Comparing Interaction in the Real World and CAVE Virtual Environments

Alistair Sutcliffe
Brian Gault
Centre for HCI Design, Dept of Computation
University of Manchester Institute of Science and Technology (UMIST)
P.O. Box 88, Manchester M60 1QD, UK
+44 (0)161 200-3315
a.g.sutcliffe@co.umist.ac.uk

Terence Fernando
Kevin Tan
Centre for Virtual Environments
Business House, University Road
University of Salford
Salford
Manchester M5 4WT, UK.
t.fernando@salford.ac.uk

ABSRACT
An experimental comparison of interaction in the real world and a CAVE virtual environment was carried out, varying interaction with and without virtual hands and comparing two manipulation tasks. The double-handed task was possible in the real world but nearly impossible in the VE, leading to changed behaviour. The single-handed task showed more errors in the VE but few behaviour differences. Users encountered more errors in the CAVE condition without the virtual hand than with it, and few errors in the real world. Visual feedback caused many usability problems in both tasks. The implications for VE usability and virtual prototyping are discussed.

Keywords: CAVE, virtual reality, usability evaluation, manipulation tasks

INTRODUCTION
There have been few studies to compare the differences between interaction in virtual environments and comparable real worlds domain, in spite of the concern of presence in VE research [18]. In mixed reality environments users transfer between real and virtual worlds [4]; ethnographic studies on such interaction suggest that users find the transition between real and virtual worlds difficult [6, 14]. Virtual prototyping is advocated as an important application for VR, in which the object is to discover usability problems that are attributable to the designed artefact rather than the VE. Hence user behaviour in the VE with virtual prototypes should be as natural as possible. The first motivation for this paper was to discover and qualify the deficit between real world (RW) and VE interaction.

Few usability evaluation studies have been conducted on CAVE environments, and interaction therein normally assumes presence of a virtual hand. However, in contrast to head-mounted displays, CAVEs enable users to see their real hands. If presence and naturalness are to be encouraged for virtual prototyping, then interaction using one's own hand should be an advantage. This formed a second motivation for this study.

Several studies have pointed out usability problems in VEs [9, 13]. VE applications have interaction styles radically different from standard GUIs, as illustrated in the work of [5] and [16]. These differences concern the ability to manipulate virtual objects interactively, and the VE's immersive capability. However, evaluation methods have not been specifically designed for VEs, which poses many problems not encountered in GUI interaction. Some guidelines and generic design issues for VEscan be found in [15], and the more extensive guidelines in [9]. A preliminary VE evaluation method and usability heuristics were proposed by [19] while [12], developed a VE usability questionnaire based on standard GUI principles; however, to our knowledge no evaluation methods have been proposed for CAVEs. The third motivation was to develop evaluation techniques for VEs.

This paper focuses on investigating the usability problems in fully immersive VE applications and comparing them with user behaviour in the real world; it identifies interaction deficits imposed by the VE technology, and the strategies that users adopt to deal with such problems. In the next section we describe the method and experimental design of the study. We then report and analyze results, and conclude with a discussion of the implications of our findings for the usability and design of VEs.

MOTIVATIONS AND MATERIALS
The experimental design had three independent variables: task, user experience and real world versus virtual world (with and without virtual hand) in a 2 x 2 x 3 partially

factorial design (task order was fixed), with dependent measures of errors and task completion times. Choice of tasks was motivated by [10]'s kinematic chain model of haptic interaction, which predicted high cognitive loading for synchronized two-handed interaction, so a double-handed task was selected to create conditions of high performance difficulty, with a second easier single-handed task. The experimental domain was a simple manipulation task of moving chess pieces to reposition them into a target pattern. In the two-handed condition this involved picking up, moving and placing two pieces simultaneously; while in the single-handed condition, subjects were instructed to pick a piece in one hand then pass it to the other hand before placing it in the correct location. The motivation for the hand-to-hand piece-passing was to study the usability problems inherent in visual feedback substituted for haptic interaction in the VE design. Pilot tests demonstrated that subjects had great difficulties in completing the double-handed, synchronous task in the CAVE condition. We changed the original randomized design to study learning effects between the tasks, which were held in the same order: double-handed first, single-handed second. Errors were hypothesized to be higher in the doubled-handed condition for which no learning was possible; in that condition we investigated adaptive strategies. In the single-handed condition, learning was possible where subjects had adopted a similar pattern of behaviour in the previous task. We investigated whether adaptive strategies reduced errors and helped performance in the second task.

