance experiments have been performed with transparent projection screens for various visual search tasks both with monoscopic [4] and stereoscopic [5] display conditions. The effects of limited field of view in AR display systems have been widely researched. Ellis et al. [6] provide a review of relevant literature related to field of view constraints.

Fig. 1. The transparent projection screen demonstrating a simple overlay (left) and during the refraction experiment (right).

Our proposed transparent projection screen (Figure 1) consists of two sheets of glass enclosing a Holographic Optical Element (HOE), through which projected light is directed towards the user principally through diffractive effects. As the HOE is polarization-preserving, passive stereo techniques can be used to produce artificial depth in the rendered images. The only equipment needed by the user for a stereoscopic display is a pair of light-weight, polarizing, glasses. This is substantially less intrusive than a traditional HMD where the user is encumbered by signal and power cables and device weights commonly over 1 kg.

Another advantage of a rigid, externally mounted display is that the position of the screen can be precisely measured. It is therefore possible to perform calibration and registration more accurately than in the case of HMD since there is no unpredictable or difficult to measure equipment slippage usually associated with HMDs.

Moreover, theoretically, multiple viewers can be accommodated in a projection screen if they are individually tracked and if multiplexing techniques are implemented for separating the distinct display viewpoints.

However, there are some design issues with this type of system that need to be addressed. (a) The screen is sensitive to high ambient light as the HOE refracts, diffracts and diffuses some light coming from other directions than the projectors, decreasing overall contrast. (b) The HOE also shows slight inhomogeneity and refraction effects at certain positions and angles. (c) The reduced contrast due to optical imperfections could affect human visual acuity. (d) Depth rendering using the screen may show some biases that would interfere with operational use.

This paper reports on the results of preliminary investigations of the impact of the enumerated issues. Issues (a) and (b) are addressed in the following Screen Properties section. Issues (c) and (d) are addressed in the Human Performance section.

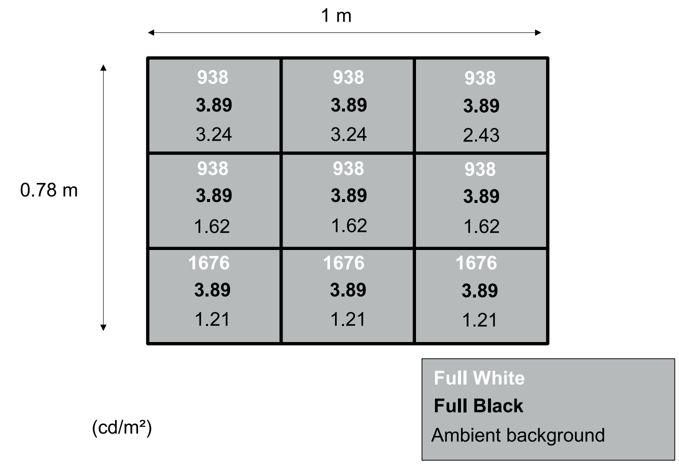

Fig. 2. Screen luminance values, measured at a perpendicular angle

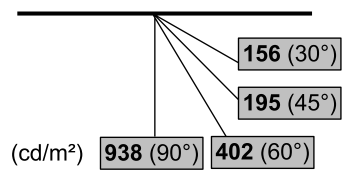

Fig. 3. Screen luminance values, measured at various azimuthal angles

II. Screen Properties

A. Luminance and Contrast

The luminance levels were measured using the light sensor of a Canon 350D digital SLR camera. The method of luminance measurement was validated against a known luminance on a laptop and determined to be accurate to approximately 13%. See Appendix for details on luminance and contrast calculations.

1) Setup: The screen was located in a darkened room 2 m from a pair of projectors mounted below at an angle of ~38°. The projectors have a maximum luminance of 6500 ANSI lumens, XGA resolution and a 1300:1 contrast ratio. The projector images were adjusted and keystoned to fill the entire screen. The camera was mounted on a tripod 50 cm from the screen on the opposite side. Luminance values were measured in three different projector conditions, (a) projecting a white image, (b) projecting a black image, and (c) projectors switched off to measure ambient conditions. In each condition nine measurements were made in a 3x3 matrix covering the entire screen. The optical axis of the camera was always perpendicular to the screen, only the tripod position and height was adjusted. In condition (a) additional measurements were made from varying azimuthal angles to the screen, with the camera always pointing towards the center of the screen at a distance of 76 cm. For each measurement the camera's recommended aperture and exposure time were recorded for calculating luminance values.

