A First Generation Mutually-Immersive Mobile Telepresence Surrogate with Automatic Backtracking

Norman P. Jouppi, Subu Iyer, Wayne Mack, April Slayden, and Stan Thomas
Hewlett-Packard
1501 Page Mill Rd.
Palo Alto, CA 94304 USA
norm.jouppi@hp.com

***Abstract*—Mutually-Immersive Mobile Telepresence uses a teleoperated robotic surrogate to visit remote locations as a substitute for physical travel. Our goal is to recreate to the greatest extent practical, both for the user and the people at the remote location, the sensory experience relevant for business interactions of the user actually being in the remote location. The system includes multi-channel bidirectional video and audio on a mobile platform as well as haptic feedback. The user-controlled surrogate can autonomously backtrack along the path it has previously been driven if it loses wireless connectivity.**

Index Terms*—*Telerobotics, Odometry, Haptics, Video Conferencing, Multi-Channel Audio.

I. INTRODUCTION

We are developing Mutually-Immersive Mobile Telepresence[10] as an alternative to short-duration business travel. The business travel we would like to replace ranges from commuting through crowded urban areas (e.g., Los Angeles freeways) to airline travel across the continent or around the world. Business travel can have several unintended consequences, such as environmental degradation. For example, one person traveling roundtrip coast-to-coast in the United States on a modern jet airplane contributes over half a ton of CO_2 to the atmosphere, and CO_2 is responsible for about 81% of global warming due to greenhouse gas emissions[3].

Several technologies have been proposed as alternatives to business travel. These include technologies such as audio conference calls and video conferencing. Why have alternatives such as current video conferencing technologies not replaced more business travel? One hypothesis is that it is because such technologies are *not immersive*. What are some of the aspects of physical business travel that make experiencing remote locations immersive?

- Wide visual field
- High resolution visual field
- Gaze is preserved (people can make eye contact)
- Both remote participants and users appear life size
- Remote colors are perceived accurately
- High-dynamic range audio
- Directional sound field
- Mobility at the remote location

In Mutually-Immersive Mobile Telepresence our goal is to provide all of the above benefits of physical travel in an immersive way without physical travel and at lower cost.

In contrast, users of traditional video conferencing are limited to a single camera position. Typically a single video stream is provided, and this stream is incapable of providing both the full field of view of normal human vision and the resolution of the human visual fovea at the same time. Only one audio channel is provided, with only a limited dynamic range and no directional information. Presentation of people at life size is important for recognizing facial expressions and gaze, but current commercial video conferencing does not preserve gaze or present user's faces at life size.

Mobility at the remote location is a key enabler of casual meetings. Casual meetings have been documented as important for effective collaboration, communication, and innovation[13][18]. Office floor plans of research labs are often designed with this in mind, providing casual open seating areas and informal meeting areas such as kitchenettes. Discussions in hallways and during outings are often a key factor in the success of conferences and offsites. None of these are supported by commercial video conferencing.

There are a number of previous and ongoing projects that are related to our work. Buxton[2] describes a number of early telepresence efforts. Paulos and Canny's terrestrial Personal Roving Presence (PRoP)[15][14][6] allows the user to maneuver through a remote environment using a robotic platform. More recently, InTouch Health[9] has introduced a product with videoconferencing on a robotic platform. PRoP and the InTouch Health systems both carry video monitors for the display of the user's face and a single video camera for acquiring images to be sent to the user. This enhances traditional videoconferencing with mobility.

Our work builds upon the mobility enhancements of this previous work while addressing auditory and visual shortcomings of traditional video conferencing. This provides a more immersive experience for both the user and the people they are visiting, creating a much stronger sense of presence for both the user and remote participants.

II. FIRST GENERATION IMPLEMENTATION

The system consists of a teleoperated robotic surrogate at a remote location and the user at an immersion room. To complete construction of a first-generation prototype system quickly with a small number of people, we have used off-the-shelf components wherever possible. The following sections describe our first-generation prototype.

