

Haptic Tele-operation of Industrial Equipment

Chris Gunn* and DingYun Zhu*
ICT Centre, CSIRO, Canberra, Australia.
{chris.gunn, dingyun.zhu}@csiro.au

Tele-operation involves taking a worker a considerable distance away from the machine-material interface and providing them with controls that allow the operation to happen over a network communications link. This can involve distances from a few metres to thousands of kilometres. Tele-operation can provide the advantage of providing the worker with a more comfortable, safer work environment, as the tele-operation control centre can be chosen independently of the location of the material being worked. One challenge in this field is to provide a worker with similar observation capabilities to that which exists on site, allowing them to interact with precision and safety. For this, it helps to provide good vision and auditory awareness, but also some form of touch or force response can be beneficial. With the addition of haptics, it is possible to utilise the operator's sense of touch in controlling equipment, whether large industrial machinery or small laboratory equipment. Another challenge is to compensate for any latency that is introduced by the network, as such latencies can mean that the operator only observes a delayed view of the work face.

We present a three dimensional, haptic interface for controlling a mechanical arm. The arm can act as a placeholder for any multi-linked industrial device since the interface is generic enough to be applied to any stationary tele-operated industrial arm. We show that the use of haptics can re-introduce the sense of touch of a tool on objects that is often absent with joystick-style controls.

Introduction

Before the invention of tools, humankind manipulated objects through the use of their hands, feeling the properties of the materials through the sense of touch and resistance to deformation. We understand some of the properties of a material by visually observing its behaviour under manipulation, by feeling the pressure of the surface on our skin and by comparing the movement of our fingers and hands on it, relative to the forces we are applying.

With the discovery of tools, this direct sense of touch of object-on-skin was removed, leaving in its place only the sense of resistance to movement: how much the hand (and a tool) moved relative to how much force the user was applying to it. The human brain can determine the success of the tool action by comparing the visual appearance before and after the action, but more precise manipulation is possible if the proprioceptive knowledge of the hand's position and motion is used as well. A craftsman can feel the wood paring away under a woodworking plane and judge the depth of the cut without having to visually observe it. Adding vision improves accuracy.

Powered tools further isolate the worker's senses of force and touch from the work surface. The mediation of motors, levers, cables and pullies decreases the sense of touch, often introducing other forms of resistance (such as internal friction) or artificial compliance (for example power assisted brakes). Connections such as these can shield the worker from unpleasant work environments and give them powers beyond that of human abilities. However, they can also reduce the ability to sense and understand what is happening at the tool-material interface. A power saw operator can cut through surfaces without any real knowledge of the material's properties. Since the energy to produce the cut is typically coming from the electrical power supply, not the worker's muscles, a direct, real-time understanding of the tool-surface interaction can be mostly absent.

However, the benefit of having greatly increased power available is seen as a worthwhile trade-off against the loss of direct tool-surface experience.

Tele-operation further removes the worker from the work materials. Typically tele-operation involves locating the worker at some location away from - and out of site of - the workplace. To compensate for this lack of vision, cameras and video transmission equipment are often installed. Controls, such as levers and joysticks are provided which relay the operator's hand actions to the distant machine, via a communications network. Both the communications network and the video stream encoding and decoding introduce a delay, or latency, between what the operator sees, and what is actually happening at the time. Also, video vision is of poor quality (even with high definition) compared to direct line-of-sight. The sense of touch and resistance is usually missing, further reducing the operator's understanding of the material and actions being undertaken.

Re-introducing the sense of touch and material resistance into a tele-operation scenario is now possible with the advent of haptic user interfaces. In this work, we investigate the use of a haptic user interface to a mining task, examples of which are rock breaking (figure 1), tuyere punching and accretion cleaning in metallurgical furnaces, explosive placement and rock drill operation. A prototype system was developed which could provide the basis of future tele-operated mining systems.



Figure 1. A mining rock breaker

Tele-operation in Mining

Despite the problems with tele-operation, there can be several compelling reasons to operate equipment from a distance. The mining industry in Australia has a particular interest in tele-operation in an effort to provide a more comfortable and safe environment for the mine workers. Many of the Australian mines are located in very remote, dry, dusty and hot locations. Because of this, the fly-in-fly-out (FIFO) system is in operation in many mines. This involves flying workers in to a mining camp from one of the cities for a period and then flying them back to their city for a time off. This is common because most Australian cities are located in much more pleasant climates and surroundings than the locations of the mines.

