Human performance in space telerobotic manipulation

Philip R. Lamb

University of Canterbury

A thesis submitted in partial fulfilment of the requirements for the degree of Master of Science in Psychology at the University of Canterbury

Supervised by Dr Dean H. Owen

2002

Acknowledgments

I offer my appreciation and thanks to all those who participated in and contributed towards the success of this research. In particular, I would like to thank Dr Dean Owen for both inspiring me to pursue study in the field of Human Factors and for acting as a mentor and advocate for ideas relating to the present research.

I would also like to acknowledge the financial support provided by a University of Canterbury Research Award. I also acknowledge those who assisted at various stages in the development of the simulator, particularly Gino van den Bergen, and Troy Dawson and the OpenGL engineering team at Apple Computer, who provided portions of the code.

Most importantly, thanks to my family and friends for their tireless encouragement, financial support and patience especially during my long periods of absence throughout the months of laboratory work.

Contents

List of figures

List of tables

Abstract

An experiment was conducted to evaluate the potential for improved operator performance in a space-based telerobotic manipulation task when the operator's control interface was based around an egocentric rather than exocentric frame of reference (FOR). Participants performed three tasks of increasing difficulty using a virtual reality-based simulation of the Space Shuttle Remote Manipulation System (SRMS) under four different control interface conditions, which varied in respect of two factors, virtual viewpoint FOR (fixed versus attached to arm) and hand controller FOR (end-effector referenced versus world referenced.) Results indicated a high degree of interaction between spatial properties of the task and the optimal interface condition. Across all tasks, the conditions under end-effector-referenced control were associated with higher performance, as measured by rate of task completion. The mobile viewpoint conditions were generally associated with lower performance on task completion rate but improved performance with respect to number of collisions between the arm and objects in the environment. Increased head movement and higher number of errors in arm motion indicated that the mobile viewpoint suffered from confounding uncontrolled keyhole effects. No correlation between performance and prior 3D simulation experience was observed. There was a significant effect of gender on performance in line with results from the field. The requirement for telemanipulation interfaces to represent critical kinematic limitations in the interface emerges in discussion of origins of performance differences between conditions. The results provide support for the partial application of an egocentric telepresence control interface to space-based articulated manipulators. Different factorings of ego- and exocentric FORs in order to alleviate poor performance under the mobile viewpoint are discussed along with implications for other space-based telemanipulation applications and fruitful approaches to further studies.

1 Introduction

1.1 Robotic manipulation tasks in space

Freed from the usual constraint imposed by gravity, bodies in space may adopt any orientation with respect to each other without internal disturbance. However, in the case of manipulation tasks in space, the orientation of bodies with respect to each other is of prime importance. Manipulation and assembly of complex and large-scale structures under zero-gravity conditions outside Earth's atmosphere are essential elements in ongoing progress in human exploration and colonisation of space.

Remote manipulator systems (RMS) have been an essential enabler of space manipulation and assembly operations. The best known, most successful, and most extensively studied application of a spaceborne RMS is the space shuttle RMS (SRMS), known colloquially as “Canadarm”. The SRMS has been used extensively throughout the 18-year history of the space shuttle program, successfully performing a large variety of manipulation tasks, and has been the subject of a program of ongoing study, upgrades and human-in-the-loop evaluation both within NASA and other research institutes (Nguyen & Hughes, 1994). (An overview of the SRMS system appears in Appendix A.)

However, despite the application of remote manipulator systems and other engineering techniques designed to simplify space-based assembly operations, such as use of prefabrication, modularity, and automated deployment, assembly tasks still require substantial extra-vehicular activity (EVA) by spacewalking astronauts to complete. The seemingly routine nature of travel into space does little to remind us of the extreme hostility of that environment and just how fragile and poor a replica of the terrestrial ecology the “bubble” of the space vehicle and space suit actually are, and EVA is costly, dangerous, and requires highly trained personnel. In the construction of the International Space Station, a new remote manipulator system, the SSRMS, has been deployed and others are planned to substitute further for some of the requirements for EVA in space manipulation tasks, and thus reduce cost, increase efficiency, and allow a broader range of skilled personnel into space (Ruoff, 1994).

1.2 The changing nature of remote manipulation tasks in space

Hitherto, the SRMS has provided the overwhelming majority of proven knowledge in the field of space robotics. However, the established paradigms for control of space manipulators are less well proven in the domain of multi-body assembly tasks, such as the ongoing construction of the International Space Station. Indeed, use of the SRMS in multi-body assembly tasks and the advent of the SSRMS represent the beginning of a new epoch in the tasks to which telerobotic manipulation is being applied, with a drastic increase in the complexity of manipulation operations.

Previously, the SRMS performed tasks involving only the shuttle and a single external body, and could naturally adopt the body of the space shuttle as a fixed plane of reference. However, in multi-body assembly tasks, there is frequently no natural plane of reference. Tasks may involve multiple external bodies, none of which will necessarily have a predetermined orientation relative to the shuttle or any other part of the space station.

Additionally, the SRMS operators rely on a mix of line-of-sight operation out spacecraft windows, closed-circuit camera views, and verbal information relayed from EVA astronauts, whereas in multi-body assembly tasks using the SRMS there is unlikely to be a line-of-sight to the manipulator workspace, and in the case of the SSRMS there is no line-of-sight, operators relying solely on camera views and information relayed from EVA astronauts.

The SRMS *manual* control interface remains functionally unchanged since its 1970s design. In part, this is owes to parsimony: the SRMS has proved effective in the tasks it was originally designed to do. The SRMS and SSRMS control interfaces are both based around resolved motion rate control (Whitney, 1969), which is designed to exploit alignment in orientation of the control coordinate system (the axes in which the operator expresses control actions) and the task-space coordinate system (the axes in which the end-effector moves). Such an alignment is known in human control terms as *direct correspondence*. The astronaut views the manipulator workspace from the aft flight deck, a fixed position relative to the shoulder of the arm. Control commands for motion of the SRMS end-effector are interpreted relative to the axes of the body of the space shuttle. Thus, when the astronaut is able to view the SRMS manipulator workspace out of the shuttle windows, there is direct correspondence between his or her control movements and the resultant visual motion of the end-effector.

However, in using the SRMS and SSRMS in multi-body assembly tasks, direct correspondence is the exception rather than the rule. There is a body of research that indicates a decline in performance and an increase in error rate when direct correspondence is violated.

In studying the coordination between control actions and their displayed effects in a remote manipulation task, Smith and Stuart (1993) varied the position and orientation of the camera in the pick-and-place style task using a master-slave telerobot and concluded that left-right reversal and up-down reversal both caused major performance impairments, as measured by time to complete the task. Similar results emerge from other fields. Holden, Flach and Donchin (1999) studied coordination between movements of a surgeon and a camera manipulated by the surgeon in a simulated laparoscopic surgical task. They found that changes to either the camera orientation or the surgeon's orientation disrupted performance in the pick-and-place task. However, when the position of camera and surgeon changed together, skilled performance was maintained. They suggested that skill in such remote manipulation tasks depends on consistent mapping between the virtual hands and eyes, but not on the particular visual or motor orientations.