A. Model 1 Surrogate

Figure 1 shows the model 1 surrogate in use. We have designed the surrogate to be approximately the form factor of a sitting person. This overall size is dictated by the size of the four PC systems it contains, each one with five or more PCI/AGP/ISA cards. The surrogate consists of a base, mezzanine level, and a head. The head height shown in the figure works well in meetings (higher head heights cause too much of an obstruction for other remote attendees). It also works well when moving, since the higher the head the more unstable the platform becomes.

Fig. 1. The model 1 surrogate in use[11].

Early in the project we experimented with remote manipulators on the surrogate. Our goal was to be able to press elevator buttons and eventually to be able to open some doors. We controlled the manipulators with a Sensable Technologies Phantom Desktop. However we found controlling the arms to be difficult and unimmersive, especially with non-binocular vision. Moreover, some remote participants were apprehensive of a surrogate with manipulators, since either they had seen too many movies with bad robots or knew industrial-strength robot arms could be dangerous. Perhaps other cultures which view robots more favorably may be more accepting of surrogates with manipulators.

However, in our target application of business interactions between people, other than the shaking of hands little manipulation is actually required. In many cases electronic equivalents for manipulation exist. For example, handouts can be transmitted electronically and shared whiteboard technologies can also be used instead of actual manipulation. Instead of using teleoperated motion for gestures, we believe gestures such as pointing or waving can be better effected by video transmission of the user's own gestures at their location. Thus for now we have removed manipulators from the design of the surrogate.

1) Base: The base of the unit contains two 12V 100AH sealed lead-acid batteries totaling 2.4 Kilo Watt-Hours. This battery capacity is currently enough to power the surrogate for three hours. Charging of the batteries is accomplished by connecting the surrogate to an external battery charging station. More than ninety percent of the power is consumed by the electronics, so as technology scales with Moore's Law the battery life will increase.

The base also contains electric wheelchair motors, motor controllers, actuator controllers, and wheels. We originally interfaced through a serial bus to control a wheelchair motor controller, but this proved to be unreliable. We have recently switched to a RoboteQ AX2500 dual-channel motor controller[16] operated in closed loop speed mode.

We have chosen a configuration with symmetric drive wheels in the center and caster wheels symmetrically placed in the center of the front and back. This configuration allows the surrogate to turn in place, something that is natural for people to do. In contrast other wheel configurations, such as those standard in an automobile, require either a large turning radius or the execution of multipoint turns. These types of turns require conscious thought and are not consistent with our goal of providing immersion. The wheelchair motors are connected to the drive wheels by clutches, so that surrogates can be manually pushed around during service or deployment.

The surrogate suspension system is designed to be able to traverse terrain compliant with the Americans with Disabilities Act (ADA) [4]. This act dramatically simplifies and constrains the environment that the surrogate must operate within. Because of this, the surrogate does not have to be able to climb stairs. Some other examples of helpful constraints from the ADA specification include that curbs must have ramps at crosswalks and hallways must be at least 3 feet wide.

2) Mezzanine: The mezzanine level of the surrogate contains four ATX standard PC systems. Each system drives the LCD panel on one side of the head. Three systems also have multiple Viewcast Osprey-1000 H.261 video capture and compression cards running in CIF resolution (352 by 240) as well as a Viewcast Osprey-100 video capture card for remote backdrops (discussed later). Each PC uses a graphics accelerator based on a NVidia Riva TNT-2. This accelerator is capable of texture mapping multiple video streams over an entire screen with alpha blending at video frame rates, including downloading new textures for each frame. The fourth PC system acts as a router, connects to a wireless local-area network (WLAN), and controls the motion control electronics and sonar. All PC systems in the surrogate can be accessed by pointing an infrared wireless keyboard at their corresponding display. The windows in the lower center of each body panel are for infrared keyboard sensors.

