



ELSEVIER

Computer Networks and ISDN Systems 28 (1995) 209–219

COMPUTER
NETWORKS
and
ISDN SYSTEMS

Beyond the Web: manipulating the real world

Ken Goldberg ^{a,*}, Michael Mascha ^b, Steven Gentner ^a, Jürgen Rossman ^c,
Nick Rothenberg ^d, Carl Sutter ^e, Jeff Wiegley ^a

^a Department of Computer Science, University of Southern California, Los Angeles, CA 90089, USA

^b Department of Anthropology, University of Southern California, Los Angeles, CA 90089, USA

^c University of Dortmund, Dortmund, Germany

^d Department of Visual Anthropology, University of Southern California, Los Angeles, CA 90089, USA

^e Center for Scholarly Technology, University of Southern California, Los Angeles, CA 90089, USA

Abstract

This paper describes a WWW site that allows users to “leave the Web” and interact with the real world. An interdisciplinary team of anthropologists, computer scientists and electrical engineers collaborated on the project, designing a system which consists of a robot arm fitted with a CCD camera and a pneumatic system. By clicking on an ISMAP control panel image, the operator of the robot directs the camera to move vertically or horizontally in order to obtain a desired position and image. The robot is located over a dry-earth surface allowing users to direct short bursts of compressed air onto the surface using the pneumatic system. Thus robot operators can “excavate” regions within the environment by positioning the arm, delivering a burst of air, and viewing the image of the newly cleared region. This paper describes the system in detail, addressing critical issues such as robot interface, security measures, user authentication, and interface design. We see this project as a feasibility study for a broad range of WWW applications.

Keywords: Robot; Remote control; Mercury Project; USC; Random token; Telerobotics

1. Goals of the project

The WWW provides a multi-media interface that spans all major platforms. Thousands of sites have been set up in the past year. Our goal with this project was to provide public access to a teleoperated robot, thus allowing users to reach beyond the digital boundaries of the WWW.

Such a system should be *robust* as it must operate 24 hours a day and it should be *low in cost* (we had an extremely limited budget). It is worth noting that the manufacturing industry uses the same criteria to evaluate robots for production. Thus our experience with RISC robotics (see below) proved helpful.

Our secondary goal was to create an evolving WWW site that would encourage repeat visits by users to collectively solve a puzzle. This paper focuses on the details of the implementation. The Mercury Project is archived with supporting examples and statistics at:
<http://www.usc.edu/dept/raiders/>

* Corresponding author. E-mail: goldberg@ieor.berkeley.edu. <http://www.usc.edu/users/goldberg>. This work was supported in part by NSF Awards IRI-9123747 and IRI-9457523.

2. Related work

The first “teleoperated robots” were developed over 40 years ago. The basic objective has always been to develop systems capable of working in inhospitable environments (such as radiation sites). Teleoperation began with very simple mock-ups in nuclear power plants [6], progressing to more versatile setups for teleoperation of robots in space [7]. Over the last 20 years, the development of intuitively operable teleoperation tools has continued to play an important role in the development of robotics in general. The basic objectives have remained the same, even though the methods and technical limitations have changed. See also [4].

Today, sophisticated “Telerobot Operator Control Stations” [3] are equipped with stereoisometric-displays, “force reflecting hand controllers” and comprehensive video graphics support. The development of teleoperation stations is currently being pushed further with the help of latest graphics workstations to provide so-called “telepresence”. Modern telepresence systems, considered to be pushing the frontier of research in this field, are defined as follows [1]: “At the worksite, the manipulator has the dexterity to allow the operator to perform normal human functions. At the control station, the operator receives sufficient quantity and quality of sensory feedback to provide a feeling of actual presence at the worksite”.

The Mercury Project does not achieve this level of telepresence but provides a limited level of teleoperation. One of our goals was to provide “teleoperation for the masses”. Instead of developing a highly sophisticated, multi-million-dollar testbed, we opted for a simple and reliable end-effector on a commercial robot. Combined with an intuitively operable man-machine-interface, the system gives all WWW users access to teleoperation. In the Discussion section, we describe a number of other WWW sites that offer interactive capabilities.