Results such as this call into question the utility of the established SRMS operator-manipulator interface. The degree of performance decrement caused by violation of principles of direct correspondence in the SRMS and SSRMS remain unclear, especially since these systems remain relatively untested in multi-body assembly tasks. However what is clear is that such tasks require a great deal of costly training in their preparation, significant ground-based and EVA support resources during their execution, and remain highly awkward for even the most capable operator.

The question is posed then as to what might be done to improve the interface to enable operators to perform remote manipulation tasks for multi-body assembly to the same standard as previous, simpler manipulation tasks. Important dimensions of operator performance include executing desired procedures with minimum time to completion, maximum positioning precision, minimising control effort to reduce expenditure of control-related resources (electrical power, reaction system propellant) and most importantly, avoidance of situations that set the occasion for errors which have potential to compromise the safety of crew or the integrity of the spacecraft, such as collisions between manipulated objects and spacecraft.

1.3 Approaches to the operator-manipulator interface problem

There have been a number of academic and engineering approaches taken towards the operator-manipulator interface problem posed by multi-body assembly tasks. Some of these have sought to compensate for the poor human-machine interface by substituting more automation in place of the human operator. However, application of a total-automation strategy to telerobotic manipulation tasks in space is neither practical nor generally feasible. In the complex sociotechnical system of space exploration, operation is frequently at the limit of known practice, and thus human skill and problem solving form a critical component of the system. Additionally, the dynamic nature of the system is such that it frequently lacks the predictability required for high levels of automation. Thus, the focus here is instead on approaches that offer promise for the improvement of the human-machine interface.

One major approach to overcoming the limitations of traditional human-machine interfaces has been to apply the use of virtual reality technology to the interface. Applications of this approach build "virtual environments" which the human operator explores and interacts with, using now-familiar virtual reality technology such as computer-rendered graphics, head-mounted displays, position trackers, and tactile input devices (Carr & England, 1995). Early advocates of virtual reality technology (e.g. Rheingold, 1991) hoped that in such an approach the human operator would experience a sense of immersion within the given data representation. In cases where the physical layout of a remote location was represented, it was predicted that such an immersion would lead to a sense of *telepresence* (Minsky, 1979, September), that the operator would experience a sensation of being physically present at the remote location.

This prediction has been the implicit motivation for much of the virtual environment research originating from the engineering disciplines (Sheridan, 1992). Yet, Lumelsky (1991) raises a number of practical objections to such an approach. Lumelsky observed that many of the systems actually built had low overall efficiency and awkward interaction between the operator and the machine. He also observed that the teams testing such systems preferred to have the manipulator controls located very close to the master, seemingly contradicting the researchers' claims that these systems provided a suitable platform for teleoperation. In two experiments, he demonstrated superior performance of a motion-planning algorithm over a human operator in a simple two-dimensional motion-planning task. Lumelsky proposed that telepresence interfaces were not supporting their claims, but were rather engaging the operator in continuous and demanding real-time control, despite the operator being ill equipped to perform such control. Lumelsky hypothesised that human operators lack the ability to deal with the task of motion planning of a robot arm, even in a relatively uncluttered environment, and that such a task should be left to automation.

Lumelsky's proposed approach, and indeed most previous research in the human factors of telemanipulation performance, has taken the perspective that in a conditionally-stable system involving both automated control and a human operator, unpredictability and unreliability in the system is attributable to flaws in the performance of the human element. Seeking unconditional stability for the entire system, such approaches treat the human element simply as a subsystem of the larger system (Flach, 1990a) and attempt to identify and model the weaknesses of human element in systems engineering terms so that these may be replaced by automatic control. These approaches propose automation in order to reduce the unpredictability and unreliability of the system.

An alternative strategy, motivated by an ecological approach to human factors (Flach, 1990b; Vicente, 1995) is to consider the both the task and the tools available for interacting with the task as an embedding environment for the human actor, and to take the combined actor-environment system as the unit of analysis. This *ecological* approach recasts the problem, from one of identifying the weaknesses of the human element of the system, to one of discovering the critical sources of information that enable skilled human performance and identifying when these sources of information change, become unavailable, or become unspecific to the aspect of the system under control (Owen, 1990). The ultimate aim of the ecological approach is to reduce the effort of information acquisition and control action from the operator to such a trivial level that the operator is freed from the requirement of continuous control and is able to adopt a supervisory role. This ecological approach to supervisory control is perhaps closer to the original intention of those who conceived the term

(Sheridan, 1994). In this conception, the role of the human does not just *change* from direct operation to planning and evaluating operations of lower-level automation, but rather these activities – planning and evaluating – *emerge naturally* as a consequence of reduced operator workload.

1.4 Application of an ecological approach

In the case of the astronaut operator performing telerobotic manipulations using the SRMS, there has to date been little research identifying the critical sources of information that enable skilled performance. Yet, knowledge of these sources is both an essential first step in designing any automated system and essential in deciding how best to make sure these sources of information remain available to the astronaut operator in the more challenging environment of multi-body assembly and manipulation tasks.

Reaching for, grasping, and manipulating objects in the environment is an activity fundamental to most human action. However, reaching, grasping, and manipulation tasks performed telerobotically in space with the SRMS are very difficult, requiring extensive training and costly support both from the ground and from EVA personnel. Therefore, there is a poor match between the astronaut operator's natural skill in these activities and the support provided by the SRMS control interface for expressing control actions in the same terms. For the astronaut engaged in a telerobotic manipulation task in space, the perception-action interface with the environment is at variance in a number of ways with the perception-action interface typically experienced when engaged in a direct manipulation task on earth using the hands.

In any remote manipulation task, displays, controls, and the remote manipulator itself become an extension of the actor's ego (Sheridan, 1994). Although the SRMS display and control systems provide information to the astronaut about the environment, this information is mediated and transformed by properties of the cameras and monitors. Extra information may be available too, through sensors such as temperature gauges, warning lights and the like. Similarly, the astronauts control actions are mediated and transformed by properties of the control hardware, software, and actuators.

Additionally, the zero-gravity space environment creates conditions that are not typically experienced on earth. Some sources of information are unavailable, for example information from surface interactions or collisions that are usually conveyed by sound. Other information becomes informative about a different aspect of the environment than is usual for the astronaut when normally resident on earth. For example, straight-line motion of a body in space has different physical meaning to the same straight-line motion on earth, owing to the different dynamics of zero-gravity.