The experimental paradigm consisted of two consecutive tests for each subject:

- All subjects were exposed to the control condition in the real world. The task end state and requirements were the same as the experimental conditions, namely to reorganize chess pieces placed randomly, into a desired pattern.
- Subjects then repeated the task in the experimental condition, performing the double-handed and single-handed tasks in that order. The presence or absence of the virtual hand was randomised following conventional experimental design.

A fully immersive CAVE system [7, 8] was equipped with shutter glasses to give the users stereoscopic views, and pinch gloves for manipulation. Head-tracking devices mounted on the users' shutter glasses controlled the CAVE viewpoint according to the users' body and head movements.

The application displayed 12 chess pieces: black and white king, queen, bishop, knight, rook, and pawn, and a board with minimal background (see figure 1). The user tasks were to move the chess pieces from a random layout to the target arrangement shown in the CAVE display. Haptic feedback for piece selection was substituted by colour changes, as follows. When a chess piece was selected by the user operating the pinch glove, it changed colour (from white/black to yellow) in response to the user's action (figure 1). The piece then moved in tandem with the user's hand until it was released. Once the piece was released it reverted to its original colour.

Figure 1. Selection of a chess piece.

When a chess piece was placed on the chessboard, the square on which it had been placed changed colour from white/blue to dark red (see figure 2). When the user released the piece, the square reverted to its original colour.

Figure 2. Placing a chess piece on the chessboard.

When an already selected chess piece (coloured yellow) was correctly positioned to be gripped for selecting by the other hand, it changed colour from yellow to blue. The user could then release the piece with their first hand and the colour changed back to yellow (= selected) (figure 3).

Fourteen students from UMIST and Salford Universities (1 female and 13 male) took part in the study; they were equally divided between VE expert and VE novice users. The average age of the subject group was 29 years, all were right-handed, with good eyesight, and none wore spectacles. Each subject performed the task to arrange the initially scattered twelve chess pieces (as illustrated in figure 4) into the final arrangement (illustrated in the VE, figure 3).

Figure 3. Passing a chess piece from hand to hand.

Figure 4. Initial set-up of chess pieces.

When they had completed the real world control task, each subject had their inter-ocular distance (IOD) measured to calibrate the VE application to their individual depth perception. Each subject was then introduced to the CAVE, fitted with a pair of stereoscopic shutter glasses and a pair of pinch gloves, and allowed a 10-minute period to familiarize him or herself with the VE chess application. At the end of this period, each subject was asked to complete the two tasks in one of two different conditions: (a) VE application with the virtual hand represented, or (b) VE chess without the virtual hand. In both conditions a precise bounding box provided collision detection when the user's hand intersected with the surrounding volume of each chess piece. When the virtual hand was not present the user had to infer the offset between their real hand and its virtual position in the CAVE environment.

The subjects were asked to report verbally any difficulties they experienced while completing the tasks; they were timed and videotaped during their performance. At the end of the study, each subject was questioned during a short debriefing session about the problems encountered during interaction, and their perceptions of the real world and VE application .

The following hypotheses were proposed:

- Two-handed interaction would take longer and produce more errors than single-handed interaction because it required more conscious motor control.
- More errors and longer task completion times would be observed in the VE than in the real world because of the limited haptic feedback in the VE.
- In the VE the presence of the virtual hand would reduce task completion times and errors while the absence of the virtual hand would make task completion more difficult.
- In the VE expert subjects would perform better (shorter times, fewer errors) than novice users in both conditions.
- The VE would increase the subjects' cognitive workload, leading to changes in their behaviour and task completion strategies when compared to the real world.

RESULTS

Task times and errors

Task times were analyzed using a repeated measures ANOVA, with user experience, task type, and task environment as the independent variables. The analysis indicated a significant difference for task ($p < 0.001$) which reflects the quicker mean completion times in the real world environments. In addition, there was a significant interaction between task and environment ($p < 0.01$), and a weaker effect for task and subject experience ($p = 0.05$). Mean task completion times (see table 1) were much faster in the real world task compared to the CAVE, as might be expected, and the double-handed task times were quicker than the single-handed task, in the virtual environment. However, there was little difference in completion times between the tasks in the real world, which indicates that the virtual environment slowed performance and imposed a different additional load on the subjects for both task. Experts were quicker than novices in the single handed task but differences were inconsistent in the double handed task.