3. User interface and environment design

The interface design for the system was challenging due to the limitations of the HTML/HTTP environment, as well as network traffic considerations. An effective system was created within such limitations by carefully designing the physical environment for the robot, and by fine-tuning the user-machine interface. For example, the initial idea of a live video feed from the camera was dropped in order to maintain compatibility with all visual clients on the Web. (Although we could have implemented some custom clients¹, we decided to stay within the limits of HTML/HTTP to reach as large a user base as possible, making this a truly global system.) In addition, initial simulations using a robot fitted with grippers (simulated in VIRTUS WALKTHROUGH) revealed a high degree of

¹ There are two possible fixes to this problem, One is to release specially modified clients that set up a two-way communication, the second is to use some other software to display the current system on the user's client workstation. Since many clients are used to view the WWW, making modifications would be difficult, especially since they are being updated all the time. Even if source code could be obtained for every major client, changes would have to be made to every release of all these be on each release of these applications, The second possibility is to write a separate program to run on the clients' workstation. The problem here is to write a robots client that can be released for enough platforms to be useful. Since this would be an esoteric piece of the system, it is not likely that other sites would customize the software for different systems like is done for the major systems. One technique is to use the X windows protocol to display a client application on the users workstation running an X server (weather, movies). We felt that this would be a limited audience, however. It also may compromise security from the user's point of view. Both these approaches may be attempted in version 2.0 of the system to allow more enhanced use of the system for some users. The HTTPD protocol could be extended to allow these sort of connections, though—maybe we need a new protocol for passing media only back that doesn't have all the hooks into system calls like X Windows and Display PostScript.




Fig. 1. The control panel.

complexity in control functions², not suitable for the anticipated 5–10 seconds per frame page loading time, a 2D Mosaic window and a naive/untrained user.

The team chose instead to use a simple environment which would allow relatively easy control of the robot. Here the analogy taken from real world archaeology—using a dry-earth environment and compressed air bursts—allowed us to simplify the robot control dramatically. Thus users could be quickly trained in the operation of the system, through a simple “Operator’s Orientation” and a “Level 1 Clearance Test”.

Even with a simplified system (see Fig. 1), users are still able to choose between fine and gross movements of the arm. Fine pitch movements are executed by clicking in the camera image, with the robot moving to center the arm over the X, Y -coordinates of the click-point. Crude navigation is provided by clicking on a schematic picture of the robot and its workspace, with the robot moving to center the arm over the click-point. Two buttons allow navigation in the Z -axis (between “up” and “down” positions), with a button to blow air only active when at the $Z = 0$ (i.e., “down”) position.

Other features of the system were designed to balance functionality with user needs. All HTML documents sent to the clients are carefully designed to minimize network traffic in order to get a high refresh rate. For example, control panel functions are clearly distinguished from text-based information documents. The “Operator’s log” was implemented to create a forum for collaborative efforts to solve the puzzle/problem regarding the underlying logic which links the artifacts. (The “Operator’s log” is readable throughout the system but only writeable after completing an operating session.) A **second entry path**³ was also created to the system, which provides a “back-story” explaining the project while also hinting at possible “real world” uses of the system.

4. Access to the robot

Most of the HTML documents seen by the user on our site are generated by a script running on the WWW server. Using a random token scheme described below, the system tracks each user as he or she proceeds through the interface and generates appropriate HTML documents. This allows the system to discriminate between

² 3D control of a robot needs: 3 dimensions of spatial movement, 3 dimensions of orientation and 1 to 3 dimensions of gripper control.

³ <http://www.usc.edu/dept/raiders/story/mercury-story.html>

“observers” and “operators” so that it presents only accessible options to each.

To operate the robot, the user must read the information on how to use the control panel, and then complete a level-1 clearance test to get a password. Since only one person can operate the robot at a time, the system maintains a queue of pending operators. A typical user will enter his/her password, and then add him/herself to the queue. Each time the update button is clicked, the system updates the queue and returns a current status page. When the user’s turn arrives, the screen returned is the live operators’ control screen.

5. System architecture

Fig. 2 shows a block diagram for the system. We start with an overview that necessarily glosses over many interesting details.

At one end are WWW clients from around the world; at the other end is a robot arm combined with a camera. The robot and camera provide an updated image of the environment, which is combined with a schematic of the robot arm/workspace and control buttons to produce the final GIF image that is sent to users.

At any given time there may be dozens of clients interacting with the system. Since there

can only be one operator at a time, one challenge is to keep track of which client is the operator.