Thus, in order to successfully perform remote manipulation tasks with the present SRMS, there are a number of new potentially informative sources which an inexperienced operator must become attuned to and a number of old potentially uninformative sources to learn to ignore. Yet requiring suppression of natural perception-action skill in activities so fundamental as reaching for, grasping and manipulating objects, constitutes a significant barrier to acquisition of skill with the SRMS and the successful performance of telerobotic manipulation tasks in space.

1.5 Overcoming complex spatial relationships

At the heart of the variance between the SRMS control interface and the everyday environment in the affordance of reaching, grasping and manipulating is the spatial relationship between the human operator and the manipulator. The results from Smith and

Stuart (1993) indicate that there must be spatial coordination between the operator's control actions and their viewed effect. As suggested by Holden, Flach and Donchin (1999), changes in camera view must be accompanied by a coordinated change in control axis.

In the present SRMS deployed on the Space Shuttle, although the SRMS operator is able to select between several different camera views, the control frame of reference is not linked to the selected camera view because this would produce a number of inconsistencies. For example, if the operator were making control actions to the arm to move the end-effector in a straight line, and while doing so switched to a different camera perspective, the control positions the operator was holding at the time of the switch would be referenced to the new frame provided by the new camera perspective. The direction of end-effector motion would thus change at the time of the camera switch. Linking control frame of reference and camera view in this way is thus undesirable.

One potential solution to this problem is to link not only control frame of reference to changes in viewpoint, but also the reverse: to link movement of the viewpoint to control action itself. In the case of the SRMS, this would suggest that the operator's viewpoint should be fixed to and move with the manipulator arm. In this way, the operator's control actions would change both the viewpoint and the control frame of reference in a coordinated manner.

Whereas at present in the SRMS, the operator's physical space must be adjacent to the manipulator arm's physical space, the use of virtual reality technology to create a virtual environment allows the operator's physical space to be arbitrarily located. At present, the operator is informed about the state of the manipulator's physical environment via viewing directly through the windows, via closed circuit television (CCTV), and via sensor displays (Figure 1). Recreating this situation in a virtual environment, the operator's line of sight is no longer available, and the operator is informed via CCTV and sensor displays only. The manipulator space, thus viewed, is *real* in the sense that it is a true representation of the manipulator environment. It is however *virtual* in the sense that the information on which the astronaut acts is mediated and transformed by the computerised head-mounted display.

It is clear that with the creation of a virtual manipulator space, the spatial relationship between the operator's physical space and the manipulator's virtual space may be selected arbitrarily. There are a number of potentially useful configurations of this relationship. Of particular interest are the *exocentric* case in which the manipulator virtual space is viewed from a fixed external location, and the *egocentric* case in which the manipulator virtual space is collocated with the operator's physical space. [See McCormick, Wickens, Banks, & Yeh (1998) for a fuller discussion of exo- and ego-centric frames of reference.]

In the exocentric configuration, the participant views the SRMS workspace from a fixed position with respect to the shoulder of the arm. Head movements within a normal range for a seated person allow a small degree of movement about the fixed viewpoint. To draw a parallel to the present-day situation aboard the shuttle, the exocentric case is analogous to a situation in which the SRMS operator is seated at a control station on the aft flight deck and is viewing the manipulator workspace through the payload bay windows.

In the egocentric configuration, the participant views the SRMS workspace from a mobile position tethered to the manipulator end-effector. Again, head movements within a normal range for a seated person allow a small degree of movement about this tethered position. To draw a parallel to the present-day situation aboard the shuttle, this is loosely analogous to a situation in which the SRMS operator is seated at a control station on the aft flight deck and is viewing the manipulator workspace through the camera mounted on the manipulator end-effector.

An exocentric virtual environment preserves the present-day real-life spatial relationship between the operator and the manipulator. In an egocentric virtual environment, the manipulator virtual environment would be located such that the position and orientation of the operator remains fixed relative to the end-effector. In the latter situation, control translations and rotations would cause an equal and opposite change in the position and orientation of the manipulator virtual environment relative to the operator. It is hypothesised however that the net phenomenal experience of the operator would not be motion of the virtual environment about him or her, but that he or she is in fact translating and rotating through a fixed and stable environment.

Thus, whereas an exocentric virtual environment would create a situation in which the SRMS control task is analogous to a reaching and grasping task performed with the hands, an egocentric virtual environment may effectively change the analogous task to control of self-motion.

In a sense, the empirical investigation of the utility of different spatial relationships between operator and task is a generalisation to three dimensions of the results from studies of control-display compatibility in two dimensions.