Behaviour and errors

Videoed subject behaviour was categorized as: (i) task action (indicating subject movement of the chess pieces); (ii) errors (indicating mistakes made by the subjects in the movement of the chess pieces); and (iii) body moves (movement of the subject's body from place to place particularly in the VE), as shown in table 2.

Table 1. Mean task completion time in seconds.

Task	RW	VE+hand	VE no-hand
Double-handed:			
Novice	44.14	291.43	531.43
Expert	45.86	308.57	428.57
Total	45.00	295.71	480.00
Single-handed:			
Novice	43.71	488.57	1028.57
Expert	40.29	308.57	745.71
Total	42.00	394.29	887.10

Table 2. Subject behaviour categories.
Behaviour types: *TA* = Task Action, *TE* = Task Errors, *BM* = Body Moves.

Behaviour Type		Category Description
TA	DP	Double Pick Up of Chess Pieces
	DM	Double Movement of Chess Pieces
	DD	Double Put Down of Chess Pieces
	SP	Single Pick Up of Chess Piece
	SM	Single Movement of Chess Piece
	SD	Single Put Down of Chess Piece
	PP	Pass Chess Piece from Hand-to-hand
TE	H	Hesitation
	I	Incorrect Positioning of Chess Piece
	WP	Wrong Hand Pick Up of Chess Piece
	WM	Wrong Hand Move of Chess Piece
	WD	Wrong Hand Put Down of Chess Pc
	C	Correction of Earlier Error
	FS	False Start
	V	Invalid Movement of Piece
	RS	Reselect Chess Piece
BM	BL	Body Move Left
	BR	Body Move Right
	BF	Body Move Forwards
	BB	Body Move Backwards
	L	Body Lean
	AR	Body Arm Reach

Most of the categories are self explanatory, apart from Hesitations, which were recorded when subjects did not carry out any purposeful activity for 10 seconds; and False Starts, which were scored when subjects attempted to pick up a piece but then never completed the action. When a subject's action was observed to be unsuccessful, e.g. the subject was unable to pass a chess piece from hand to hand, an error was scored by adding the prefix *U* to the action, e.g. an unsuccessful pick piece was recorded as UPP. Behaviours were transcribed so sequences could be analysed as well as total frequencies.

ANOVA analysis of the total frequency of task actions (see table 3) was significant for task ($p < 0.001$) with a weak interaction for task and environment ($p = 0.025$) with the single-handed task taking nearly twice as many bouts, reflecting the additional actions necessary for passing pieces.

Table 3. Mean frequencies for task actions.

Task	RW	VE+hand	VE no-hand
Double-handed:			
Novice	29.86	35.42	38.00
Expert	24.86	35.57	35.00
Total	27.36	35.50	36.50
Single-handed:			
Novice	60.00	58.43	59.42
Expert	58.29	57.57	58.71
Total	59.14	57.00	59.14

Task errors showed a weak significant main effect for task ($p < 0.05$) with a weak interaction between task and environment ($p = 0.05$). Task errors (table 4) showed an inconsistent picture: means for the double-handed task were higher in the real world condition; the single-handed errors were higher in the VE conditions. Experts had lower averages than novices.

Table 4. Means for task errors.

Task	RW	VE+hand	VE no-hand
Double-handed:			
Novice	4.00	1.86	2.57
Expert	2.14	1.14	1.71
Total	3.07	1.50	2.14
Single-handed:			
Novice	2.14	4.57	7.43
Expert	0.71	2.57	4.42
Total	1.43	3.57	5.93

Usability errors did not apply in the real world. Frequencies of usability errors were significantly different between the VE environments ($p < 0.05$) and showed an interaction between environment and experience ($p =< 0.01$). Higher error averages were observed in the no-virtual hand conditions and in the single-handed task, while novices had higher average errors than experts in both VE conditions. This may have been caused by the more complicated manipulation of selecting and releasing pieces when passing them between hands in the VE, and the inexperience of novices in the VE manipulations.