2) Results: The results of the luminance measurements are shown in Figure 2. The numbers in white font show the luminance values when projecting a full white image, numbers in black and bold font when projecting a black image, numbers in thin black font when projectors are turned off. Figure 3 shows the measurements when projecting a white image but varying the azimuthal angle from a 90° normal to the screen.

The contrast values were calculated using the full white and full black values subtracted by the ambient light values from Figure 2. Figure 4 shows the contrast values calculated with the Michelson formula.

3) Discussion: The luminance values when projecting a white image are higher in the lower three measurements (Figure 2). Since the projectors are mounted below, the throw distance is shorter and the incident angle is higher in the lower part of the screen, which are likely causes of the higher luminance values.

The ambient luminance values increase in the upper measurements. This is likely due to the fact that the room (basement) has windows close to the ceiling. They were covered with blinds, but the blinds did not block 100% of the outside

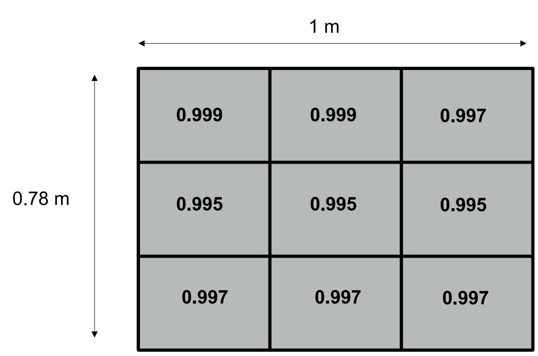

Fig. 4. Contrast values (Michelson)

light. As the luminance was measured in the upper parts of the screen, more of the windows were visible in the camera frame, resulting in the higher luminance values.

When using this screen technology in the normal daylight conditions of a control tower, the ambient light could approach 30,000 cd/m^2, and contrast will thus be reduced. Contrast levels with the system we used would be approximately 1/3 of that needed when comparing to aircraft HUDs, but this could be easily compensated for with tinted window films that reduce the amount of incoming daylight or through the use of brighter projectors.

B. Inhomogeneity

As previously described the screen consists of two sheets of glass and a HOE. The HOE is constructed to affect the light in ways similar to a prism. The HOE is made of dichromated gelatin (DCG). The purpose of the DCG is to modulate the refractive index so that the HOE diffracts incoming light into wave fronts which by constructive interference produce parallel rays so that the HOE in effect refracts the rays without spectrally separating the light. There are at least two theories on how the DCG simulates refraction. A complete description is still missing and much of the production process is adjusted by trial and error [7].

Since DCG is sensitive to humidity and changes in temperature, we are curious to see how homogeneous the DCG is across the HOE. Distortions in the lens and screen of display equipment are significant sources of registration error [8] and inhomogeneities in the HOE would disturb spatial registration used in an AR system.

1) Setup: To investigate the DCG we mounted a laser pointer on a stand and studied the projected laser point on a sheet of paper on a perpendicular wall at 27.75 m distance. We used the sheet of paper to sketch the outline of the projected laser point. First we sketched the projected laser point when the laser did not intersect the screen. This registration was used as a reference point. Then the screen was moved laterally into the laser beam. The screen was successively moved so that the laser beam would pass through the screen at 10 discreet points approximately 10 cm apart along a horizontal line over the screen. At each point the outline of the projected laser point was sketched on the sheet of paper.

2) Results: As the screen was moved laterally the projected laser point moved significantly to the left compared to the reference point. However, approximately 5 cm from the edge of the screen the projected laser point moved back and assumed roughly the same location as the reference point. Not until the laser point reached the other edge of the screen did an offset on the sketch paper become noticeable again. In short, the first and the last projected laser points were clearly offset from the rest of the projection points as seen in table I. This leads us to believe that the properties of the DCG are different in the edges of the screen compared to the center portion of the screen.

