Ubiquitous Telepresence

Benjamin G. Zorn

CU-CS-824-96

University of Colorado at Boulder
DEPARTMENT OF COMPUTER SCIENCE

Ubiquitous Telepresence

Benjamin G. Zorn

Department of Computer Science
Campus Box 430
University of Colorado
Boulder, CO 80309–0430 USA

CU-CS-824-96 December 1996

University of Colorado at Boulder

Technical Report CU-CS-824-96
Department of Computer Science
Campus Box 430
University of Colorado
Boulder, Colorado 80309

Ubiquitous Telepresence

Benjamin G. Zorn

Department of Computer Science
Campus Box 430
University of Colorado
Boulder, CO 80309–0430 USA

Telephone: (303) 492-4398, FAX: (303) 492-2844
E-mail: zorn@cs.colorado.edu
URL: http://www.cs.colorado.edu/homes/zorn/public_html/ut/Home.html

December 1996

Abstract

This paper describes a technology vision for the future that I call *ubiquitous telepresence*. The main idea of ubiquitous telepresence is to allow users to project their physical presence anywhere on the internet. Specifically, users are given the ability to explore and interact with many remote physical environments (e.g., all sites on the internet) through remote sensing and manipulation devices. My vision for implementing ubiquitous telepresence is to build robotic or remote-controlled units at such low cost that they become as ubiquitous as television sets.

Ubiquitous telepresence extends the concept of virtual reality by providing users with a physical embodiment of their presence in a physical location. Ubiquitous telepresence is related but separate from ubiquitous computing in the model of how telepresence is provided. In particular, ubiquitous telepresense focuses on providing telepresence by defining interfaces, software, and hardware for standard low-cost robotic devices (ubots) that are connected to the World Wide Web. Instead of assuming the ubiquity of many different computing devices permeating an environment, ubiquitous telepresence requires only the presence of a single inexpensive ubot to realize its power.

This vision statement describes the motivation of ubiquitous telepresence, an assessment of the technology required to realize it and the areas of research required, and its relation to existing, competing technologies. I also attempt to suggest future directions that this technology may take.

1 Introduction

Many recent developments, such as the growth of the internet, the extremely high public visibility of the World Wide Web, the rapid development and deployment of new programming language technology (e.g., Java), and the dizzying evolution of consumer electronics products suggest that technology is changing at an unprecedented pace. The presence of high-bandwidth access to enormous amounts of information on the internet has greatly facilitated everyone's ability to find significant information, learn about new ideas, and share ideas with a very large community.

This document describes a new technology concept that I call ubiquitous telepresence (UT). As the name implies, this technology enables users to maintain a physical presence remotely in any one of many places. For example, the technology could allow a user anywhere in the world who has access to the internet to tour any of the largest museums in the world at any time.

How could such a technology work? The central component of this technology is a very low-cost mobile robotic device that can be controlled via commands over the internet. I call this device a "ubot" to distinguish it from traditional mobile robots, which are commonly autonomous vehicles. I propose that such ubots can be manufactured even with existing technology at sufficiently low cost that they would provide many of the capabilities that this vision statement outlines.

The purpose of this vision statement is to motivate why ubiquitous telepresence is needed, to describe the basic capabilities of the technology, to identify existing research that is related, and to discuss the implementation of this technology and its current feasibility.

2 Motivation

The investigation of UT technology is motivated by the presence of the internet and the capabilities it provides. While the internet is currently being most used for storing, sharing, and transferring information, its physical substrate enables the proposed implementation of UT. The capability of projecting one's physical presence remotely provides tremendous advantages that are not realizable in the current context of either the internet or virtual reality systems. The field of telerobotics, which is closely related to the proposed vision, has been very successful in the limited domains in which has been applied. The most notable applications involve very expensive units that are specialized for particular jobs (e.g., spacecraft, submersibles, and nuclear waste cleanup robots)[1]. This proposal builds on the successes of these specialized applications, and investigates the possible uses of much lower-cost and more widely available devices. If UT is successful, then the related

[1]For an overview of such devices, see http://nctn.oact.hq.nasa.gov/STI/Innovation21/Telepresence.html

telerobotics technologies will certainly greatly benefit both from expanded research interests and the cost reductions enabled by mass production.