These issues can be circumvented if the mine workers could operate the machinery directly from a control centre in the city in which they live, avoiding the need to travel to and from the mine. As well as having a more pleasant work environment, they could quite likely live a more conventional lifestyle, along with perhaps a more amenable personal and family life.

Another motivation to provide a tele-operation system for heavy machinery, is that of safety. With workers on site, controlling large machines in a harsh environment, there is a risk that accidents will occur, and because of the size and power of the machinery, these can be very serious or fatal. If humans are removed from the work site, the risk of injury and death is also removed. However, sufficient safety mechanisms need to be put into place to ensure that the system does not introduce other risks to the mining machinery itself. Damage to specialised mining machinery can

be very expensive to repair, but often the cost of the disruption to the output of the mine caused by an accident can be much more significant.

A safe and efficient tele-operation interface requires good vision of the tools and materials being worked on. It is generally accepted (Kitchin, 1997) that vision is the primary sense in understanding the spatial arrangement, size, shape and surroundings of an object. Typically vision is provided to a tele-operator via one or more video feeds from cameras at the work site. Standard two dimensional video can lack depth cues however, but this can be compensated for by providing two or more camera views at an angle to each other. An alternative to this is to provide a three dimensional video view, and allow stereopsis within the human visual system to create a mental 3D model of the scene. Raw video streams typically require a large amount of data for each frame of video. Video frames need to be displayed at approximately 20 frames a second or above to give the impression of smooth motion. Slower frame rates may be acceptable but actions and motion appear jerky and not continuous. To achieve the required frame-rates over a typical network bandwidth, the video data often needs to be compressed, or encoded, before transmission and decoded at the receiving end. The transmission time on the network adds to this to produce the overall latency.

An alternate method of providing a visual representation of the actions happening at a mine is to create a 3D graphical model of the mining components from the machinery design drawings and to update the positions of the moving parts in real time as interaction occurs. With mining equipment this can be achieved by installing sensors on the moving components and transmitting values representing the angles and positions of those components. The values are received at the control centre and are used to redraw the graphic representation of the machinery on a screen or other display in front of the operator. This has two benefits. Firstly, the data that needs to be transmitted is very small. The angle of a hydraulic arm can be represented within 32 bytes of data, whilst it would take many times more than that to transmit the pixels of the video frame showing that hydraulic arm. The second advantage is that a computerized 3D model can be rotated by the operator, allowing a viewpoint to be selected from any angle. This cannot be done with a simple video stream coming from a fixed camera.

However, there is a significant disadvantage in using a graphical model for tele-operation in a mining scenario. The material being worked on is typically rock or earth, the shape of which is not known before-hand. It also changes shape (crushing, breaking) during the mining process. It is unlikely that sensors could be successfully attached to the rock to detect its motion, so that any sensing of its size, shape and position needs to be gained visually. Three dimensional models of abstract shapes, such as rock, can be acquired through image analysis methods (Haneberg, 2006), but currently this does not happen at refresh rates fast enough for the application. In our work, we used a combination of a video view along with a moving 3D model to allow (close to) real time observation along with the ability to rotate the viewpoint onto the model. A dual joystick control was used to drive the actual components of the machinery and this arrangement was able to control a large piece of machinery at a mine site from control centre 1300 kilometres away (Duff, 2009).

A Haptic mining interface

A haptic interface can provide an operator with a sense of touching a 'virtual' (or computer generated) object. The device used in this experiment was a Sensable Omni (Sensable, 2010) (figure 2), which provides the sense of touch by pushing or pulling on the user's hand as they hold the end effector. In a tele-operation scenario, the haptic interface can be programmed to reproduce some component of a real interaction occurring at some other place. The forces that the mining tool is experiencing can be detected, and these numerical values can be transmitted to the control centre. The user interface then can represent these forces via the haptic device, to the operator's hand. Because the transmission is computer mediated, those forces can be scaled or altered in any way, to suit the interface.

With tele-operated mining equipment, the operator often needs to position large, multi-linked hydraulic arms. This positioning is a learned skill and it is often beneficial to locate it accurately on a particular part of a rock – be it a flat segment, a groove, crack or crevice. These features can be visible to a greater or lesser degree in a tele-operator's view, depending on the quality of the viewing system installed, but any ability to feel the texture and irregularities of the rock can enhance the operator's ability to efficiently work the rock. We implemented the haptic interface to allow the operator to feel the rock that they are about to work on. The tip can be scratched along the rock's surface, surface features can thus be felt as well as perhaps seen, with the aim of enabling a more accurate positioning of the tip.