3) Discussion: The path traced stretched laterally from one edge of the screen to the other. It was only in the edges that inhomogeneities were clearly visible. We are speculating that this is due to the fact that the substrate is exposed to air humidity at the edges in its current setup or has been exposed to uncontrolled changes in temperature and humidity during the production process.

TABLE I
OFFSET BETWEEN PROJECTED LASER POINT AND REFERENCE POINT

Point	Screen location (cm)	Offset (mm)	Visual angle (arcmin)
1	5	-37	-4.58
2	15	-12	-1.49
3	25	-6	-0.74
4	35	-9	-1.11
5	45	-6	-0.74
6	55	-9	-1.11
7	65	-6	-0.74
8	75	-9	-1.11
9	85	-9	-1.11
10	95	12	1.49

We could immediately use this result in the next experiment on refraction (see section below) as it proved that varying intersection points between screen and laser beam is an independent variable. The intersection point must remain fixed or else it will influence the results.

In general terms, the results show that normal usage through the center of the screen is not a problem, as registration errors around or slightly above 1 arcmin are hardly detectable. However, the inhomogeneities in the fringes of the screen could be problematic when tiling several screens for a larger field of view, as the error approaches 5 arcmin and the inhomogeneity is not continuous across the edge. A more controlled experiment, preferably over several screens, is needed to further describe this phenomenon.

C. Refraction

A ray of light entering from air, passing through glass and exiting to air should only exhibit a slight parallel displacement proportional to the incident angle. Normally this would not be a problem, but since the data is overlaid on one side of the screen one might experience a registration error between the data layer and the real world due to refraction. Deering has previously described how Snell's law can cause positional errors in displays with thick glass surfaces [9]. The greater the angle the more significant the displacement. Moreover, the effects of the HOE from an oblique viewpoint are not known.

1) Setup: To further study the refraction in the azimuthal plane of the DCG we made use of the same laser, mounted on a stand, projecting a point on a sheet of paper on a perpendicular wall at 27.75 m distance, as visible in Figure 1. The sheet of paper had a 0.5 cm reference grid. As in the previous experiment, the screen was initially taken aside so that the unobstructed laser point could be recorded as reference.

Nine azimuthal angles were marked on the floor. Plumb lines from the screen ensured that the screen could be positioned at these angles varying from 0° to 80°, where 0° meant a perpendicular intersection between laser beam and screen. The screen would rotate around the center of the radial so

TABLE II

OFFSET AND VISUAL ANGLE AS A FUNCTION OF INCIDENT ANGLE

Point	Incident angle (°)	Offset (mm)	Visual angle (arcmin)
1	0	0	0.00
2	10	0	0.00
3	20	0	0.00
4	30	5	0.62
5	40	5	0.62
6	50	5	0.62
7	60	10	1.24
8	70	17.5	2.17
9	80	37.5	4.65

that the beam always would intersect the same point on the screen so as to avoid interference with the aforementioned inhomogeneities.

For each angle the location of the laser point projection was recorded.

2) Results: As the screen was inserted into the beam the laser projection point was the same as the reference point. This indicated that the portion of the screen where the laser intersected was homogeneous. As the screen was rotated clockwise, the laser projection point did not move until the screen reached 30° where the point would be displaced 5 mm to the right compared to the reference point. It would remain at 5 mm offset until the screen was oriented at a 60° angle where the offset would increase to 10 mm.

3) Discussion: The resulting refraction distortion is 0.62 arcmin at 30-50°. At 60° the distortion, 1.24 arcmin, should start to become noticeable, and it approaches 5 arcmin at 80°. If the user stays within 50° from the normal the distortion should not contribute significantly to the resulting registration error. In this region the distortion is less than 1 arcmin in visual angle, which is normally defined as human visual acuity.

III. HUMAN PERFORMANCE

A. Visual Acuity

An experiment was performed to measure visual acuity loss in the HUD. This would give figures on the extent which the HOE (without any projected graphics) actually degrades human vision. The task was to determine visual acuity by having subjects read lines of letters on Snellen eye charts in various display conditions, until two errors on the same line were reported.

1) Setup: The experiment was performed as a mixed design, with two independent and one dependent variable. The independent variables were display condition (within subjects) and ambient light level (between subjects). The display conditions were "no screen" (the subject viewed the eye charts directly), "screen" (the subject viewed the eye chart through the screen), and "glasses" (as "screen" but with the addition of polarized glasses). Ambient light was either on or off (meaning the ceiling lights which provided normal indoor fluorescent office illumination were switched on or off).