Forty Polaroid ultrasonic sensors ring the surrogate. The sensors are used as an ultrasonic sonar ranging system as part of the surrogate's collision avoidance and navigation assistance system. They are contained in square projections on the surrogate skins (as shown in Figure 1). In an office environment, desks and tables are common but they have only a small profile area, reducing potential direct ultrasound reflections. By placing an upper ring of ten sensors close to the average table height, we maximize the direct reflection from the edge of the tables. This is especially important when the surrogate is close to a table, since the sensors have a relatively narrow field of view (e.g., 25 degrees). A lower ring of ten sensors is meant to detect low obstacles, such as trash cans.

Although the main part of the ultrasonic sensor's emission is above the maximum frequency of human hearing, when the sensors fire it is possible to hear a distinct click from the overall waveform envelope. Since the sensors are fired many times per second in sequence, the clicking can be quite distracting in quiet environments such as closed offices or conference rooms. Thus we only enable firing of the ultrasonic sensors when the surrogate is actually in motion.

The front panel of the surrogate includes four LED panel meters. Two of the panel meters monitor the voltage of each of the main batteries, while the other two monitor the current used by the motors and electronics, respectively. This information is also read out from the meters and relayed to the user's surrogate control display. By monitoring the voltage across each battery individually, the health of each battery can be ascertained.

3) Head: The head of the surrogate has three levels: LCD panels, video capture, and audio capture. Linear actuators are used to control the tilt and pan of the surrogate's head. The cameras and microphones are fixed to the head, so that when the user pans or tilts their surrogate head their visual and auditory views move as well. This produces the same effect for both the user and the remote people they are visiting as if the user was physically present and moving their head.

The head of the surrogate is built with four LCD panels at right angles to each other. Panels on the front and both sides of the head display live video of the user's head. The LCD panels have special wide view coatings which allow them to be viewed over a range of almost 180 degrees both horizontally and vertically. This is in contrast to laptop LCD panels, which can only be viewed over small angles without losing contrast or appearing to have inverse video artifacts. The surrogate's head gives the appearance of an orthographic projection of the user's head. In the long run, as flexible display technologies become available, it would be better to use a display with a shape closer to that of the user's head. The surrogate has two small speakers mounted under the LCD panels of the head, and a subwoofer inside its body.

Eight compact CCD board cameras are mounted directly on top of the displays. To reduce gaze errors, we would like the cameras to be behind the display of the user's eyes on the LCD panel, but this is impractical since the LCD panels are lit. The next best place to mount the cameras is directly above the LCD panels, since this introduces the least error between the surrogate and what would be seen if the user is actually present at the remote location. Five of the board cameras are used to acquire video to send back to the user. The three remaining cameras acquire remote backdrops (described later). The cameras are concealed behind a ring of neutral density plastic to avoid distracting remote participants.

Above the camera level the head contains four short shotgun microphones at right angles to each other. The shotgun microphones are directional, and capture sounds coming from the left, front, right, and back of the surrogate. The top level of the head also contains a WLAN antenna.

B. Immersion Room

The Immersion Room (see Figure 2) uses two BarcoGraphics 6300 projectors with the widest possible lenses in a rear projection configuration. The projectors are housed in custom-designed “hush boxes” to reduce their already low noise levels. Five videos from cameras in the surrogate's head are composited together for projection onto the screen. The screen is bent at an angle of 127 degrees at the center where images from the two projectors meet, and has an aspect ratio of 16:6. Moreover, it fills close to 120 degrees of the user's field of view. In order to capture an eye-level view of the user to preserve their gaze, we make a small slit in the projection screen and mount a camera with a pinhole lens on an extension tube behind the screen. Because the projectors have wide angle lenses and the screen is bent, there is plenty of room behind the bend in the screen to mount a camera. Since the pinhole lens is mounted on a narrow tapered tube, only a small dark spot is present on the screen. This camera allows people to make eye contact with the people they are visiting remotely.

Fig. 2. The immersion room.

We use two light boxes and three reflectors to pleasingly light the user's head while preserving display screen contrast.