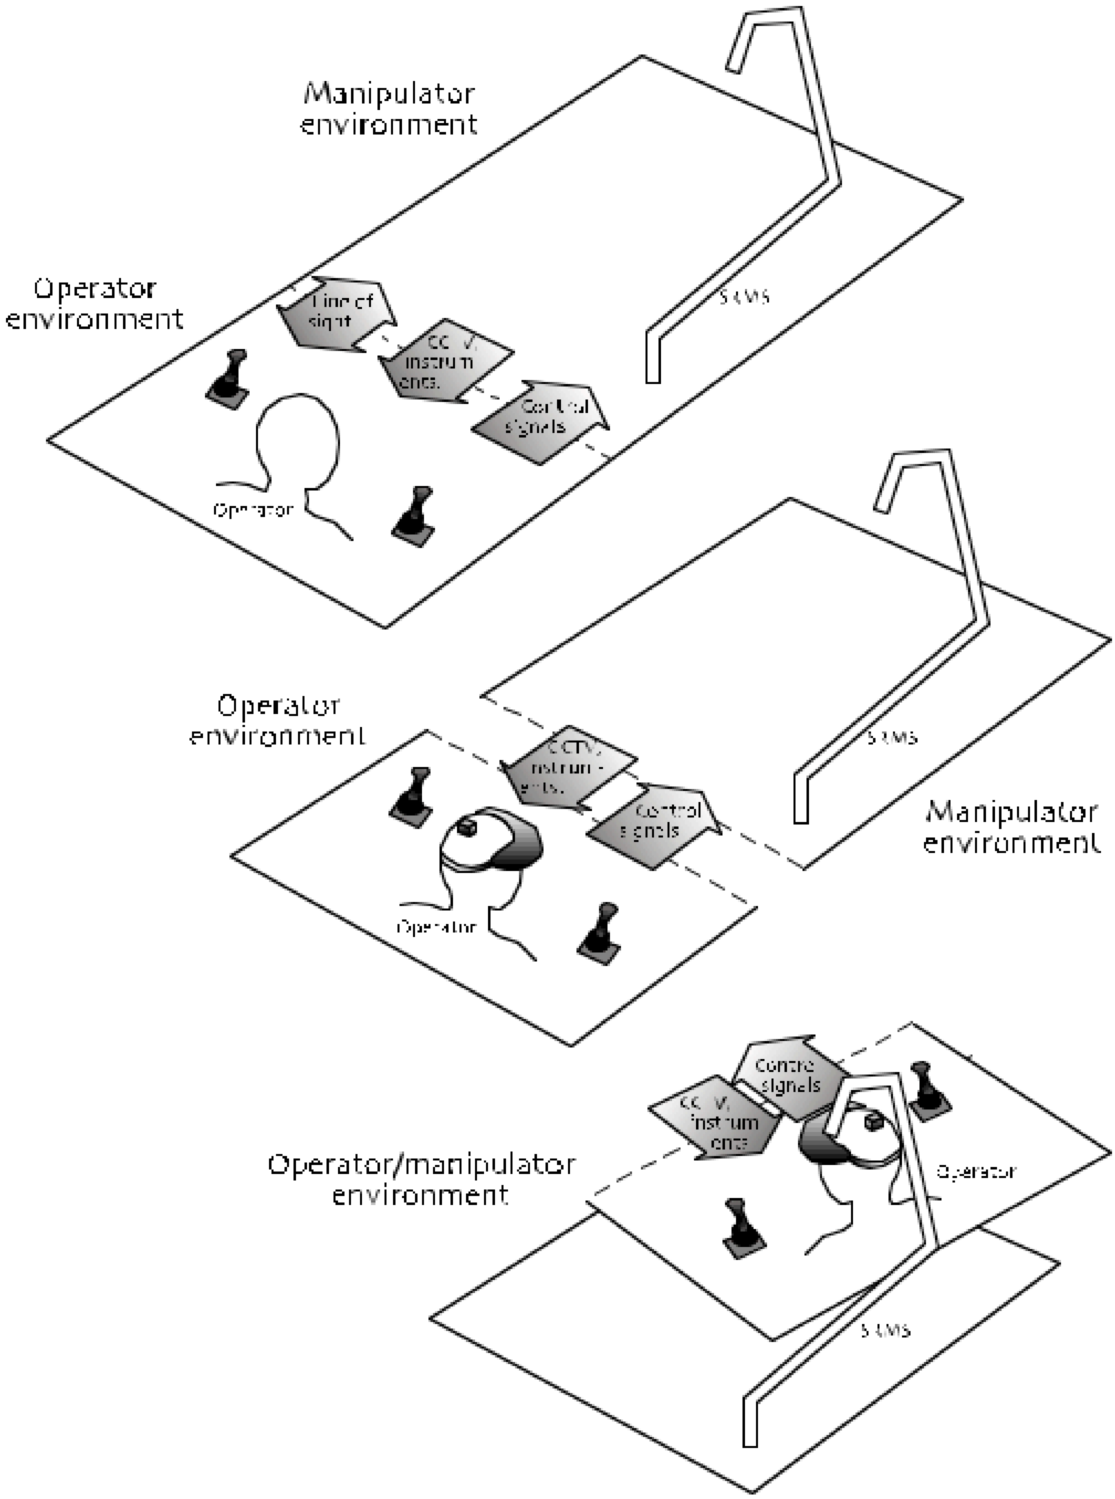

Figure 1. Relationship between operator and work environment. In the actual space shuttle (top), the operator is informed about the state of the manipulator environment via windows, CCTV, and sensors. In the exocentric virtual environment (middle), this information is mediated by a head-mounted display, but the spatial relationship between operator and manipulator is preserved. In the egocentric virtual environment (bottom), the operator environment is collocated with the manipulator environment.

1.6 Experimental approach

The aim of the present research was to determine the effects of varying certain aspects of the human interface to a telerobot on the performance of the human operator using the robot to perform a manipulation task under zero-gravity conditions in a space environment. The particular telerobot considered was the space shuttle robotic manipulation system (SRMS).

To investigate this empirically, the experimental hypothesis was that manipulation tasks performed telerobotically would show improved operator performance when the operator's actions in the work environment were made with respect to an egocentric frame of reference rather than an exocentric frame of reference. Furthermore, it was hypothesised that the performance advantage of the egocentric frame of reference would increase with increasing spatial complexity of the manipulation tasks.

To evaluate these hypotheses, an interactive simulation was designed to replicate a space-borne work domain similar to the space shuttle RMS. The simulation, implemented using virtual reality technology, provided both exocentric and egocentric operator environments and a representative set of zero-gravity robotic manipulation tasks. The simulation served as a platform for empirical evaluation of the research question and for additional investigation of other human factors and cognitive ergonomics issues that arise in the application of a telepresence system to zero-gravity robotic manipulation tasks in a space-borne work domain including, but not limited to, simulator sickness effects.

The overarching aim was to draw conclusions with respect to understanding human performance in teleoperation, theories of telepresence, the design of interfaces to telerobotic systems, and understanding human perception and action in manipulation tasks of high spatial complexity.

1.7 Independent variables

As the purely exocentric and purely egocentric operator environments differed in more than one dimension, and as each was potentially of different utility for different manipulation tasks, experimental comparison of operator performance in each environment was broken down into three independent variables. These were evaluated in a factorial experimental design:

1. Viewpoint frame of reference uncoupled versus coupled to manipulator end-effector motion. This factor varied the way in which the participant viewed the SRMS workspace. In the *fixed* category, the participant viewpoint remained stationary near the forward bulkhead of the cargo bay, a short distance from the shoulder of the manipulator arm. In the *mobile* category, the participant viewpoint was located near the manipulator end-effector and moved with movements of the manipulator arm so as to maintain a fixed distance and orientation with respect to the end-effector.
2. Control frame of reference aligned with body of shuttle versus aligned with manipulator end-effector. This factor varied the coordination between the axis of movement of the hand controllers and the corresponding axis of translation and/or rotation of the manipulator end-effector. In the *world-referenced* category, hand controller movements were aligned with the body of the space shuttle. In the *self-referenced* category, hand controller movements were aligned with respect to the current orientation of the manipulator end-effector.
3. Task difficulty. This factor, through variation of the position of the object to be grasped, varied the degree to which the task showed (a) loss of a natural reference plane in the environment (e.g. level ground), (b) greater distance between the start and end points of the manipulation and (c) greater change in orientation of the objects' axes between the start and end points of the manipulation.

The combinations of experimental factors (1) viewpoint frame of reference and (2) control frame of reference created four distinct experimental conditions, as indicated in Table 1.

Table 1. Names of Experimental Conditions.

		Factor (1) Viewpoint frame of reference:	
		Fixed near SRMS shoulder.	Mobile attached near end-effector.
Factor (2) Control frame of reference:	Aligned with body of space shuttle.	Fixed, world-referenced.	Mobile, world-referenced.
	Aligned with arm end-effector.	Fixed, self-referenced.	Mobile, self-referenced.

2 Method

2.1 Participants

45 participants were recruited from the student population at the University of Canterbury, 35 males and 10 females. Ages ranged between 17 and 36 years, with a median age of 20.