The Mercury system is comprised of three communicating servers. The first, call it *A*, is a standard Mosaic server (NCSA httpd v.1.3, currently running on a Sun SPARCserver 1000, with SunOS Release 5.3. When the Site is requested by an observer, the most recent image, which is stored on server *A*, is simply returned.

The database of registered users is handled by another server, call it *B*. In our case, server *B* runs on the same machine as server *A*. The database server is custom programmed for this project, but performs fairly standard database functions.

When a client request comes in, server *A* communicates with server *B*. If that client is an operator, server *A* must then communicate with a third server, call it *C*, that controls the robot. Server *C* runs on a Pentium-based PC and communicates with servers *A* and *B* via the Internet. Server *A* decodes the ISMAP *X* and *Y* mouse coordinates, and sends them to server *C*.

On server *C*, a custom program decodes these coordinates into a robot command and verifies that the command is legal, e.g., within the robot workspace. If it is, this command is then converted into a robot command format which is sent to the robot over a serial line. Once the robot move is completed server *C* uses the CCD cam-

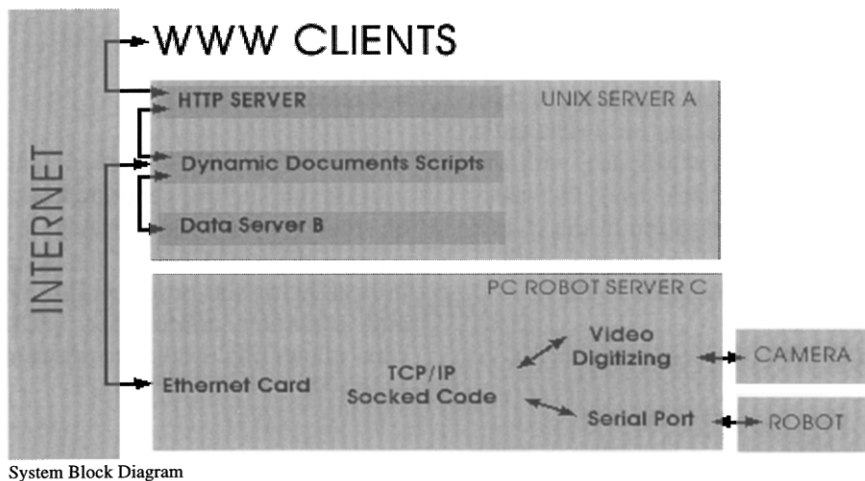


Fig. 2. System block diagram.

era to capture a stable 8 bit 192×165 image of the workspace.

Using a simple set of equations for inverse kinematics server *C* then generates a schematic view of the robot in its new configuration. This schematic is combined with the camera image, and the up, down, and air control buttons to form a new composite image. Server *C* then compresses this image into GIF format and returns it to server *A*, which updates the most recent image and returns it to the operator client.

6. Subsystems

6.1. Random tokens for cache avoidance and user tracking

Following some complex and unwieldy tests, we implemented a random token scheme for tracking users as they use the system. Each time a URL is returned, a large random number is added to the path (which the NCSA HTTPD 1.3 server splits into the `PATH_INFO` environment variable). This “token” serves two purposes:

The first is to prevent the WWW client from caching the robot view. When a document is requested a second time during a session, it is much faster to swap in a local copy of the document rather than going back over the net to retrieve it a second time. Most implementations of Mosaic support such caching at various memory levels. However in our case we want to repeatedly retrieve the URL containing the robot image because it is updated continuously. In brief, we don't want users to cache this URL. The random token makes each request look different and tricks the client into retrieving a fresh version of the document.

The second use for the token is to identify operators. When an operator logs in with a correct password, the system begins tracking him/her as he/she moves from viewing the robot to being on the operators' queue to operating the robot. Since the same script is used for all views, the token allows the system to customize the result for every user depending on his/her position in the system.

6.2. Scripts

The robot view screen is controlled for the most part by one script at the HTTPD server. Each call to the main script has a token attached to the URL. The token is decoded by the WWW server, and placed in the `PATH_INFO` environment variable. The main script then checks the token with the database server to determine the status of the user. Each check of the database generates a system update to keep the queue moving. The user's status is used to generate the custom system status page.

The robot image itself is only changed by the operator when he or she makes a move. Each image is date and time stamped, so WWW clients that cache the image will only retrieve the image when it changes (since its filename will be different, due to a different time stamp).