In the remainder of this section, I discuss the suggested capabilities of the proposed technology, and then describe a number of possible applications.

2.1 Proposed Capabilities

As mentioned the central hardware component of UT technology is the ubot, an inexpensive mobile robotic device. I also will refer to the ubot as an "agent" device because it represents a physical embodiment of software agents. For the device to attain the level of ubiquitousness that I envision, the consumer cost would have to be comparable to current home PC's (i.e., in the range of $1000–$3000). The ubot is not intended to be a completely autonomous device, as a human would be controlling it remotely. Providing the ubot with different levels of autonomy (e.g., autonomous movement) would reduce the level of interaction required for the user to operate one. The issue of how much autonomy the ubot requires needs to be investigated as part of this research. The essential capabilities of the ubot include:

1. Mobility—telepresence is highly limited if the "agent" device cannot move at all. Tethered mobility represents one possible solution, but the ideal ubot would be highly mobile and untethered. Future directions of this work will explore the issues of mobility in different environments (e.g., underwater, in the air, and in space).
2. Video output—a very simple and very important capability that the ubot can provide is remote viewing. Vision is essential if the ubot is being controlled entirely from a remote site (i.e., the ubot is completely non-autonomous).
3. Communication link to internet—For the ubot to be controlled remotely, it must have a two-way communication link to the internet (or a machine connected to the internet).
4. Two way audio—the ubot should have both sound sensors and speakers. Such technology is currently inexpensive, and will enable applications that require human/ubot interaction.
5. Additional sensor technology—technology that enables the ubot to measure different things is both inexpensive and useful. For example, infra-red autofocus devices currently use on cameras used to measure distances could be valuable. Such devices could be used either to provide the remote operator with additional information, or used by the ubot itself to give it greater autonomy.

With these five capabilities, the ubot is capable of performing a significant number of the proposed applications. If possible within the constraints of the budget, additional capabilities would also be valuable:

1. Manipulators—ideally, the ubot would be able to physically manipulate its environment. Such functionality would fully support the notion of complete telepresence, in which the remote agent is capable of doing everything the controller is capable of (and more, considering the device is a machine). The major limitation in this category is cost/benefit ratio of the different possible manipulators. In the scenarios below we outline the kinds of manipulation that would be useful.

With the capabilities outlined, the ubot can be controlled from a remote internet site, and the data from the ubot's sensors (video, audio, etc) can be fed back to the controller.

3 Applications

In this section I outline several possible applications of UT technology. At the conclusion of each example, I will discuss the some of the research issues that are required to realize the application described. The initial scenario illustrates many of the advantages of UT technology, contrasts available alternatives, and points out many of the technical challenges currently present.

3.1 Tour an Art Museum

Scenario An art museum has a number of ubots that are connected to its internet server. A user at an arbitrary site can connect to the museum server, "check out" a ubot, and start controlling it. The user can then move the ubot about the museum, examining the entire collection at whatever degree of resolution the video camera provides. The museum can charge for such uses, supporting the cost of purchasing the ubots.

Discussion This application strongly illustrates the differences between UT and competing technologies. For example, one alternative would be for a PC user to purchase a CD-ROM containing many pictures of the artwork in the museum. The UT solution has many advantages over the CD-ROM solution. First, the museum collections change, and the CD-ROM user would be unable to see new exhibits without buying a new CD-ROM. Second, the images on the CD-ROM are static, limited, and of finite resolution. If a person wanted to see a piece of art in more detail (or from a different angle), they would be unable to with the CD-ROM solution. Finally, since the CD-ROM is limited in capacity, it is likely that many parts of the collection would not be able to fit on a CD-ROM, thus they would be generally unavailable.

Another alternative is for the museum to place images from its collection on the WWW (as the Louvre has done, see http://www.paris.org/Musees/Louvre). This solution addresses the problem that the collections change, and if the curators are ambitious, they can also make a point to put an image of every work of art on-line. Realistically, taking photographs of every piece would require a large-scale effort that would be very expensive (especially for museums that are less well-funded than the Louvre). In addition, the problem of resolution and viewing angle remains. Thus, the UT solution is less expensive (e.g., just the cost of the ubots) and more flexible.