It was expected that a small proportion of participants would find the task prohibitively difficult due to its demanding requirements on motor skill and spatial perception/action coordination. It was also expected that some participants might be susceptible to simulator sickness effects to a degree that would preclude them from completing the session. The experimental session was structured so that participants who fell into these categories could be objectively identified during the initial training phase of the experiment and could be withdrawn from the experiment at the completion of the training phase, without undue offence or disappointment to the participant. (In the event, no participants were withdrawn at the juncture of training and testing, however 6 participants withdrew at various stages during the test phase owing to the onset of intolerable simulator sickness effects. The data of one further participant was withdrawn after completion of the experiment due to a malfunction in the automatic data recording during the experiment.)

As compensation for the time involved in participating in the study, participants were paid $10. Participation complied with requirements of the University of Canterbury Human Ethics Committee.

2.2 Materials and apparatus

2.2.1 Interactive simulator

To provide a test bed for different control and display configurations, an interactive simulation of an SRMS control station and workspace was developed by the experimenter. A schematic of the equipment implementing the simulation appears in Figure 2.

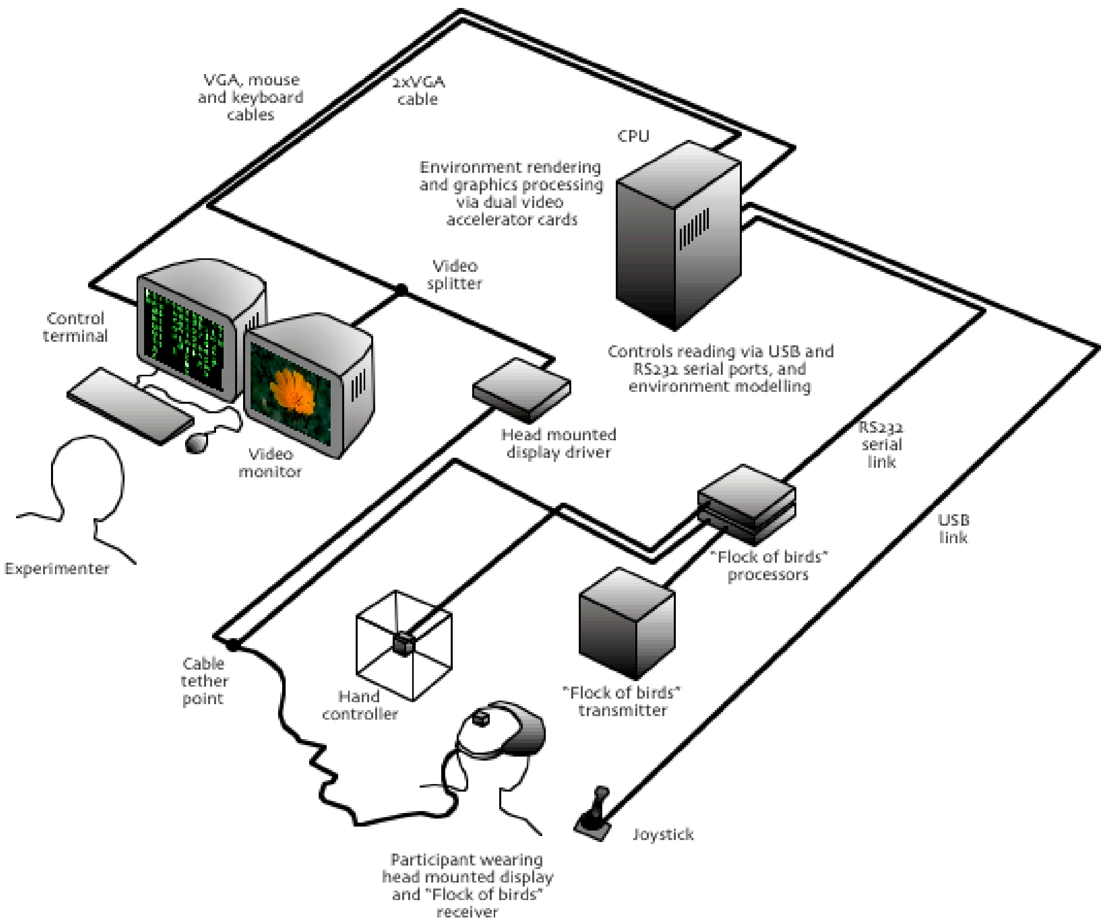

Figure 2. Schematic of equipment used in implementing the simulation.

A three-dimensional image of the simulated environment was presented to the participant through a tethered stereoscopic head-mounted display (HMD.) The HMD was a Virtual Research Systems model V8, and contained two 3.3cm 640x480 pixel active-matrix fullcolour liquid crystal display elements mated with optics to produce an image focused at optical infinity. The image presented occupied a 60° field-of-view diagonally and had 100% stereo overlap. The refresh rate of the displays was 60 Hz.

A six-degree-of-freedom tethered Ascension Corporation "Flock of Birds" system tracked the position and orientation of the participant's head in the actual environment. Position and orientation information was read from the head tracker at a rate varying between a minimum of 60 and a maximum of 100 samples per second, depending on the computer CPU load.

The participant issued position and orientation control actions through two displacement hand controllers. The left-hand controller, used for positional input, employed a second Flock of Birds six-degree-of-freedom tracker together with a supporting piece of equipment constructed by the experimenter, consisting of a enclosing cradle for the sensor suspended in the centre of an open cubic lattice of 30 cm each side, with a centre-return force provided by elastic bands. The controller is shown in Figure 3. The participant grasped the cradle and could independently apply a displacement by pushing or pulling it forwards, backwards, up, down, left or right. When released, the cradle returned to the central resting position. Only the displacement of the sensor from its resting position was read. (The orientation of the sensor, although available, was not used.) The hand controller's tracker was connected to the computer via the same serial I/O channel used for the head tracker.

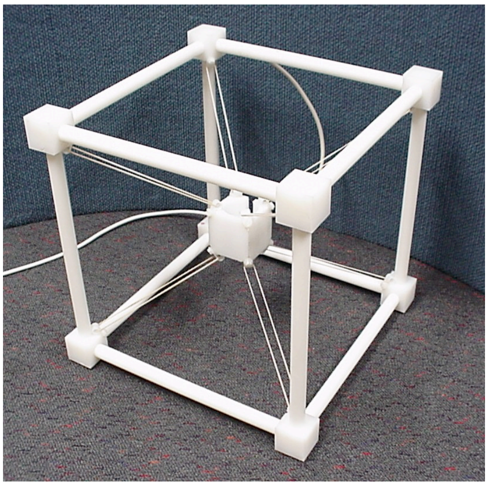

Figure 3. Oblique view of left-hand controller

The right-hand controller was an off-the-shelf three-degree-of-freedom joystick, Logitech Corporation model Wingman Extreme Digital 3D. It provided three rotational degrees of freedom: forwards-backwards, left-right, and a twist movement about the axis of the stick.