Three usability errors accounted for over 95% of the observed subjects' critical incidents. First, unsuccessful selection of a piece (USP errors, 50% of total) was caused by perceptual difficulty in positioning the virtual hand (or their real hand in the no-virtual hand condition) in the correct position so the bounding box intersected with the chess piece. Once the intersection had been achieved most subjects then pinched the glove to successfully select the piece; however, in 65% of USP errors the subjects failed to complete selection either because they forgot to pinch the glove or because they moved their hand before pinching, which then inhibited selection. The second most frequent problem was the unsuccessful pass piece (UPP), accounting for 39% of the overall usability errors. In these cases the user correctly selected the first piece, then moved the left hand to an appropriate position for hand-to-hand passing, but then released the right hand before selecting with the left. This resulted in an unselected piece floating in front of them. The final usability error, unsuccessful put piece down (UPD, 6%), arose because the user could not select the square on the chess board or, once having done so, did not release the glove pinch in time, consequently disabling the place piece operation. Novices accounted for 74% of the total usability errors, and had higher scores for each type; furthermore, experts had no UPD errors.

Table 5. Means for usability errors.

Task	VE+hand	VE no-hand
Double-handed:		
Novice	2.57	12.57
Expert	1.00	3.14
Total	1.79	7.86
Single-handed:		
Novice	9.43	15.29
Expert	2.14	8.14
Total	5.79	11.71

The majority of errors were therefore caused by perceptual problems in manoeuvering the virtual hand, causing the collision detection algorithm to be triggered and the appropriate selection highlighted. However, the UPP errors were caused by a complex set of manipulations that had little mapping to the real world task, e.g. gripping the selected piece with both hands before releasing it with the right hand in hand-to-hand passing. A fourth problem, which did not lead to any specific usability problems, was when users experienced difficulty in trying to position their virtual hands to select pieces. We refer to these incidents as the "cross hands" problem, when users found themselves in a contorted position with their real hands crossed over in a vain attempt to reposition their virtual hands. Their solution was to abandon the movement and start the approach to select a chess piece again.

Sequences of behaviour were analysed by casting transition frequencies for all dyadic combinations of behaviour categories (i.e. frequencies where A followed B, B was followed by C, etc.) in matrices and then constructing behaviour network graphs for each subject. Individual subject behaviour network graphs were inspected for significant commonalities or differences; however, no particular common patterns of behaviour were observed among the subjects for either task of the VE condition. Behavioural frequencies were summed for novice and expert subjects by task and VE condition, and group-level diagrams created. To test for the more significant transitions an expected value was calculated for each cell in the matrix by dividing the total transition frequencies by the number of cells, having eliminated the zero diagonal. The expected value was then used in the Binomial test to calculate the z distribution for sample sizes where $N > 25$.

In the following network diagrams only frequencies above 2% of the overall total are reported. The abbreviations (SP, SD, etc.) are spelt out in table 2. Dashed arrows are non-significant, normal arrows show transitions that were significantly more frequent than the expected value at $0.01 < p < 0.05$, and bold arrows for $p < 0.01$.

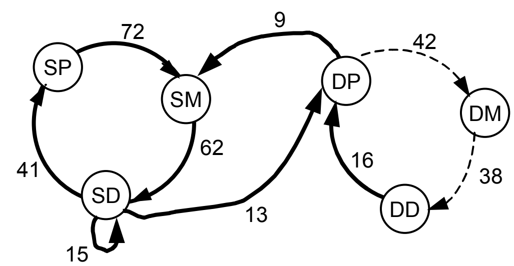

Figure 5. Behaviour network for all subjects: real world double-handed task. Transition frequencies are shown on the arcs (i.e. SD was followed by SP 41 times)

The subjects showed two patterns for the first double-handed task (figure 5): one cycle of picking, moving and placing pieces concurrently, and a separate, more significant cycle of single-handed operation which did not conform to instructions. Only one subject completed the task without error, while three subjects failed to perform any two-handed interaction at all. The expert-only pattern (not illustrated) was similar to figure 5 with the addition of transitions from Corrections and Hesitations, which, although non-significant may indicate the difficulty of the double-handed task; corrections in the single-handed task were caused by subjects changing their minds on piece selection.