Another available solution would be to put the entire collection into a virtual reality space that would include very high-resolution images of the paintings and very detailed 3-dimensional models of all the sculptures. This VR solution supports highly flexible interaction with the collection, but at a potentially high cost. In particular, the storage space to maintain such a collection, and the network bandwidth to transfer pieces of the collection to a remote viewer may be prohibitively high. In addition, the cost of capturing the museum as high-resolution images and 3-D models is currently very expensive.

Finally, there is the ubiquitous computing approach. Using this technology, there might be a computing device (presumably with a camera attached) associated with every work of art. This would reduce the need for mobility because the computing devices would be ubiquitous. However, this approach has a significantly higher cost then UT, and could require the museum to invest in a large scale renovation of the facility. Using UT, the facility itself does not change, but instead a single new kind of device is placed within it.

Beyond any existing technology, UT ubots with manipulator arms would provide remote users with the capability to explore beyond what is available on display. Many museums have vast collections in back rooms that are not available to the general public, and are often greatly underutilized due to the inconvenience of traveling a long distance to view them.

Issues This application requires minimal technology beyond that already described (including no manipulators). Probably the most significant issues are those of mobility (e.g., how do the ubots move around in the museum) and of cost (how much does each device cost, and how to provide enough that all the potential users in the world can be satisfied).

There are a number of important research problems that arise for this particular application. In particular, it is likely that given resource limitations, users would probably be required to "share" ubots. After all, you couldn't have 1000 of these devices walking around the Louvre at one time, except perhaps after hours. The different issues that arise in the context of sharing (e.g., providing the illusion of not sharing, ceding control, etc) will require research in systems and collaborative computing.

Another interesting and important issue is one of safety and security. If ubots are moving around among human beings, how does one prevent them from hurting people, each other, themselves, or the art work? Some of these issues can be resolved by carefully designing the hardware, while others will require "software overrides" that prevent harmful situations.

3.2 Archeologists or Paleontologists Collaborate on a "Dig"

Scenario An archeological (or paleontological) site in a remote location has been discovered. A team of scientists, each of whom resides in a different place, wants to collaborate on the project. A number of ubots are brought to the site, and each scientist controls a ubot, directing the device toward a location they are interested in. The scientists are enabled to act independently in their investigation, but collaborate in a virtual space that brings them, their computer data, and the physical artifacts together so that they can share their ideas, organize the information and physical artifacts, and create new information. The ubots in this scenario could be augmented with manipulators consistent with the kind scientific investigation being performed. Scientists local to the site could also be integrated into the collaboration space.

Discussion This application illustrates the possible integration that can occur when UT technology and VR technology are both utilized. No existing technology can provide the capabilities suggested in this scenario. To effectively realize this vision a number of significant research issues must be considered. First, how is the virtual reality collaboration space structured? How are diverse sources of information that can be valuable to the scientists integrated into this space. How are the physical artifacts being investigated integrated into the virtual reality collaboration space?

How do the scientists interact with each other? Can the scientists collaborate to control the ubots? Just as significant advances in VR research have been enabled by low-cost VR I/O devices (e.g., head-mounted displays, data gloves, etc), possible scenarios such as the one above will be enabled by the existence of low-cost UT technology.

3.3 Other Applications

A large number of other applications are possible given the technology outlined. Here, we provide a brief description of a few such applications to provide a sense of the scope of the intellectual and commercial opportunities available.

Visit the Ocean Floor (or any Remote Environment) Submersible ubots could be tethered to a boat and controlled remotely via the WWW. Users could rent a ubot and spend a few hours exploring a remote ocean reef. In addition to this recreational use of ubots, they could also be very valuable for scientific applications. Biologists could use ubots to conduct population surveys in many different remote environments (ocean floor, arctic tundra, etc.) at relatively low cost.

Ubot Competitions Competitions between ubots could be held by bringing a number of ubots together at a site and allowing their controllers to operate them remotely.

Tour a Home Remotely If a person is relocating, they may want tour a home in a distant city. With UT technology, the real estate agent would simply have a ubot stationed in the home being viewed, and this would allow the interested buyer to explore the home at will, even flushing the toilets and testing the faucets if simple manipulator capabilities were available.