Due to the client-server architecture of the World Wide Web HTTPD protocol, the robot system (server) has no way to contact the client except at the client's prompting. From the user's point of view, once he or she gets the robot view screen, there is no way for the server to keep sending updates automatically as the robot is moved by the operator. The screen updates must be driven by the user. Since the user must trigger each update, we wanted to provide a button for doing so, since each web client handles reloading the page differently. Some sites have a “reload” hypertext link to the same page, but this doesn't work for any client that caches pages. If a page is being viewed, hitting reload will just re-display the page from the cache, thus not obtaining a new view from the system. Asking the user to disable his/her cache is also problematic, since not all clients allow this option.

One attempt was made to use a mini-form, since the submit button always calls a script and is not cached. That scheme was eventually dropped, since passing registered user identification information to the server via hidden fields only worked on some clients. Using the random token allows for an elegant interface.

Since the robot can only be controlled by one person at a time, a registration scheme was implemented to allow the server to track operators

as they move on to the waiting queue and progress to controlling the robot. Since the server only knows the IP address of each user, some user information had to be incorporated into the HTML robot view document itself for re-transmission to the system when the user hits “reload”. There are various techniques used by many sophisticated web systems to accomplish user identification between document requests, but we found some problems in many of the standard solutions. In the end, the random token served excellently as a means of identifying registered users.

A preliminary attempt was made to use a small form to identify the user. Hidden fields could hold the user id, but once again, many clients do not implement the hidden field attributes so the interface is cluttered by unnecessary fields. Putting the user’s id information into the ACTION field of the form tag is also client dependent. Unfortunately, some clients strip that data before adding the encoded field information.

Since random tokens were already being passed with each update, the system was extended to track the tokens of each registered user. Each time the script is called, the token is exchanged for a new one, and the database is updated with the new token for registered users. One side effect is that the user cannot use the client re-load button, since this will not use the new URL (it is embedded in the update HREF).

6.3. The data server

The data server (*B*) is a custom Perl script that handles all of the database work for the project. It continuously runs as a TCP/IP listener, waiting for database transactions from the other system scripts. The data server runs as a single process, handling requests serially to maintain internal data integrity. Typically, transactions are very short, since the data is kept in main memory. The data server could be replaced by an off-the-shelf transaction based DB system in the future. A time-out is set to close the connection if there is too much time elapsed between commands. This was implemented because some WWW clients would crash in the middle of a document

request, leaving the system waiting for the connection to be closed.

6.4. Internal network interface

The networking functionality required by the project was defined by two factors. On one hand, the camera that we purchased required a PC-based platform running an MicroSoft DOS or compatible operating system to run on server *C*. On the other hand, the expected load of client requests required a machine capable of more heavy networking duties such as a Sun workstation (server *A*). Currently server *A* is located across campus from server *C*.

These servers are connected via Ethernet. Each machine has its own IP address and resides in the usc.edu domain. Communication is achieved using a socket connection between the two machines. The implementation on server *A* was done using the standard BSD socket functions provided with the SunOS 4.1 operating system and Perl. On server *C* we used a publicly available socket package called Waterloo TCP and Borland C. The Waterloo TCP package was obtained from the ftp site dorm.rutgers.edu in the file /pub/msdos/wattcp/wattcp.zip.

With this software server *A* can request a socket connection to server *C* to establish a connection. The first step in obtaining a new image is for server *A* to write a command consisting of thirty bytes which encodes the (*XY*) coordinates of the ISMAP event. After server *C* completes the moves and generates the new image, it writes the size of the new image to server *A* so that server *A* knows exactly how many bytes to expect. Server *C* then proceeds to write the entire image to the socket and waits for the socket to close to insure deliver of the data. Once server *A* has read all the specified bytes it closes the socket. Server *C* is now ready and waiting for another socket connection. Server *A* is free to continue processing the Mosaic actions of the current users.

Current throughput is approximately 20 Kbytes/second, which is poor compared to the 0.5 megabyte per second rate that can be achieved between two Sun workstations in close proximity on the campus network. At this time we feel that

the delays are being imposed by the MS-DOS operating system because of its inability to support networking operations and its lack of multi-tasking abilities, which necessitates busy waiting cycles in the PC software to obtain concurrence between the robotic/camera operations and the networking duties.