The novice pattern showed an even stronger bias towards single-handed operation, reflecting the difficulty experienced in synchronizing doubled-handed manipulation even in the real world. Incorrect position errors indicated task knowledge problems with selecting the appropriate square for the piece.

In contrast, the single-handed task was much easier and more natural for all subjects (figure 6); furthermore, both experts and novices showed similar patterns. Experts approached a near perfect cycle of picking up a piece, moving it, then passing it from hand to hand before moving it to the required square and placing it (6 black and 6 white pieces required 12 moves x 7 = 84).

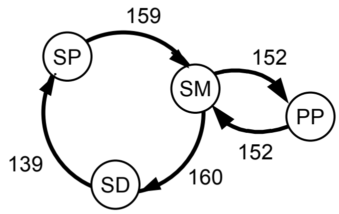

Figure 6. Behaviour network for all subjects: real world single-handed task.

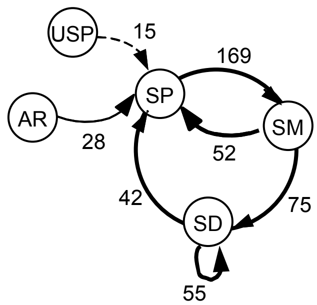

Figure 7. Behaviour network for all subjects: VE, double-handed task with virtual hand.

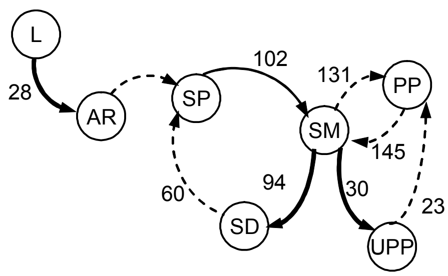

Figure 8. Behaviour network for all subjects: VE, single-handed task with virtual hand.

The subjects' behaviour pattern in the VE for the double-handed task reverted to single-handed mode, even for experts (see figure 7), but differed from the real world pattern with more transitions between piece selects and moves. This suggests that the additional cognitive load of the virtual environment made double-handed synchronization nearly impossible. The expert behaviour patterns also showed some movement and piece selection errors, possibly caused by perceptual difficulties. Novices followed experts in reverting to single-handed performance, but showed more leaning and arm reaching movement, which may reflect their lack of perceptual experience in the CAVE. Perceptual problems may also have been responsible for their high frequency of piece selection errors.

In the single-handed task (figure 8), the subjects showed similar behaviour patterns in the CAVE and the real world, with the addition of arm reach body movement to select pieces. Experts experienced few errors in this task condition. Novices, in contrast, showed little body movement apart from arm reaching, so the experts' experience may have made movement more natural for them. Novices also experienced more usability errors, especially with passing pieces which required them to select with both hands before releasing with the first.

In the double-handed task without the virtual hand (figure 9), user reverted to a single handed strategy; furthermore the cycle of select, move and place was disrupted by piece selection errors. Novices tended to move more, and encountered more errors than experts in trying to pick up and pass pieces. This indicates that experts as well as novices found adjusting to the real-hand cue confusing, until they learned the perceptual offset between the VE depth and their hand. This is indicated by the movement actions preceding the USP errors.

The single-handed task patterns without the virtual hand showed many piece passing errors (figure 10). While selection errors were more common in the double-handed task, in the single-handed task placement errors were more frequent. The experts' pattern showed less movement compared with the double-handed task and more piece-passing errors. Novices showed frequent place and select piece errors, with several reselect piece actions demonstrating their difficulty in operating without the virtual hand; in addition, they showed more movement than in the double-handed task, indicating more perceptual exploration.

In summary, the behaviour pattern analysis showed that the double-handed task was too difficult for subjects in the virtual environment, and both novices and experts reverted to single-handed operation. The patterns of both experts and novices for both conditions were consistent for each task, while the main novice-expert differences were more usability errors from the novices and more movement by the experts, although this was not a consistent trend.