Unplug Your Iron Remotely Ubots could be used for inexpensive security devices. In particular, home owners could use them to inspect their homes while away on vacation. With the appropriate manipulators, ubots could also be used to unplug a iron or coffee pot accidentally left on while away, or even used to feed pets.

This section has outlined the capabilities and possible applications of UT technology. In the following section, I discuss how such technology might be implemented using existing components.

4 Related Work

Before discussing the issues involved in implementing UT technology, we mention some of the current related work in this area and describe how our vision relates.

As mentioned, a significant amount of research has related to the remote operation of robots, or telerobotics. Traditionally, these efforts have been targeted at specific applications, such as waste disposal [3], hazardous environments [2], and space exploration [8, 11]. [2]

Several research groups have investigated using the internet for telerobotics, and at least three groups have already created preliminary implementations of the ubots described in this proposal. In particular, a Kaplan, Keshav, Schryer, and Venutolo at Bell Laboratories have mounted a video camera with a wireless transmitter onto a radio controlled car and driven the car on the internet [6]. The video is transmitted over the Internet Multicast Backbone (MBONE) and the program nv was enhanced to provide local frame rates up to 15 frames per second. The paper describes some of the problems encountered in engineering the system, including issues related to the user interface and how bandwidth problems affected the usability of the device.

Vanu Bose of the Telemedia, Networks And Systems (TNS) Group at MIT constructed a similar vehicle called the TNS Video Rover [1]. This vehicle was connected to the internet and used the VuNet (a research gigabit ATM network) to transmit video locally. Problems encountered in this implementation include providing power to the ubot device (traditional R/C batteries typically last less than an hour) and the range of the wireless transmission (about 30 feet indoors).

Other work that is related to ours is that of Paulos and Canny at UC Berkeley. They have implemented two telerobotic devices, Mechanical Gaze [12] and the Web Blimp [13]. Mechanical Gaze allows remote users to manipulate a robot arm to which a digital camera is attached, and to investigate small museum pieces such as mineral specimens, live insects, live reptiles, etc. In the Web Blimp project they placed a radio controlled blimp with an attached wireless video camera on the Web.

There are direct parallels in the goals and approach taken with the Web Blimp project [13] and our ubiquitous telepresence project. On the home page for the project,[3] there is even a reference

[2]For a comprehensive overview of the work in this area, see the WWW page http://www.mcs.anl.gov/home/jebb/telerobot/telepaper.html assembled by Michael Jebb.

[3]See `http://vive.cs.berkeley.edu/blimp/`.

to "Ubiquitous Tele-embodiment" indicating that their vision and ours are very closely related (with matching acronyms). They use the term "mobot" in a way very similar to our term "ubot". They estimate the current cost of the hardware at $1000, with the envisioned cost dropping to approximately $300 when manufactured in quantity. While the idea of using a blimp is appealing, there are still important technical challenges that they face. In particular, the blimp is currently 6'x3'x3' and has a payload of about 1 pound. They have also not addressed the problem of the power requirements of the device, which currently requires frequent recharges.

We view these three research efforts as directly related to our work, and our current goal is to understand the engineering problems they encountered and attempt to move beyond the limitations of these existing systems. In particular, our current goal is to provide higher frame rates (i.e., via lower bandwidth requirements) and longer battery life for the ubot than has been provided in previous work. While we find the use of a blimp intriguing, we are currently investigating a ground-based alternative because it appears much easier to provide in a highly available manner.

Another mobile robot connected to the WWW is Xavier, created by the Learning Robots Lab (LRL) at CMU [9]. Xavier is a platform on which a large number of robotic experiments have been performed, including learning, planning, and position estimation. The focus of the LRL has been on autonomous mobile robots, and as a result, their results are only partially applicable to the UT project. Our focus will initially be on less autonomous remotely controlled devices, with greater emphasis on maintaining a low cost, and less on guaranteeing autonomy. However, if devices like Xavier can be constructed at low cost, they would be prime candidates for implementing UT technology.