These movements corresponded to pitch, roll and yaw commands respectively. When released, springs provided a restoring force returning the stick to the resting position. The joystick connected to the computer through a Universal Serial Bus (USB) connection.

The hand controllers were affixed to a standard office desk and this, combined with a swivel roller chair formed the console at which the participant was tested in the experiment. The participant's console is shown in Figure 4.

Figure 4. Simulator equipment located at the participant's console

The virtual environment was modelled using software written in C. A data flow diagram representation (Pressman, 2001) of the modelling software appears in Appendix C. The modelling software performed a range of tasks including processing of participant control inputs, processing of head position and orientation, calculation of arm inverse kinematics, task initiation, virtual object collision determination and dynamics, monitoring for task completion or termination, and participant virtual viewpoint determination.

The virtual environment was rendered using software written in C, which issued commands to a standard hardware-accelerated OpenGL 1.2 applications programming interface (API). A data flow diagram representation of the rendering software also appears in Appendix C. Tasks performed by the rendering software included storage of graphics models and textures, creation of OpenGL graphics primitives from virtual environment object positions, and rendering of graphics primitives from the participant's virtual viewpoint.

Two consumer-level ATi Corporation Rage 128 graphics cards provided graphics acceleration and display. Each card's VGA output was split, with one output driving one eye of the head-mounted stereo display and the other output available for the experimenter to view on a monitor.

The software was compiled for the Mac OS X operating system and ran on a single-processor PowerPC G4-based central processing unit (CPU). Modelling and rendering operations ran consecutively each frame in one process, and I/O operations and coordination of the initiation and termination of each trial ran in separate processes. An average frame rate of 75 full-stereo frames per second (FPS) was obtained, with the minimum frame rate greater than 60 FPS even during heavy computational loading from the modelling operations. This figure is well above minimum recommended 10 FPS suggested by Liu et al. (1993).

2.2.2 Simulated environment

The virtual environment was designed to replicate the salient features of the SRMS and its workspace. It consisted of correct-scale three-dimensional models of the space shuttle exterior, payload bay and manipulator arm, a visual control aid, a payload, and the backdrop.

The space shuttle exterior, payload bay and manipulator arm were all modelled as non-intersectable solids. In the case of a collision between any of the elements, a brief auditory warning was sounded, further arm motion was suspended, and then the arm backtracked for 3 s along the path taken 1.5 s immediately prior to the collision. A further beeping tone sounded during the backtracking manoeuvre.

The manipulator arm itself was a kinematically correct model of the actual SRMS. The main difference from the real SRMS was that it did not conform to joint-rate limits and the end-effector was able to move at a maximum translational rate of 1.5 m/s and a maximum rotational rate of 45 °/s, laden or unladen. (The maximum laden and unladen translational rates of the actual RMS, 0.06 m/s and 0.6 m/s respectively, were found to be too slow to allow a reasonable number of arm movements to be performed in the time available for conduction of experiment.) Dynamic properties of the arm (joint backlash, flexibility and hysteresis, and boom elasticity) were also neglected. Notwithstanding, the arm could be manoeuvred both to the edge of its reach envelope, to its joint limits, and into configurations where one or more joints approached singularity. In case of entry into one of these forbidden configurations, a different brief auditory warning was sounded, further arm motion was suspended, and then the arm backtracked for 3 s along the path taken 1.5 s immediately prior to the entry of the forbidden arm configuration. A further beeping tone sounded during the backtracking manoeuvre.

A visual aid was present in the virtual environment to provide support for determining the orientation of the hand controller axes with respect to the arm end-effector. This visual aid consisted of three lines arranged in a right-hand coordinate system and originating from the point of resolution (POR) upon and about which the controls acted. The lines were colour-coded red for the left-right / pitch axis, green for the up-down / yaw axis, and blue for the fore-aft / roll axis. The lines were 3 m in length in the virtual environment. This aid was included after a pilot experiment indicated that untrained participants had considerable difficulty understanding the effect on the controls of changing between the different control frames-of-reference. With the aid, participants could be trained to a level where they could adequately predict end-effector motion within the 15-20 minutes available for training.

The payload itself consisted of a model of a satellite, made up of a gold-textured cylinder, 3.5 m long and 1.5 m in diameter. A raised plate of the same dimensions as the end face of the manipulator end-effector was mounted in the centre of one of the end faces of the payload. It was to this plate that the end-effector was to be brought and aligned. The desired target position of the payload, which remained static during across all trials, was indicated to the participant by a semi-transparent replica of the payload model.

The backdrop consisted of an authentic Earth model that the shuttle orbited once every 56 min, and an overhead Sun that provided directional illumination. As well as adding some visual realism, the backdrop helped partially alleviate the unwanted HMD “keyhole” effect where a participant wearing a HMD has difficulty locating the imagery of interest against a uniform black background. However, in order to prevent the backdrop being used as a reference plane for the manipulation tasks, the shuttle also rotated on its longitudinal axis at a rate of 1 deg/s. This was the fastest rate that did not distract from the main visual material and was determined by a trial and error process.

All imagery was displayed in stereo, shaded, and textured to reduce depth ambiguities. The stereo imagery was produced using an asymmetrical viewing frustum, with the centreline of the two eyes converging at a distance of 15 m. The eye separation selected was 8 cm. An example of the imagery produced by the simulator appears in Figure 5.