On the desk in front of the display is an infrared wireless keyboard, joystick for audio control, Immersion Corporation force-feedback joystick, and a LCD panel. The LCD panel displays the status of the system and surrogate. The walls of the immersion room are painted chromakey blue.

1) Audio Telepresence: In the audio telepresence component[12], we attempt to recreate as accurately as possible both the remote sound field for the user, as well as the user's contributions to the remote sound field as if they were actually present. We transmit four channels of near CD-quality audio from the remote location to the user. Each audio channel is currently sampled at 40KHz with 16-bit resolution. A single channel of 40KHz 16-bit audio is transmitted from the user's lapel microphone to the surrogate. The high dynamic range of the system allows users to whisper back and forth with remote participants. This is a key enabler for private conversations in public spaces, overcoming a limitation in existing technologies[13].

We ring the user with four Bose Acoustimass speakers and a subwoofer for accurately recreating the remote sound field. The user can control the relative gain of the four speakers by steering with a joystick. This is useful for amplifying the voice of a single person speaking in a noisy room and enhancing the "cocktail party" effect available with multichannel audio.

We have developed our own software to acquire, compress, decompress, and output the audio data. This software also reduces the gain of the microphones when the speakers are active in order to prevent echoes in the audio system. True echo cancellation would be difficult given the changing acoustics experienced as the surrogate moves about, combined with variable network delays through the internet.

2) Haptic Collision Avoidance and Navigation Assistance: A research-quality force-feedback joystick[8] is used for driving the surrogate. The basic controls are similar to that in a wheelchair joystick. Forward and backward joystick motions denote forward and reverse movement, while motions to the right denote clockwise rotation of the surrogate and motions to the left denote counterclockwise rotation, with combined motions resulting in combined movement.

A surrogate status and control window (shown in Figure 3) is displayed on a LCD panel mounted on the desk of the user. In normal operation the user should not need to watch the panel; the combination of video from the remote location and haptic feedback through the joystick should suffice for driving the surrogate. The large white circle with red diamonds shows the distance to the nearest object at 45 degree intervals as returned from the ultrasonic distance sensors, similar to a radarscope. The white circle on the right gives an indication of both the actual forward and angular velocity of the surrogate based on data from the surrogate wheel encoders.

We use a variant of artificial force reflection control[7] to aid the user in controlling the surrogate. If the speed requested by the user pushing the joystick would cause the surrogate to impact an obstacle within 4 seconds, a "wall" sensation is presented haptically to the user in the direction of the obstacle and the surrogate slows down so that any impact will be 4

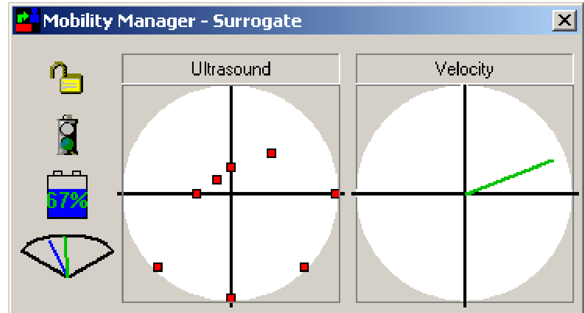

Fig. 3. The user's surrogate status and control window.

seconds away. If the user continues pushing on the joystick, as the surrogate gets closer to the obstacle it will continue to slow down so that the time to impact remains 4 seconds. If this continues further, the surrogate will eventually stop just in front of the obstacle when the time to impact in 4 seconds results in a speed below the minimum speed of the surrogate.

We originally tried presenting a force that was proportional to the distance to the object in the direction of motion. However, in order to prevent the surrogate from "running away" when the user lets go of the joystick, a significant spring-type force must be output at all times to center the joystick. This force uses up a significant amount of the force range. Furthermore, if additional small forces are presented to the user representing objects at moderate distances, these forces will be small compared to the spring force when the user is pushing the joystick far from center. This makes all but the most extreme cases hard to perceive by the user. Thus, we now use an interface that only outputs a large additional "wall" force when an impact is imminent within 4 seconds.