Figure 2. Sensable Phantom Omni haptic device

A weak point in the design described, is the need for some sort of force sensor on the machine itself. There are several possibilities. The first is to mount a force sensor on the device close to the working tip. This would be prone to damage through vibration or contact with the rocks and is impractical. Another alternative would be to detect the oil pressure being applied to the hydraulic rams driving the machinery. This may be feasible but it could be complicated converting the various hydraulic pressures of the machine's segments into a total pressure being applied by the tip, as there are numerous combination segment angles and a proportion of the hydraulic pressure in any one configuration is applied to counter-act the force of gravity. Also there would be no differentiation between a force against a rock and a force of acceleration due to the device starting to move.

Fortunately there is a third, quite simple option, which takes advantage of the fact that sensor on the segments of the machine can report back the angles of those segments and therefore the position of the tip, in real time. The haptic user interface provides the operator with the ability to set a target position in 3D space, which is where they want the machine's tip to go. They do this by simply moving their hand held haptic device while looking at a representation of the scene. Within the scene they can see the target move as their hand does, as it is graphically overlaid onto the view. As the haptic device moves, the target destination is continually streamed to the machine's control system at the mine site and the machine starts to move autonomously towards the target. This movement continues until the real machine tip coincides with the target position. Such a system could be extended to incorporate path planning as described by Verner (Verner 2001).

However, if the target is placed by the operator to be just below the surface of a rock, the tip collides with the rock before it reaches the target and, in fact, may never reach the target. In this case it would come into equilibrium pressing against the surface of the rock, leaving a discrepancy between its actual position and the target. If this discrepancy is continually streamed back to the operator, the discrepancy can be portrayed as a force on the haptic device. The further the target is 'embedded' into the rock (i.e. the harder the operator is pressing against the haptic force) – the stronger is that reactionary force.

Also, as the operator drags the target towards a rock, there is likely to also be a temporary discrepancy in the position of the target and machine tip, which will also create a force via the haptic device. This force provides a sensation of pulling the tip of the machine on a spring, and when the tip connects a solid object, like a rock, that spring stretches. As the machine tip is dragged along the surface of a rock, this virtual spring can fluctuate in length and the variations are felt remotely by the operator. The result is a sensation of feeling the cracks, bumps and grooves of the rock.

The strength of the haptic force, as well as whether to apply it at all, is completely under the control of the software driving the system. In certain circumstances it may be appropriate to turn

off the haptic feedback altogether. One configuration in the system allows the operator to place a target with no forces at all and then only turn the forces on once the machine tip is within a certain distance of that target. The operator then has the option of quickly setting a target, attending to some other task while the tip travels there, then taking back direct control with force feedback for the rock breaking or drilling operation.

As mentioned earlier, image analysis allows the creation a 3D model of the rock. It is quite straight-forward to also represent this model haptically, so that the operator can now touch the virtual representation of the rock, independently of any forces coming back from the machine itself. We provided a mode which retrieves a 3D model of the rocks at the operator's request, and then allows the operator to explore the surface haptically, searching for a good location to attack it. Once at the target, the virtual haptic feedback is turned off and the real haptic feedback is turned on, so that the operator can then feel directly the interaction of the machine on the rock.

Implementation

The control mechanism installed at the mine-site is described in (Duff, 2009). The operator's end was developed in C++ and runs within Windows XP on a high end dual processor PC. It has both joystick and haptic interfaces. There are three video streams coming from three separate cameras at the machine site as well as a 3D model of the machinery that is regularly updated from tilt sensors mounted on the mining machine arm segments. The 3D model of the rocks is generated using CSIRO's photogrammetry software (CSIRO, 2007). The haptic interface made use of the H3D open source haptic development environment (H3D, 2010).

Testing

One problem encountered when testing mining applications is that the mining equipment is often very large and expensive and is likely to be fully in use for production. It can therefore be hard to get access to it for testing of new interfaces. We overcame this by first developing a software simulator for a particular mining machine: a rock-breaker. This should behave in the same way as the real thing when receiving control signals. We then developed and tested the operator's joystick interface against this simulator. We were then able to test the joystick interface for a short period to the actual machine on a mine during a pause in operations. Unfortunately real mine-site testing was not possible at a later date when the haptic interface was available. However, since the joystick interface was developed with repeated testing against the simulator, and this interface then worked as expected with the real machine, we were fairly confident that the simulator accurately mimicked the real thing. The haptic interface was then able to be tested against the simulator and displayed the ability to feel objects, and control the machine as described above.