Qualitative data and subject comments

Subject comments were grouped into comments about hardware, depth perception and presence, and interaction problems. Five novices complained about interaction, in particular passing pieces where the colour change was confusing; three commented on depth perception, which was worse for more distant pieces; and four subjects commented that they had to concentrate because interaction was difficult. Problems with the shutter glass led to three subjects' complaints. Only 12.5% of the comments were positive, and all these were about the VE presence. Four experts complained about the depth perception, although four also commented favourably on the quality of presence, with two subjects sharing both views. Four subjects found the mapping and orientation between the virtual hands and their real hands to be difficult because their real hands "blocked" the virtual hands and they could not see if the chess pieces were being selected or positioned correctly. This problem was highlighted when passing a chess piece. Most subjects found movement within the CAVE to be natural; however, four subjects experienced a time delay or lag in updating of the VE when moving within it. 50% of the subjects remarked that they experienced problems with depth perception, when the perceived position of a chess piece did not correspond to where they had reached to select it. This was particularly a problem for the novice subjects, whereas expert subjects moved around more in the VE to improve their perception when selecting distant chess pieces.

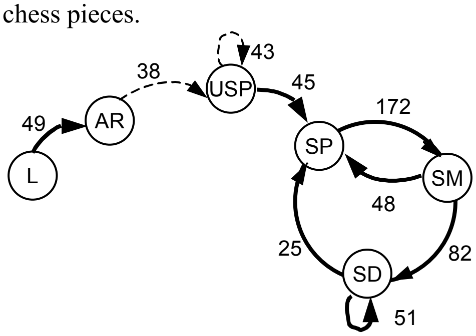

Figure 9. Behaviour network for all subjects: VE, double-handed task, no virtual hand.

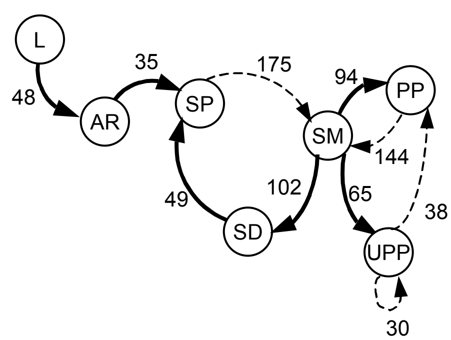

Figure 10. Behaviour network for all subjects: VE, single-handed task, no virtual hand.

DESIGN IMPROVEMENTS

The VE design could be improved for piece selection (USP error) in several ways. First the 3D perceptual depth could be improved which was a particular problem in corners of the CAVE where distortion is more severe. Secondly, the bounding box- collision detection mechanism could be improved to map more accurately to the users virtual hand. Third, the selectable status might be cued more effectively using a handles metaphor as used in GUI drawing applications with colour change, this could leverage users existing mental model for affordances. Alternatively a 'snap to' auto-selection function could be triggered when the glove was pinched, or even on collision detection. The extent of automated support for interaction will depend on the demands for naturalness of the VE. Finally haptic feedback for grip could be provided by vibration emitting devices in the glove fingers, with movement being calculated from an effective grip by thumb and fingers.

Many of the above improvements could cure the other pass piece (UPP) and place piece (UPD) errors. For UPP errors adding simulated gravity could have cued rapid learning not to drop a piece before it had been picked by the other hand, furthermore pressure feedback by finger tip pressure sensor would give better haptic perception of grip by both hands. The cross hands problem might be cured by improving depth perception or by interacting via the users real hand as in the no virtual hand condition. However, this condition generally caused more errors than the virtual hand, but a more accurate collision detection with a wider range bounding zone and a visual cue for a nimbus [2] might make this interaction technique more effective

DISCUSSION AND CONCLUSIONS

Reviewing our research hypotheses, the hypothesis 1 (double handed task longer than single handed) was not upheld, since the double-handed interaction was quicker and induced fewer task and usability errors than the single hand, part from the real world where the opposite results were found. A possible explanation is that usability problems in the VE made the single handed task longer and more error prone. However, there was considerable variation in subject behaviour in the double-handed task, and in many cases subjects were arguably performing single-handedly.

The second hypothesis (more errors in VE than in real world) was upheld, as more errors and considerably longer times were observed in both VE conditions. This difference was not surprising; however, the quality differences in user behaviour were more interesting. Our experiments demonstrated the difficulty of carrying out complex manipulations in CAVE environments, especially when the task demanded synchronization and double-handed operation. These tasks are known to be difficult in real world conditions [8], so the degraded performance we observed between the real and virtual world operation of the same task strongly suggests that CAVE environments impose a considerable cognitive burden on users. Furthermore, our subjects changed their task strategies in the face of these difficulties, so CAVE environments might induce a completely artificial pattern of user behaviour. This has important implications for virtual prototyping uses of CAVEs, where the observed problems are more likely to be artefacts of the virtual environment than interaction problems with the designed artefact.