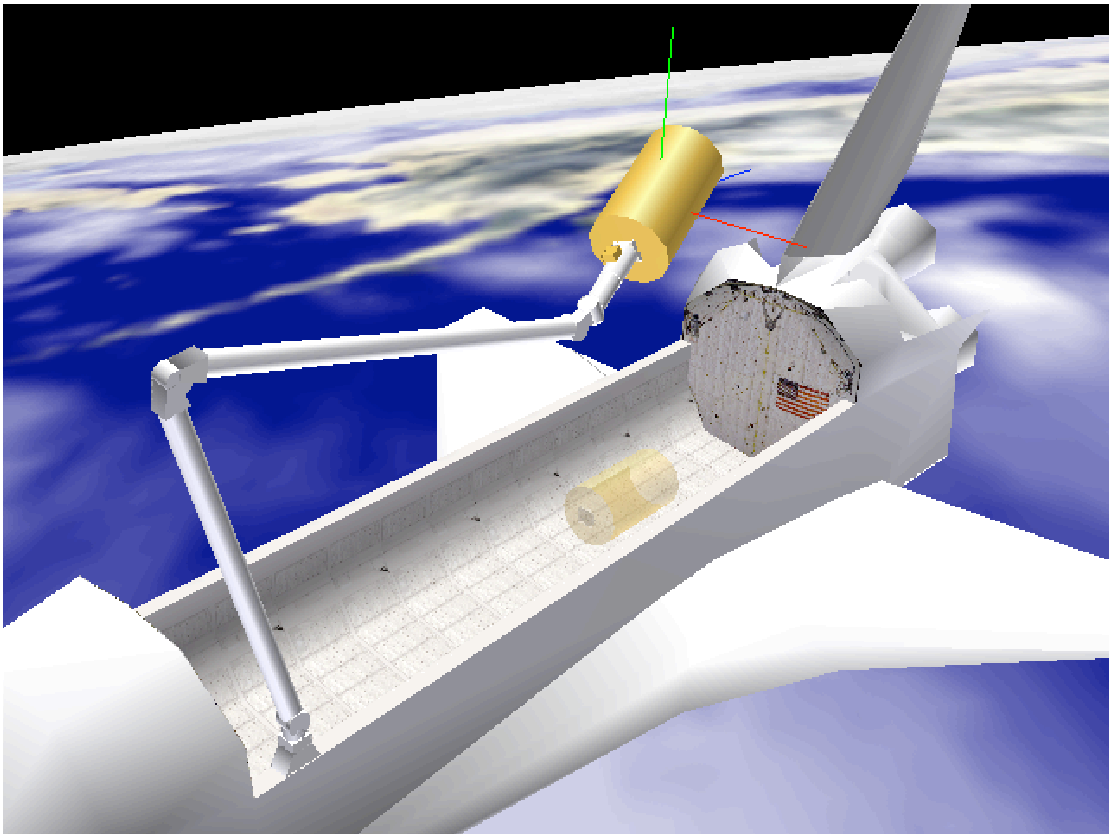

Figure 5. Example simulator imagery showing an overview of the virtual environment during a payload manoeuvre.

2.2.3 Manipulation tasks

Each manipulation task in the experiment was essentially a three-dimensional pick-and-place task. The unloaded manipulator arm began from a rest configuration aligned with and just above the port-side longeron. Participants were required to move the manipulator end-effector from its initial location and orientation to bring it to and aligned it with the payload grapple fixture. At that point, the payload became automatically latched to the arm, an audio message confirming the latch was played, and the arm was loaded. Participants were then required to move the loaded arm so as to overlay the payload on top of the translucent payload target. At all times, collisions between any parts of the arm, the payload and the shuttle body were to be avoided, and movement of the arm to place it in a forbidden configuration were also to be avoided. Participants were instructed during the training phase to deliberately make collision and forbidden-configuration errors so that they would know what to expect in case of error during the testing phase.

Although visual fidelity between the SRMS simulation and the actual Space Shuttle was desirable, the same did not apply to the task requirements of typical on-orbit SRMS operations, since most of these were considered peripheral or irrelevant to this study and were therefore neglected. Neglected SRMS task requirements included manoeuvring the shuttle so as to place a free-flying satellite to be grappled within the reach envelope of the RMS, opening of payload bay doors, release and swing-out of the SRMS from its stowed position, illumination of the payload bay, and powering up of the SRMS and control console, and selection of the correct SRMS operating mode. In the manipulation tasks in the experiment, it was implicit that the shuttle was on orbit and at the correct altitude, velocity and attitude with respect to any free-flying satellite payload to be grappled. A final neglected aspect of typical on-orbit RMS operations was management of payload systems and subsystems. For the purposes of this experiment, the payload was treated as an inert mass.

The training phase of the experiment (see § 2.3 below) presented one distinct manipulation task, then participants encountered a further three distinct manipulation tasks during the test phase. Each task was differentiated by a different initial payload position, with the initial resting position of the manipulator arm and the position of the final payload target remaining the same throughout the experiment. Each task therefore had a different level of difficulty due to the difference in complexity of the motion of required to move the manipulator arm from its resting position to the payload position and the laden arm to the payload target. The tasks were presented in order of increasing difficulty. (Although the difficulty can only be determined posteriori by an appropriate comparison of participant performance between tasks, the difficulties were ordered based on the findings of the pilot study.)

The initial payload positions for each task are displayed in Figure 6. Task T was used only in the training phase of the experiment. Tasks 1-3 were used only in the test phase.

T. The initial payload position was directly above the target, rotated 45° about its vertical axis.

1. The payload was located forward and above the target off the port wing leading edge, rotated 45° about its vertical axis.

2. The payload was located near the payload bay centreline and above the forward bulkhead, rotated 90° about its horizontal axis.

3. The payload was located off the starboard wing just aft of the forward bulkhead, and rotated 135° about its vertical axis and titled downwards approximately 30°.

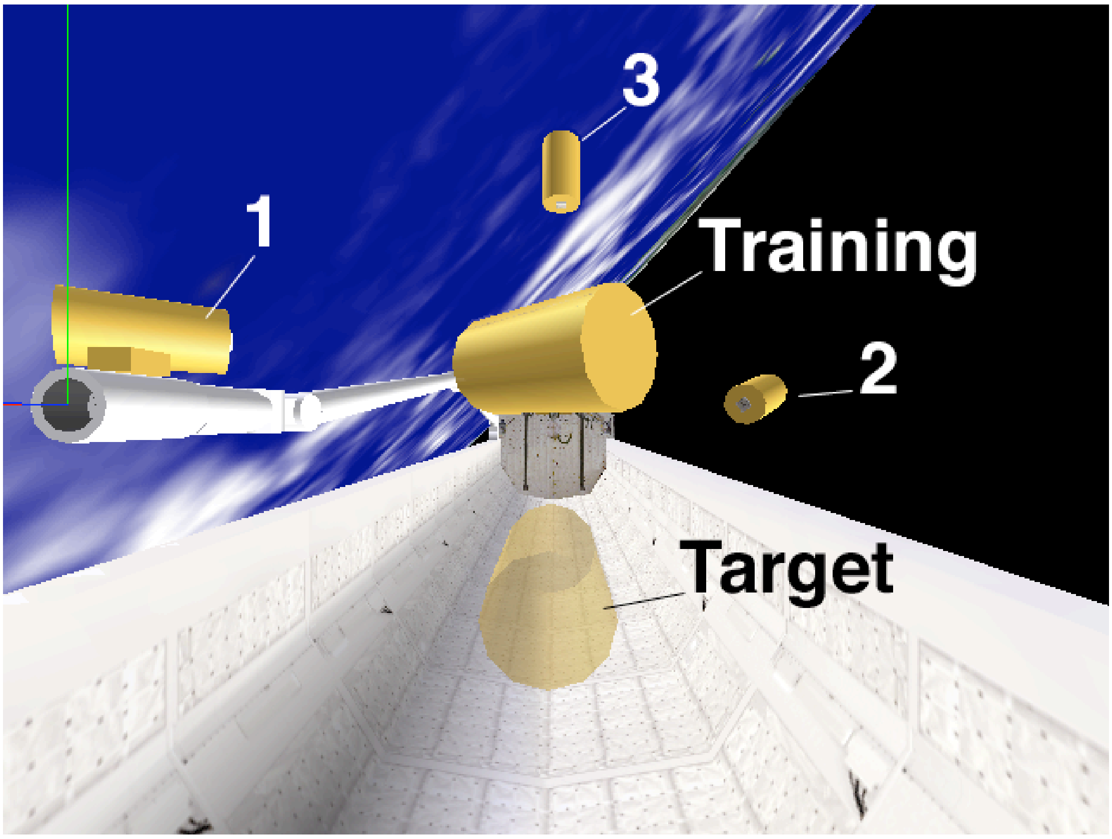

Figure 6. Composite view of the initial arm position, payload locations for training task and test tasks 1-3, and payload target position, viewed from the rear payload bay bulkhead using an artificially wide-angle lens.