The eye charts were illuminated at all times and had a luminance of 156 cd/m^2 measured from the subject viewpoint. The addition of the screen reduced the luminance level to 74 or 62 cd/m^2, depending on whether the ambient light was on or off respectively. Thus the ambient light increased the luminance of the screen by 12 cd/m^2. The addition of glasses reduced the luminance to 30 or 25 cd/m^2. Thus the addition of the screen alone decreased eye chart luminance by about 55%, screen combined by glasses about 80%. (Luminance was measured in the same way as in the Luminance and Contrast section.)

Twelve subjects performed the experiment, 6 with and 6 without ambient light. Subjects were nested within the lighting condition and crossed with all other independent variables. The display condition was counterbalanced throughout the experiment. The subject was placed on a chair 4 m from the eye charts (designed for 4 m viewing distance, the 6/6 line measured 5.8 mm in height), with the chin on a chin rest. Polarizing glasses were added to the chin rest in the "glasses" condition, but the eye height remained constant in all trials. The subjects were told to read from line 4 and down, and were stopped after two errors were reported on the same line. The previous line was recorded as their acuity. If one error was made on the line before the line with two errors, acuity was considered to be half-way between the line of one error and the previously passed line.

2) Results: The result line number is quantized and based on a logarithmic scale. Therefore the line numbers were converted to a linear scale, visual resolution in arcmin, more suitable for averaging. The mean visual resolution values are shown in Figure 5, where the impact of the independent variables are visualized. Further statistical analysis using ANOVA showed that the display condition had a significant effect on the results ($F(2,22) = 4.264$, $p \leq .029$). The "no screen" condition showed a 0.15 arcmin better resolution than the "screen" condition, corresponding to approximately 1 line on the eye chart. There was no significant difference between the "screen" and "glasses" conditions. Ambient light variation showed no significant effect.

3) Discussion: It was surprising that ambient light showed no significant effect on the results. Possibly this was due to the high contrast in the eye charts, combined with the rather low screen luminance increase with ambient light (12 cd/m^2). A future experiment will involve the much greater outdoor luminance conditions, that are similar to those that would be experienced in a control tower.

B. Depth Matching

An experiment was performed to test the user's depth matching ability with stereoscopically rendered objects seen in the HUD. The purpose was to test user performance in matching the depth of real to virtual (rendered) and virtual to real objects in various conditions at range of 3-10 m. This experiment would thus indicate if there is any overall judgment biases associated with depth rendering on the HOE screen.

Even though the distances covered in this experiment are much shorter than the distances observed from a control tower, they would give preliminary measures of the possibilities and

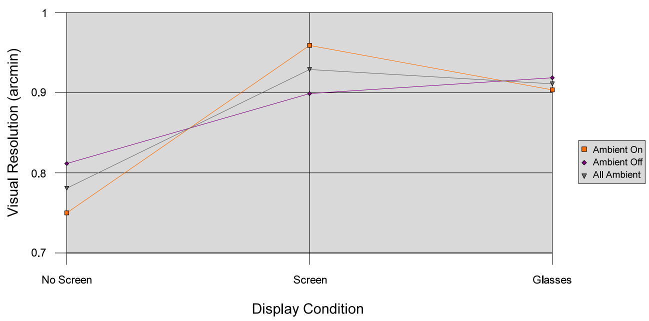

Fig. 5. Mean visual resolution values per display condition

Fig. 6. The setup for the depth matching experiment showing the real object on rails (left) and the virtual object on the screen (right).

limitations of depth rendering in this type of large format HUD.

1) Setup: The experiment was performed as a within-subjects design, with three independent and one dependent variable in a fully crossed repeated measures experiment. The independent variables were screen distance, target distance and target type. Screen distance is the distance in meters from the subject to the screen, which was either 1 or 2.5 m. Target distance is the distance from the screen to the target object, which was either 4, 6 or 7.5 m. Target type is the type of rendered target, which was either real (a real object) or virtual (an object rendered on the screen). The matching task was performed using a cursor. If the target was real the cursor was virtual, and vice versa. The dependent variable measured was cursor distance, the distance from the screen to the cursor. The initial position of the cursor was randomized such that it would be initialized both in front of the target and behind.