The third hypothesis (virtual hand better than no virtual hand) was upheld, as significant differences were found in both tasks. Interaction without the virtual hand proved to be more error prone than interaction with the virtual hand presence, even though our subjects said they preferred the no-hand condition. Our initial hypothesis was that the no-hand condition would be easier than the virtual hand condition; however, it appears that interaction without the hand gave the users worse perceptual problems. They had to judge the location of a virtual bounding box within the 3D virtual world; unfortunately, the mapping between the depth dimension in the CAVE and the location of the subject's real hand was far from perfect, so our users had to discover perceptual offsets from different viewpoint angles. This led to increased selection errors. However, the virtual hand condition did cause the cross-hand problem, which did not occur when no virtual hand was present. Interaction with no virtual hands in CAVE environments may therefore be advantageous, but only if improved perceptual mapping with the real user's hand and collision detection in the virtual environment can be provided.

The fourth hypothesis (expert better than novices) was partially supported. Experts did show fewer errors than novices, but most differences were non-significant, and task completion times were not consistently quicker. This may be seen as a positive indication that novices can learn to interact in CAVE environments with minimal training, as indeed they did in our study. Alternatively, the persistence of errors shown by experts suggests that perceptual problems do not diminish with training, and these are deeper-seated design defects in the current technology. The behaviour differences between the subject groups indicated that experts got used to CAVE environments and moved around more naturally, whereas novices tended to be rooted to the spot. Experts' use of movement may compensate for perceptual inaccuracies in the CAVE.

The final hypothesis (VE would change behaviour compared with real world) was partially supported. Although we took no direct measure of cognitive workload, it appears that the CAVE environment did induce a radical change in subjects' behaviour, and caused many of them to abandon the cognitively more demanding doubled-handed task. A possible explanation for the additional load is the lack of haptic feedback; poor visual feedback for collision detection imposed an additional cognitive load in maintaining synchronized action. The usability errors we observed were frequent but they could be attributed to a small number of causes. First were perceptual problems in selecting objects. Selection problems caused by difficulty in judging 3D depth in VEs have been reported in other studies [1, 11, 17, 19]. Techniques to make collision detection and bounding boxes more visible to users, e.g. the nimbus concept [2, 3] might be one way to improve interaction, although such feedback can be intrusive and impair the user's sense of presence [18].

Clearly this study has limitations in the generalizability for haptic interaction, and its implications for virtual prototyping. Nevertheless, we argue that the difficulties we found in a relatively simple task do not augur well for more complex manipulations in VEs. In our future work we will investigate the effects of visual and haptic feedback on more complex manipulations in maintenance tasks, to explore the limitations of what can be achieved in virtual prototyping.

ACKNOWLEDGEMENTS

We would like to thank students and researchers in The Centre for Virtual Environments, University of Salford, and Dept of Computation, UMIST for their participation.