There has also been a proliferation of mechanical devices attached to the WWW for some time (e.g., the Amazing Fish Cam, etc), and some of the most popular Web pages contain pointers to these devices [15, 14]. Currently, these devices include many fixed cameras, some mobile cameras (including a head-mounted camera at MIT [10]), coffee pots, pagers, games, soda machines, and a few robots with attached cameras.

Of the non-mobile projects that are most closely related to ours, Ken Goldberg has developed two tele-operated robots, the Mercury Project [4] and the WWW Tele-Garden [5], at least one of which has been available on the internet almost continuously since August 1994. The Mercury Project is a remotely operated robot that could direct blasts of compressed air on dry earth, allowing remote operators to view and "excavate" a location. The Mercury project was the first to permit tele-robotic manipulation of a remote environment. The Tele-Garden project extends this model further, allowing remote participants to plant, water, and view seedlings in a remote garden. The emphasis in these projects on the experience of the participants is significant, as is the goal of providing high reliability and availablility of the site. The major difference between these existing projects and our own is our emphasis on mobility and the associated problems and constraints.

`sniffle` (Simple-Networked Interface the the Full Functionality of Laboratory Equipment) is another closely related project [7]. `sniffle` enables WWW clients to remotely control an arbitrary device connected to an analog/digital I/O board installed in a PC. `sniffle` has been applied to the problem of remotely controlling an oscilloscope device (and obtained data back from it) and is also used to pan and tilt a digital camera that provides real-time video feedback. One of the most important contributions of the `sniffle` work is that its author carefully documented the system he built, enabling others to duplicate his efforts rapidly. While `sniffle` provides the basic infrastructure needed by UT, it does not in itself consider issues related to mobility or possible applications and research issues related to ubot devices.

4.1 Related Technologies

Ubiquitous telepresence is related but separate from ubiquitous computing [16] in the model of how telepresence is provided. In particular, ubiquitous telepresense focuses on providing telepresence by defining interfaces, software, and hardware for standard low-cost robotic devices (ubots) that are connected to the World Wide Web. Instead of assuming the ubiquity of many different computing devices permeating an environment, ubiquitous telepresence requires only the presence of a single inexpensive ubot to realize its power.

5 Implementation Issues

The purpose of this section is not to fully describe a working UT architecture, but to raise some of the important technical issues that need to be addressed to achieve UT technology. Subsequent papers will provide many more details about our actual implementation and the technical problems that we faced and solved.

- The combination of mobility and vision raises the important implementation issue of how to provide enough bandwidth for effective video on both the wireless link and across the internet. One interesting aspect of this problem that is highlighted by our application is that of the tradeoff between resolution and frame-rate. Because the ubots will be performing different kinds of tasks at different times (e.g., "move to the next room" versus "take a close look at this Van Gogh painting") a flexible solution to the video bandwidth problem is highly desirable.

- How autonomous will the ubot be? To achieve optimal leverage of the human user's time, the ubot would be capable of performing such general commands as "go to THAT room" (where THAT is perhaps indicated by the controller via a mouse click). The ubot would then need to be somewhat autonomous for simple motion tasks. In any event, it will need to have enough embedded intelligence to interact safely with the humans and objects around it.

- How will multiple users control a ubot or set of ubots? None of the current internet devices have particularly interesting ways of handling multiple users, but as the number of such devices increases, problems related to collaborative use of remote devices will need to be solved.

- How will users interact with the ubots (and other users)? Issues of what the user-interface to the device looks like need to be addressed. Again, existing systems provide some solutions to this problem, but currently none support problem domains as complex as how to maneuver around a building.

- How to prevent unauthorized use or unsafe use of the ubots? This problem starts sounding like science fiction when one considers that a ubot in a home with a manipulator arm could be directed to commit a crime remotely. Fortunately, controlling unauthorized access to information is an equally important problem, and with the commercialization of the internet, is getting a lot of attention already. Our hope is to use the most effective approaches to access control, and extend them into the domain of ubot control.

- How to deal with stairs? This mundane problem is very important with respect to getting the mobile robot technology to be ubiquitous. The simplest solution (also quite an expensive one) is to buy two (or how many separate floors you have). This solution requires developing a means of coordinating them, and has the obvious cost disadvantage.