Because our first generation surrogate is longer than it is wide, rotations of the surrogate can also result in collisions at the remote location. In order to model these potential collisions haptically, the horizontal motion of the front and back of the surrogate must be computed as a function of requested rotation rate, and if this results in an impact within 4 seconds the rotation rate is reduced to increase the time to impact to 4 seconds and a haptic "wall" to the right or left is presented to the user through the joystick.

We have found the artificial force reflection to be a useful aid in navigating through tight corridors or around obstacles. However, a force-feedback device with greater dynamic range (or less need of centering spring force) would be a useful enhancement and allow gradations in distance to be communicated.

3) Head Tracking and Remote Backdrops: We have developed head-tracking techniques based on chromakeying. In order to present the user's face life-sized at the remote location, it needs to fill the LCD display. But users naturally shift around in their chairs, so to prevent the user's head from falling off the surrogate's displays we capture a wider field of view at the user station and then translate and zoom with texture mapping to fill the surrogate's display with the user's head. Given the size of the head's LCD panels, this presents the user's head at roughly life size. The immersion room walls behind and to the sides of the user are painted chromakey blue. The use of the blue screen allows us to identify the bounding box of the user's head at video speeds with low overhead on a PC.

Instead of leaving the blue-screen behind the user on the displays of the surrogate, we capture a wide field view out the opposite side of the surrogate's head. We first texture map this onto the whole screen, then alpha blend the user's face over it with the alpha being set to a function of the blueness of the pixel. This generates a final image of the user's head in front of scenery from the remote location. This remote backdrop adds to the level of immersion experienced by people at the remote location.

4) Network Requirements: The surrogate currently requires about 2.7Mb/sec of wireless bandwidth. All video streams are acquired at H.261 CIF resolution (352 by 240). The views going from the surrogate to the user are compressed at 15 frames per second and at a rate of 320 Kbits/sec. The video streams from the user to the surrogate are compressed at 10 frames per second and a rate of 200 Kbits/sec. We find less bandwidth is required for the user's face than for the view presented to the user. The video bandwidths total 1.6 Mbits/sec from the surrogate to the user and 0.6 Mbits/sec from the user to the surrogate. Audio bandwidths total 0.4Mbit/sec from the surrogate to the user and 0.1Mbit/sec from the user to the surrogate. The bandwidth for surrogate motion control is negligible. We are currently connecting to the internet through an IEEE 802.11a network that supports peak data rates of between 6 and 54 Mbits/sec.

III. AUTOMATIC BACKTRACKING

Some parts of the remote location may not be covered by wireless services or may have dead zones. If the user of the surrogate drives into an area without wireless coverage, the surrogate can effectively become disconnected from the user. Since the user is not at the remote location with the surrogate, they cannot push it into an area with adequate wireless coverage. Moreover, since the connection to the surrogate from the user has been lost, they cannot use the surrogate to ask people at the remote location to push the surrogate back into an area with adequate wireless coverage. Thus it can be quite time consuming and inconvenient to regain connection to a surrogate that is driven outside areas with adequate wireless coverage.

Inadvertent travel outside wireless coverage areas is a problem that needs to be addressed before widespread deployment of surrogates is possible. This problem is more acute for surrogate deployment than in other teleoperated systems. This is because the range of WLAN access points are quite limited, and large metal objects (such as HVAC ducts or elevators) can cause significant radio shadows invisible to the user. To date, WLANs have been deployed in unlicensed radio bands where power levels and signal strengths are limited. Furthermore, as higher and higher frequencies are utilized in order to achieve higher bandwidths, radio signals behave more like light rays and shadowing by metal objects becomes more pronounced. In contrast, most prior teleoperation systems often used lower frequencies that were less affected by shadowing, used higher signal strengths, and/or traveled within restricted areas, so loss of coverage was less of a factor.