There were 12 individual combinations of the independent variables, and the sequences were counterbalanced for each subject. Eight subjects aged 25-56 years performed the experiment. Each combination was repeated 3 times, so each subject performed 36 trials. The experiment took about 1 hour per subject. Before the trials, the subject was given an instruction sheet and a stereo vision test. The subject's inter-pupillary distance (IPD) was measured using an optical stereo-camera eye tracker.

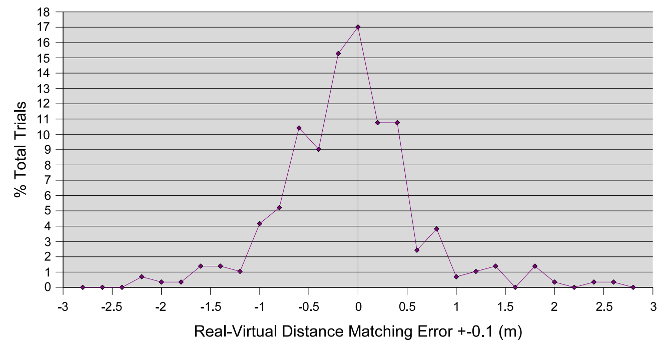

Fig. 7. Mean distance matching error

The trials were performed in blocks of nine trials with short breaks between the blocks. Between these 4 groups of trials the screen distance and target type was changed. The screen distance was changed by the user moving between two fixed office chairs. The polarized glasses were mounted on stand (with a chin rest), so the viewpoint was the same for each user. The chair height was adjusted for each user so the eyes were level with the glasses. If the target type was real, the cursor was virtual. The real object (cursor or target) was a paper bag in an approximately spherical shape or radius 10 cm, standing on a 10x20 cm cart, with wheels on an 8 m rail extending away from 2 m behind the screen as shown in Figure 6.

Each trial was initiated by the subject pressing the space bar on a keyboard. If the trial had a real target (i.e. the cursor was virtual) the cursor was rendered at the randomized initial distance. The subject changed the distance of the cursor by pressing the keyboard arrow buttons, and when finished pressed the space bar again. The final cursor position was logged automatically. If the target type was virtual (i.e. the cursor was real) the cursor distance was adjusted by pulling a pair of strings attached to the cart. When the subject was finished adjusting the distance, space bar was pressed and the cursor position was recorded by manual observation on a measuring tape in the tracks. The subject was informed that there was no time constraint of the trial.

The stereo images were produced using polarization multiplexing, using the same method as described in [5]. Two projectors, each of a maximum luminosity of 6500 ANSI lumens, each projected a linearly polarized image in XGA resolution of size 745x552 mm on the HUD. The stereo image was calibrated according to the user IPD.

2) Results: The main result of this experiment is shown in the real-virtual distance matching error distribution in Figure 7. It shows, for all independent variable combinations, the absolute error in the distance match between the real object and the virtual (regardless if they were cursor or target in the matching). Each data point contains the range of error values ±10 cm from it. The overall mean real-virtual distance matching error was -0.084 m (SE=0.043), i.e. a slight negative bias.

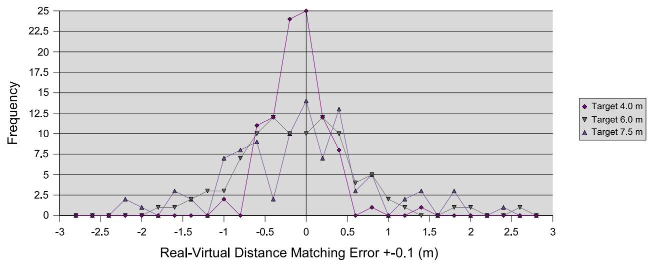

Fig. 9. Error distribution per target distance

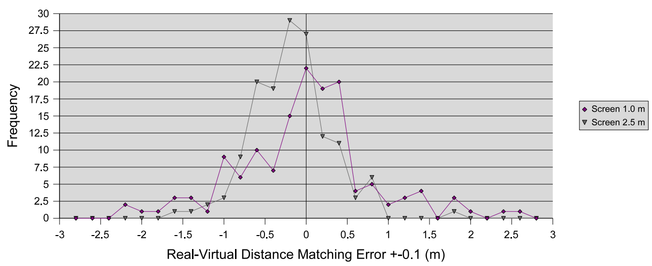

Fig. 10. Error distribution per screen distance

This means the real object was on average judged to be located 8.4 cm further than the virtual, so when comparing with the real object the virtual object distance was overestimated by 8.4 cm. In the following figures, errors below zero indicate virtual object overestimation. Figure 8 shows the error distribution per subject. Figures 9, 10, and 11 show the error impact of the different independent variables and their values.