REFERENCES

1. Barfield, W., Zeltzer, D., Sheridan, T., & Slater, M. (1995). Presence and performance within virtual environments. In W. Barfield, & T. A. Furness (Eds.), *Virtual environments and advanced interface design* (pp. 473-513). New York: Oxford University Press.
2. Benford, S., Bowers, J., & Fahlen, L. (1994). Managing mutual awareness in collaborative environments. *Computing*(12 May), 39.
3. Benford, S., Greenhalgh, C., Bowers, J., Snowdon, D., & Fahlen, L. E. (1995). User embodiment in collaborative virtual environments. In I. R. Katz, R. Mack, L. Marks, M. B. Rosson, & J. Nielsen, (Eds). *Human Factors in Computing Systems: CHI 95 Conference Proceedings, Denver CO 7-11 May 1995*, (pp. 242-249). New York: ACM Press.
4. Benford, S., Greenhalgh, C., Craven, M., Walker, G., Regan, T., Morphett, J., & Wyver, J. (2000). Inhabited television: Broadcasting interaction from within collaborative virtual environments. *ACM Transactions on Computer-Human Interaction*, *7*(4), 510-547.
5. Bowman, D. A., & Hodges, L. F. (1997). An evaluation of techniques for grabbing and manipulating remote objects in immersive virtual environments. In *Proceedings: 1997 Symposium on Interactive 3D Graphics*.
6. Craven, M., Taylor, I., Drozd, A., Purbick, J., Greenhalgh. C., Bebford, C., Fraser, S., Bowers, J., Jaa-Aro, K. M., Linterman, B., & Hoch, M. (2001). Exploiting interactivity, influence, space and time to explore non-linear drama in virtual worlds. In J. A. Jacko, A. Sears, M. Beaudouin-Lafon, & R. J. K. Jacob, (Eds). *CHI 2001 Conference Proceedings: Conference on Human Factors in Computing Systems, Seattle 31 March-5 April 2001*, (pp. 30-37). New York: ACM Press.
7. Cruz-Neira, C., Sandin, D. J., DeFanti, T. A., Kenyon, R. V., & Hart, J. C. (1992). The CAVE: Audio visual experience automatic virtual environment. *Communications of the ACM*, *35*(6), 64-72.
8. Cruz-Neira, C., Sandin, D. J., & DeFanti, T. A. (1993). Surround-screen projection-based virtual reality: The design and implementation of the CAVE. Computer Graphics. *Proceedings: SIGGRAPH93, Computer Graphics*, (pp. 135-142). New York: ACM Press.
9. Gabbard, J. L., & Hix, D. (1997). *A taxonomy of usability characteristics in virtual environments: Deliverable to Office of Naval Research, grant no. N00014-96-1-0385*. Blacksburg VA: Dept of Computer Science, Virginia Polytechnic

Institute.
10. Guiard, Y. (1987). Asymmetric division of labor in human skilled bimanual action: The kinematic chain as a model. *Journal of Motor Behavior*, *19*(4), 486-517.
11. Hix, D., Swan, J. E., Gabbard, J. L., McGee, M., Durbin, J., & King, T. (1999). User-centered design and evaluation of a real-time battlefield visualization virtual environment. In L. Rosenblum, P. Astheimer, & D. Teichmann, (Eds). *Proceedings: IEEE Virtual Reality '99, Houston TX 13-17 March 1999*, (pp. 96-103). Los Alamitos CA: IEEE Computer Society Press.
12. Kalawsky, R. S. (1999). VRUSE: A computerised diagnositc tool for usability evaluation of virtual/synthetic environment systems. *Applied Ergonomics*, *30*(1), 11-25.
13. Kaur, K., Maiden, N. A. M., & Sutcliffe, A. G. (1999). Interacting with virtual environments: An evaluation of a model of interaction. *Interacting with Computers*, *11*(4), 403-426.
14. Koleva, B., Taylor, I., Benford, S., Fraser, M., Greenhalgh, C., Schnadelbach, H., Vom Lehn, D., Heath, C., Row-Farr, J., & Adams, M. (2001). Orchestrating a mixed reality performance. In J. A. Jacko, A. Sears, M. Beaudouin-Lafon, & R. J. K. Jacob, (Eds). *CHI 2001 Conference Proceedings: Conference on Human Factors in Computing Systems, Seattle 31 March-5 April 2001*, (pp. 38-45). New York: ACM Press.
15. Mills, S., & Noyes, J. (1999). Virtual reality: An overview of user-related design issues. *Interacting with Computers*, *11*(4), 375-386.
16. Poupyrev, I., & Ichikawa, T. (1999). Manipulating objects in virtual worlds: Categorization and empirical evaluation of interaction techniques. *Journal of Visual Languages and Computing*, *10*(1), 19-35.
17. Slater, M. (1999). Measuring presence: A response to the Witmer and Singer questionnaire. *Presence: Teleoperators and Virtual Environments*, *8*(5), 560-572.
18. Slater, M., Usoh, M., & Steed, A. (1994). The depth of presence in virtual environments. *Presence*, *3*, 130-144.
19. Sutcliffe, A. G., & Kaur, K. D. (2000). Evaluating the usability of virtual reality user interfaces. *Behaviour and Information Technology*, *19*(6), 415-426.