A. Autonomous backtracking mode

To solve this problem, distance-measuring and collision-avoidance hardware that is used to assist the user in driving the surrogate during normal operation can be used to help operate the surrogate autonomously. In autonomous backtracking mode, the surrogate “follows its footsteps” back into an area with adequate wireless coverage, where the user can regain connection to it and control it again in a teleoperated fashion. The surrogate only needs to back up by a limited amount, since the communication loss can be detected within several seconds (at most) of its occurrence. During normal operation the surrogate keeps track of its relative location as the user is driving it, through monitoring position encoders on its wheels (i.e. odometry[5][17][1]). Inaccuracies in direction or distance measurements can accumulate over time, so dead reckoning is not effective for navigating over large distances or over long periods of time. However in our application we only need to reverse up to ten seconds of surrogate motion to get back to a place where an adequate wireless connection previously existed.

If the surrogate cannot reacquire communication with the user after reversing a specified distance it should halt and wait. This is because if it has returned to an area where it previously had good communication with the user without being able to restore communication, the communication loss may be due to a failure at some other point in the network. In rare cases the surrogate may not be able to return along the path previously taken due to the appearance of obstacles in the path since it was there last (e.g., a door has just been closed behind it). In this case the surrogate should temporarily suspend backtracking, and try to keep a specified distance between it and surrounding obstacles while remaining on the path.

Once the surrogate has entered backtracking mode, it must navigate along its recent path in the reverse direction. During backtracking, it is best if the surrogate does not travel at high speed. Therefore backtracking may take longer than the original motion. However, backtracking should not take too long because any undue delay increases the amount of time that the user is left incommunicado.

B. Loss and reacquisition of adequate wireless connectivity

One feature common to many WLANs is that the supported bandwidth varies with distance and signal path quality to the access point. At significant distances from the access point or if significant obstructions occur in the path to the access point the communication can downshift to a much lower bandwidth. If the surrogate requires a data rate greater than the minimum data rate for normal operation, networking failures will occur when the surrogate moves into an area supporting only a lower data rate. Such failures can consist of dropped packets or rapidly increasing latency, depending on the protocols and hardware being used. When the percentage of packets lost exceeds a threshold percentage or packet transmission delays exceed a time threshold, the surrogate can treat the connection as effectively being lost and initiate backing up.

Traffic from other wireless clients can also overload an access point being used by a surrogate. This can cause loss of packets or excessive packet delays. If these problems are transient (i.e., lasting less than several seconds), the surrogate will not need to initiate backtracking. The automatic backtracking only needs to begin when a networking problem persists for several seconds or more.

C. Ancillary features

Since the surrogate is more prone to problems when it is traveling autonomously, during the time that the surrogate is operating under autonomous control, this may be optionally indicated to people at the remote location by blinking a small yellow light and/or producing a quiet "backup beeper" type of audible warning (e.g., 1kHz sine wave, 50% duty cycle per second).

Finally, note that the addition of a default backtracking mode for surrogate navigation does not preclude the use of more standard robotic techniques such as autonomous navigation in situations where the user cannot find an alternative path with better wireless coverage or the user desires to travel to a specific destination at the remote location.

IV. FUTURE DIRECTIONS

Based on our experience using our first-generation system, we are implementing several improvements in our second-generation system. First, the oblong shape of the model 1 surrogate means that it cannot be turned in place in narrow places such as doorways. To correct this problem, we plan to make our future surrogates roughly circular in shape. Second, having the surrogate head at a fixed sitting height was awkward when the user was interacting with standing remote participants. So future surrogate designs should also have the capability of adjusting their head between a sitting and standing height. Third, advancements in electronics will allow future surrogates to be implemented using two PCs instead of four, enabling the surrogate to have a more anthropomorphic form factor.

V. CONCLUSIONS

Mutually-immersive mobile telepresence can bring immersive telepresence to ordinary public places. It leverages technologies such as computer graphics hardware, wireless networking, and the internet which are rapidly increasing in capability and decreasing in cost. Initial user feedback on our prototype system has been favorable.