3) Discussion: We had expected some statistically significant effects for some of the variables. This was not the case however, which has both positive and negative interpretations. Positive is that e.g. screen distance has no significant impact on the results, which here means that the user can come close to the screen, 1.0 m, and not perform statistically worse than at the theoretically better 2.5 m position. Also positive is that the results are not statistically different when increasing the target distance from 4 to 7.5 m. However, the negative interpretation is that the statistical power was insufficient or our measurement technique was too noisy, and this is

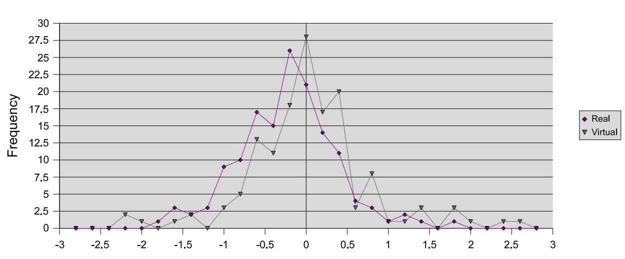

Fig. 11. Error distribution per target type

the reason why no significant effects are showing. We have identified a number of issues that should be addressed in a follow-up experiment which could more accurately pinpoint where significant effects are found:

- Change the real object to a flat shape (no depth)
- Change the virtual object to a flat shape (no depth)
- Increase the relative size difference between the real and virtual objects, so that no relative size comparison can be performed
- Include a closer screen distance for increased accommodative change and accommodation-vergence mismatch
- Increase the range of target distances into the far field to more closely match the real application
- Measure each subject's individual IPD with higher accuracy under same convergence conditions as the experimental trials

IV. Conclusion and Future Work

This paper reports the results from a set of preliminary studies on a new HUD format, aiming to provide parameters for future controlled experiments.

Refraction can be excluded as a significantly contributing source to registration error if the screen is used within $\pm 50°$ from the normal. More oblique angles can result in small but detectable distortions.

Inhomogeneities in the DCG causes perpendicular light rays to refract when they intersect the edges of the screen. It should not pose a major problem when using a single screen, but the inhomogeneities in the edges could pose a challenge when expanding the field of view by tiling multiple screens.

There is a statistically significant contrast reduction and subsequent loss of visual acuity of observed objects through the screen. Increased ambient light does not significantly affect the visual acuity in an indoor lighting condition, but we suspect this will be the case in high ambient light conditions like the control tower. One potential solution to this problem would be to reduce the amount of ambient light in the tower by mounting tinted window films. The effects of high ambient light and potential solutions will be further investigated.

Tinted window films may also be an approach to increase contrast levels in the screen, which are found to be approximately 1/3 of that needed when comparing to aircraft HUDs. Brighter projectors could be another (more costly) approach.

A slight overestimation (8.4 cm) was found in the matching of virtual objects with real using stereoscopic rendering on the screen over short and medium ranges (3-10 m). However, the results did not show any statistically significant effects between the measured variables. A follow-up experiment will involve modified target shapes and broader ranges in the measured variables. Especially the target distance range will be substantially increased in order to determine the effects of depth rendering in distances common in a tower environment.