- How to supply power? We need to explore what the power requirements of the ubots will be and what constraints the existing rechargeable battery technology will impose on the size, mobility, and lifetime between recharges of the ubots will be.

6 Summary

In this paper we have outlined a line of research that we believe will lead to the creation of low-cost mobile robotic devices (ubots) that can be easily attached to and remotely operated via the World Wide Web. We call this technology ubiquitous telepresence, and in this paper we have outlined some interesting applications of UT, including remotely viewing a museum. Ubiquitous telepresence is related to the many existing projects that attach devices to the WWW, and we described some of the most closely related projects and discuss how they relate to ubiquitous telepresence. We also have outlined some of the research challenges in implementing ubiquitous telepresence. For current information about our project and related work, please visit our WWW site at http://www.cs.colorado.edu/homes/zorn/public_html/ut/Home.html

Acknowledgements

I would like to thank the participants in the UT project Jeanine Cook, Jonathan Cook, Michael Doherty, Matthew Greene, Bryan Loughry, David Keaton, Carlos Matzahn, Christian Och, Matthew Seidl, and William Waite for their interest and support. I would also like to thank Alan Kaplan and Norm Schryer of AT&T Bell Laboratories for being so supportive of our work. Eric Paulos and John Canny also provided insightful feedback.

References

[1] Vanu Bose. The tns video rover. Work by the Telemedia, Networks And Systems Group, MIT. Available at http://www.tns.lcs.mit.edu/vs/rover/rover.html.

[2] John J. Fisher. Applying robots in nuclear applications. In *Teleoperated Robotics in Hostile Environments*, 1985.

[3] R. F. Fogle. The use of teleoperators in hostile environments applications. In *Proceedings of the 1992 IEEE International Conference on Robotics and Automation*, Nice, FRANCE, May 1992.

[4] Ken Goldberg, Michael Mascha, Steven Gentner, Juergen Tossman, Nick Rothenberg, Carl Sutter, and Jeff Wiegley. Beyond the web: Excavating the real world via mosaic. In *Second International WWW Conference*, Chicago, IL, October 1994.

[5] Ken Goldberg, Joseph Santarromana, George Bekey, Steven Gentner, Rosemary Morris, Carl Sutter, and Jeff Wiegley. A tele-robotic garden on the world wide web. *SPIE Newsletter*, Spring 1996.

[6] A.E. Kaplan, S. Keshav, N.L. Schryer, and J.H. Venutolo. An internet accessible telepresence. *Multimedia Systems Journal*, 1996. To appear.

[7] Terrence Kelly. sniffle: A simple-networked interface the the full functionality of laboratory equipment. Available at http://www.cs.princeton.edu/ tpkelly/sniffle/, September 1995.

[8] Won S. Kim. Graphical operator interface for space telerobotics. In *Proceedings of the 1993 IEEE International Conference on Robotics and Automation*, 1993.

[9] Learning Robotics Lab. Where in the world is xavier, the robot? Available at http://www.cs.cmu.edu/afs/cs.cmu.edu/Web/People/Xavier/.

[10] Steve Mann. Wearable wireless webcam and telemetry. Available at http://www-white.media.mit.edu/ steve/netcam.html.

[11] NASA. Nasa space telerobotics program. Available at http://ranier.oact.hq.nasa.gov/telerobotics_page/telerobotics.html, February 1996.

[12] Eric Paulos and John Canny. A world wide web telerobotic remote environment browser. In *Fourth International World Wide Web Conference*, December 1995. Also available at http://www.cs.berkeley.edu/ paulos/papers/www4/.

[13] Eric Paulos and John Canny. Ubiquitous tele-embodiment: Applications and implications. *International Journal of Human-Computer Studies*, 1997. To appear.

[14] YAHOO WWW Server. Interesting devices connected to the net. Available at http://www.yahoo.com/Computers_and_Internet/Internet/Interesting_Devices Connected to_the_Net.

[15] Ronald J. Vetter. Computer-contrlled devices reach the internet. *IEEE Computer*, pages 66–67, December 1995.

[16] Mark Weiser. Some computer science issues in ubiquitous computing. *Communications of the ACM*, July 1993.