As part of the project, we are investigating a number of different technologies. The audio telepresence component enables users to whisper back and forth with individual remote participants, auralize the location of remote sound sources, and utilize the cocktail party effect to improve the intelligibility of remote speakers. The 1-to-1 correspondence between participants and "people-sized spaces" provided by the surrogate helps make remote meeting attendees full-fledged participants. The mobility of the surrogate is an important enabler for casual meetings.

In order for surrogates to be widely deployed in office settings, the problem of limited wireless network connectivity needs to be addresses. We have designed the surrogate to autonomously backtrack along the path it has previously been driven by the user in the event it loses wireless coverage. This is a key enabler for the actual deployment of surrogates in real business applications.

ACKNOWLEDGEMENTS

The authors are grateful for those who helped with various aspects of the project. Dan Scales implemented the original motion control software. Summer interns Mark Newman, Keith Emnett, and Mike Pan were a great help in the initial system implementation.

REFERENCES

[1] J. Borenstein and L. Feng. Measurement and Correction of Systematic Odometry Errors in Mobile Robots. *IEEE Transactions on Robotics and Automation*, 12(6):869–880, December 1996.

[2] W. A. S. Buxton. Telepresence: Integrating Shared Task and Person Spaces. In *Graphic Interface '92*, pages 123–129, 1992.

[3] EPA Global Warming Web Site. http://www.epa.gov/globalwarming/, 2003.

[4] Evan Terry Associates. *Pocket Guide to the ADA: Americans with Disabilities Act Accessibility Guidelines for Buildings and Facilities*. John Wiley & Sons, revised edition, 1997.

[5] H. R. Everett. *Sensors for Mobile Robots: Theory and Application*. A. K. Peters Press, 1994.

[6] K. Goldberg and R. Siegwart, editors. *Beyond Webcams*. MIT Press, 2002. Chapter 9.

[7] S.-G. Hong, B. S. Kim, S. Kim, and J.-J. Lee. Artifical Force Reflection Control for Teleoperated Mobile Robots. *Mechatronics*, 8(6):707–717, September 1998.

[8] Immersion Corporation Impulse Engine 2000. http://www.immersion.com/industrial/products/impulse_engine2000.php, 2003.

[9] InTouch Health. http://www.intouchhealth.com/, 2003.

[10] N. P. Jouppi. First Steps Towards Mutually-Immersive Mobile Telepresence. In *Proceedings of ACM CSCW '02*, pages 354–363, 2002.

[11] N. P. Jouppi, S. Iyer, W. Mack, and A. Slayden. First-Generation Mutually-Immersive Mobile Telepresence. In *Video Proceedings of the IEEE International Conference on Robotics & Automation*, 2004.

[12] N. P. Jouppi and M. J. Pan. Mutually-Immersive Audio Telepresence. In *Proceedings of the 113th Audio Engineering Society Convention*, October 2002.

[13] R. Kraut, R. Fish, R. Root, and B. Chalfonte. Informal Communication in Organizations: Form, Function, and Technology. In R. Baecker, editor, *Groupware and Computer-Supported Cooperative Work*, pages 287–314, 1993.

[14] E. Paulos and J. Canny. Designing Personal Tele-embodiment. In *IEEE International Conference on Robotics & Automation*, pages 3173–3178, 1998.

[15] E. Paulos and J. Canny. PRoP: Personal Roving Presence. In *ACM CHI '98*, pages 296–303, 1998.

[16] RoboteQ Dual-Channel Motor Controller. http://www.roboteq.com/, 2003.

[17] K. Seng and L. Kleeman. Accurate Odometry and Error Modelling for a Mobile Robot. In *IEEE International Conference on Robotics & Automation*, pages 2783–2788, 1997.

[18] S. Whittaker, D. Frohlich, and O. Daly-Jones. Informal Workplace Communication: What Is It Like and How Might We Support It? In *ACM CHI '94*, pages 131–137, 1994.