2.3 Procedure

2.3.1 Training phase

Each participant read an information sheet giving a brief outline of the experiment's purpose and procedure, filled out a brief questionnaire eliciting some demographic information, completed the pre-experiment Simulator Sickness Questionnaire and signed an informed consent sheet. (See Appendix D for examples of these materials.)

Each participant initially undertook a period in the simulator that served as a familiarisation, training, and initial selection phase. The participant was seated comfortably at the control desk and introduced to the equipment. Operation of the hand controllers and the range of movements that they effect was explained and demonstrated, and the head-mounted display was fitted and adjusted to satisfaction. The participant was then given an automated flyover of the workspace. During the training phase, the participant was free to ask the experimenter questions on the task requirements and the operation of the manipulator.

The participant was trained to use the hand controllers to move the arm to accomplish a simple manipulation Task T, as detailed in § 2.2.3 above. The participant repeated the same training manipulation task under each variation of experimental factors (1) (viewpoint frame of reference) and (2) (control frame of reference) twice, making a total of eight repetitions of the task. In order to pass the initial selection and continue with the second phase of the experiment, the participant was required to take no longer than 3 min to complete each of the last four repetitions. Participants were advised to work as quickly as possible but were not advised of a specific time limit. This was designed to eliminate participants who found the task prohibitively difficult, as well as those with strong susceptibility to simulator sickness. In the experiment, this criterion did not eliminate any participants.

Following the completion of the training phase, participants took a break for 2 to 5 min, during which they could remove the HMD.

2.3.2 Test phase

On beginning the test phase, each participant was instructed that he or she would be given three more increasingly difficult manipulation tasks to be performed, with each task repeated with the four different configurations of displays and controls that he or she had already been exposed to. The participant was advised that there would be a 3-min time limit on each trial and that he or she should work as quickly as possible. The participant was also advised that the experimenter would not answer questions or provide further advice about the task until the completion of the experiment.

The participant repeated each manipulation task under all four variations of the experimental factors: (1) viewpoint frame of reference and (2) control frame of reference, before moving onto the next level of difficulty. Thus, each participant performed a total of 12 trials in the test phase.

At the beginning of each trial, the participant was informed verbally by the experimenter and by an onscreen message for 10 s which variation of factors (1) and (2) they would be operating under for that trial. At successful completion of the task or if the 3-min time limit expired, there was another 10-s pause, meaning a minimum of a 20-s pause before the start of the next trial.

2.4 Performance measures

Participant performance measures were designed to allow extraction of information relevant to the performance requirements of a typical real-world on-orbit RMS task. These requirements and the measures selected are listed in Table 2. These constitute the primary dependent variables of interest in the experiment.

Table 2. Performance Measures.

Task performance requirement	Task performance measure(s)
Minimise time to complete task.	Total elapsed time from first control action, up to fulfilment of task completion or failure criteria.
Maximise efficiency of path taken during manoeuvre.	Root mean squared (RMS) value of distance between point-of-resolution (POR) on manipulator and target throughout manoeuvre multiplied by duration of manoeuvre.
Minimise manipulation errors.	Count of number of collisions between manipulator and payload, or manipulator and orbiter, or payload and orbiter, and count of number of instances in which the manipulator is placed in forbidden configurations such as singularities or at its reach limit.
Minimise control effort.	RMS value of control excursion in all axes throughout manoeuvre multiplied by duration of manoeuvre.

With reference to the table, the RMS value of x multiplied by the duration of measurement is directly equivalent to the time integral of the absolute value of x, which is a scalar measure of *work* done by x. Work is given by the function

$$w = \int_{t_0}^{t_n} |x| \, \mathrm{dt} \quad \text{......} \quad \textbf{(Equation 1)}$$

where $[t_0, t_n]$ is the interval over which measurement occurs. For discrete time systems, the above equation becomes

$$w = \sum_{k=0}^{k=N-1} |x_k| t_d \quad \text{.......... (Equation 2)}$$

where N is the number of samples and t_d is the sampling interval, with $t_d \times N = t_n - t_0$.

In order to produce data for the above performance measures, a number of measures were recorded in real-time:

- Elapsed time (since start of trial).
- Viewpoint position and orientation.
- End-effector position and orientation.
- Payload latch state.
- Control excursion (from centre point).
- Elapsed number of collisions.
- Elapsed number of forbidden arm configurations entered.

From these, any number of secondary measures could be derived. Basic secondary measures derived are presented below and more complex measures are detailed in the results.

The following measures were extracted from data gathered during the unladen phase of each trial:

- Straight-line distance between manipulator end-effector and grapple lug on payload.
- Solid angle between major axis of manipulator end-effector and major axis of grapple lug on payload.

From data collected during the laden phase of each trial, the following measures were derived.:

- Straight-line distance between centroid of payload and centroid of payload target.
- Solid angle between major axis of payload and major axis of payload target.

2.5 Design

The experiment was a factorial design, where combination of factors (1) viewpoint frame of reference and (2) control frame of reference created four distinct experimental conditions, which were repeated in blocks across different levels of factor (3) task. The first two blocks used Task T and constituted the training phase. The last three blocks used Tasks 1, 2 and 3 and constituted the test phase.

Within each block of four trials, there were 24 possible combinations of the ordering of the four conditions. For the initial training task, participants were randomly assigned to a given ordering of conditions, and the orderings were counterbalanced between participants. The within-participant ordering of conditions also changed between tasks. As the total number of combinations required to guarantee complete orthogonality (24^5) was prohibitively high, a systematic ordering was adopted which minimised the number of participants beginning more than one task with the same condition, thus:

Participant number (modulo 24)	1	Order	1,	13,	7,	19,	4
	2		2,	14,	8,	20,	5
	..		..	..	..	..	..
	n		n,	n+12,	n+6,	n+18,	n+3

where a given order number refers to a one of the 24 particular orderings of the four conditions within one task.