Figure 3. Miniature robotic arm. Also showing joystick interface.

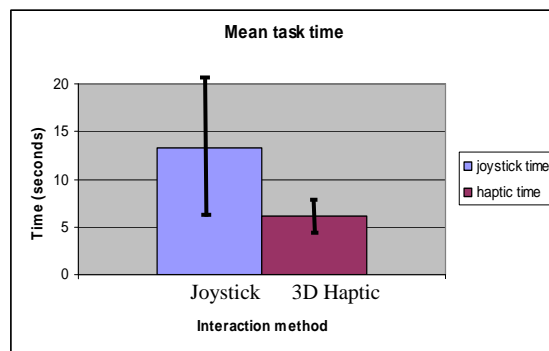


Figure 4. Task times for joystick and haptic control

To verify the use of this interface, we connected the software to a miniature robotic arm (figure 4) in the laboratory to run a user study on a subset of the system described above: comparing simple manipulation using the Phantom Omni against using a joystick. It should be noted that the configuration in the two cases differs in more ways than the simple presence or absence of haptics. With the phantom Omni the user moves their hand in three dimensions and the robotic arm copies the movement. With the joystick control the user moves one joystick left-right and forward-back for similar arm control, and moves the other joystick arm for up-down control. Our aim was to test speed and ease of use in the task of pushing a rock into a hole.

Ten subjects, selected from computer studies students and workers, were run through the task, half using the joystick first and half using the Omni first. The average time for the joystick interaction was 13.33 seconds with a standard deviation of 9.5 and for the 3D (Omni) interaction with haptics was 6.1 seconds with a standard deviation of 2.89 (figure 3). The 3D, haptic interface was also considered easier to use by 70% of the subjects. This agrees with (Aliakseyeu, 2002) in that ideally, interactions with data should look and feel as if directly manipulating the data itself, without the need to pay much attention to the interaction itself. In a tele-operation context the 'data' is the position of the tip of the robotic machine.

Discussion

A 3D haptic interface to tele-operate large machinery was found to be feasible and may have advantages over a joystick-style solution, in that it can provide an operator with a more natural user interface as well as the ability to explore the structure of surfaces by touch and to feel the response and pressure of the machine on a surface. A laboratory test showed that such a three dimensional user interface with haptic ability was superior to a more traditional joystick interface for speed and ease of use in a tele-operation task. It should be noted that this laboratory test did not involve mine or machinery operators and any new user interface would need to be extensively trialled with real equipment and the actual operators involved in the process before definitive conclusions could be drawn.

Acknowledgements

The authors would like to acknowledge the funding of the background work behind this project by Rio Tinto Plc. This work built on work by Elliot Duff in the CSIRO ICT Centre and Con Caris in CSIRO Earth Science and Resource Engineering.

References

- Aliakseyeu, D. et al.(2002). Interaction Techniques for Navigation through and Manipulation of 2D and 3D Data Eighth Eurographics Workshop on Virtual Environments, pp 179-188, 2002
- CSIRO (2010). <http://www.csiro.au/resources/earthmatters19.html> (accessed June 2010).
- Duff E., Taylor, T., Caris, C., Bonchis, A., Gunn, C., Adcock, A (2009). The Development of a Tele-Robotic Rock Breaker. 7th International Conference on Field and Service Robots, Cambridge, Ma., USA, July, 2009.
- H3D API. (2010), www.H3D.org (accessed June, 2010)
- Haneberg, William C. (2006), Sirovision. Environmental & Engineering Geoscience. Vol. 12, no. 3, pp. 283-291. Aug. 2006
- Kitchin, R. Blades. M. and Golledge, R. (1997). Understanding Spatial Concepts at the Geographic Scale Without the Use of Vision. Human Geography, 21(2). 1997.
- SensAble Technologies (2010). Phantom Haptic Device. URL: www.sensable.com, accessed 30 February 2010.
- Verna, D. (2001). Virtual Reality and Tele-Operation: a Common Framework. In Nagib Callaos, Susana Esquivel et Jamika Burge editors, SCI'2001, Fifth World Multi-Conference on Systemics, Cybernetics and Informatics , pages 499-504, volume 3, Sheraton World Resort, Orlando, Florida, USA. International Institute of Informatics and Systemics.