Appendix

Using the recommended aperture (N) when using the camera in automatic mode (no flash) and exposure time (t),

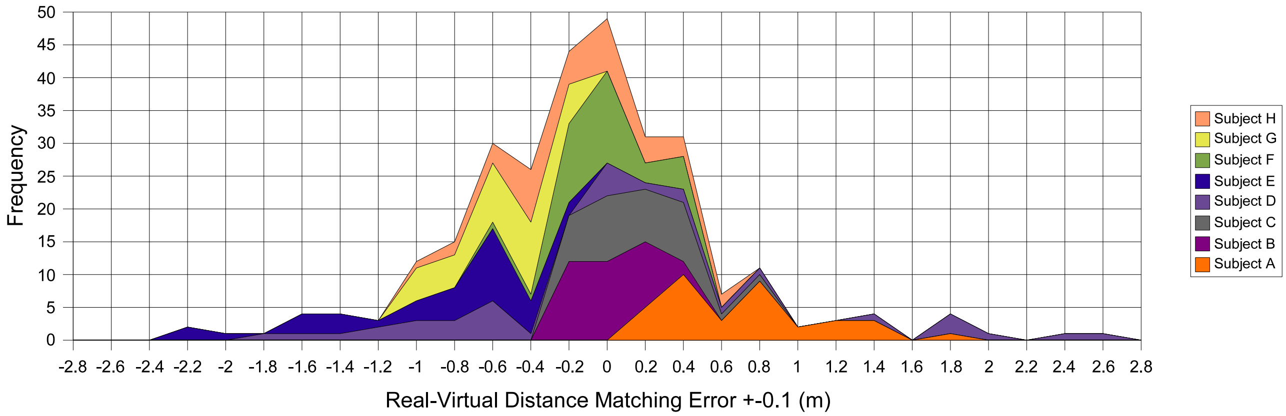

Fig. 8. Error distribution per subject

together with the ISO speed (S), luminance levels (L) in cd/m^2 could be calculated using the reflected light exposure, see equation 1. K, the reflected-light meter calibration constant, was chosen to be 12.4, a commonly used K-value in the ISO recommended range[1]. S was pre-set to ISO 400.

$$L = \frac{K \cdot N^2}{t \cdot S} \tag{1}$$

$$c = \frac{L_{Max} - L_{Min}}{L_{Max} + L_{Min}} \tag{2}$$

Contrast values are calculated with the Michelson formula, see equation 2.

References

[1] D. Weintraub and M. Ensing, "Human factors issues in head-up display design: the book of hud." CSEIRAC, Wright Patterson AFB, Ohio, 1992.

[2] S. R. Ellis, "Towards determination of visual requirements for augmented reality displays and virtual environments for the airport tower," in *Proceedings of the NATO Workshop on Virtual Media for the Military*, no. HFM-121/RTG 042 HFM-136, (West Point, N.Y.), pp. 31–1 – 31–10, 2006.

[3] R. J. Reisman and S. R. Ellis, "Augmented reality for air traffic control towers," in *Proceedings of the SIGGRAPH 2003 Conference on Sketches and Applications*, 2003.

[4] N. Furstenau, M. Rudolph, M. Schmidt, and B. Lorenz, "On the use of transparent read projection screens to reduce head-down time on the air traffic control tower," in *Proceedings of Human Performance, Situation Awareness and Automation Technology Conference*, 2004.

[5] S. Peterson and E. Pinska, "Human performance with simulated collimation in transparent proejction screens," in *Proceedings of the Second International Conference on Research in Air Transportation (ICRAT)*, pp. 231–237, 2006.

[6] S. R. Ellis, J. Schmidt-Ott, J. Krozel, R. J. Reisman, and J. Gips, "Augmented reality in a simulated tower environment: Effect of field of view on aircraft detection," Tech. Rep. NASA/TM-2002-211853, NASA Ames Research Center, NASA Ames Research Center, Moffett Field, California 94035-1000, September 2002.

[7] P. Smid, H. Hiklova, and J. Keprt, "Dichromated gelatin and its importance for optical hologram recording," *Czechoslovak Journal of Physics*, vol. 54, pp. 1461–1472, December 2004.

[8] R. L. Holloway, "Registration error analysis for augmented reality." Presence: Teleoperators and Virtual Environments 6(4) (August 1997), pp. 413-432., 1997.

[9] M. Deering, "High resolution virtual reality," in *SIGGRAPH '92: Proceedings of the 19th annual conference on Computer graphics and interactive techniques*, vol. 2, (New York, NY, USA), pp. 195–202, ACM Press, 1992.

[1]ISO 2720:1974, http://www.iso.org/iso/en/CatalogueDetailPage.CatalogueDetail?CSNUMBER=7690, Dec 18, 2006.