Our low data rate is somewhat tolerable because the time for communication between servers *A* and *C* is small compared with Internet delays between clients and server *A*. One way to speed communication would be to use different methods of image compression such as JPEG to reduce the size of the image. However this may introduce latency due to encoding.

6.5. The IBM robot and server *C*

The robot we're using is an IBM SR5427 SCARA arm, built around 1980. SCARA stands for "Selective Compliance Assembly Robot Arm". Robots with SCARA kinematics are common in industrial assembly for "pick-and-place" operations because they are fast, accurate and have a large 2.5D workspace. However, the SCARA arm can only rotate its gripper about the vertical (*Z*) axis. We selected this robot over other robots in our lab due to its excellent durability, large workspace, and because it was gathering dust in the Robot Education Lab.

The IBM SCARA robot is controlled through a 4800 baud serial port by a custom written C library constructed with reference from IBM's BASIC library distributed along with the robot. The commands sent by the library are simple instructions consisting of instruction id, length, data and checksum. The data length and content varies depending on instruction id. The IEEE floating point format is used to represent the necessary data. This command string is then sent over the serial line to the robot to issue the command.

Unfortunately IBM no longer supports this robot and we were forced to read two antiquated BASIC programs and monitor their serial line transmissions to decipher the protocols needed for serial control of the robot. The robot accepts XYZ and Theta commands using IEEE format

and checksums. Server *C* now runs on a Pentium based PC with all custom code written in Borland C.

The first step was implementing a local graphical user interface to control robot movements and monitor subsequent functions such as network flow. We chose two views of the workspace: a global schematic view for coarse motions, and a local camera view for fine motions. Note that a click on the camera image requires a different relative move if the camera is in the up or down position. To handle it, we implemented an empirical calibration program.

The major difficulty in implementing server *C* was to schedule response to the network, the local mouse, and the image capture board. At first we discussed a multi-tasking environment but, upon further study, we realized we could achieve this cooperation within a single DOS task. Another problem, inherent to DOS-based applications, is memory management. This complication was solved by careful usage of memory and by utilizing the screen itself as a memory buffer. This careful usage of memory enabled the custom written GIF encoder to use more memory which, combined with an appropriate hash function, sped the GIF encoding process up to a few microseconds.

In future versions of Mercury, we plan to incorporate a more sophisticated PC-based robot simulation system based on COSIMIR [8] from the University of Dortmund.

6.6. Camera

We are using an EDC 1000 digital CCD camera from Electrim Inc. This camera was chosen based on size and cost. Image data is sent from the camera back through a serial line into a video capture card. The picture captured is always 192 × 165 pixels with 256 shades of gray. The image size and gray shades are fixed. Focus and contrast are manually adjusted. Exposure time can be changed by software to range between 1/200th to 1/64th of a second. 1/150th exposure was used to reduced light streaking that the camera is prone to.

Although the robot's control system quickly dampens oscillation about the destination point, dynamic effects can cause image blur. Two solutions were implemented. First the robot was slowed down enough as to reduce some of the vibration but not to hinder the robot access speed considerably. Second, once the robot responds positively to an issued command, the camera captures two pictures each at 1/64th of a second. These two images are compared to determine a factor of similarity. If this factor is below some set value the image is presumed to be stable, otherwise subsequent pictures are taken until the image pair is determined to be stable. More than 5 trials results in a time-out in which case the most current image is used and the program continues. This image comparison procedure reduces movement streaks seen in pictures of moving objects.

Lighting the workspace has been problematic. The work space is primarily illuminated by standard florescent ceiling fixtures and augmented by two additional florescent lamps to reduce shadows and raise the overall ambient light levels. We tested a contrast enhancement routine to normalize the lighting of each image captured from the camera. This increased the visual aesthetics of the image but subjected it to drastic light and dark changes as the robot moved onto different objects with different light reflecting qualities. In response, a global lighting adjustment was implemented but found to reduce certain areas to unacceptable darkness. Certainly a better lighting system is required.

Due to the manual focus adjustment of the camera, the focus adjustment could not be changed between the up and down position of the camera. This resulted in a compromise focus adjustment that is not perfect for the up or down position of the robot arm, but acceptable in both positions.

To decrease compressed image size and thus increase network transfer rate the image is reduced from 256 to 64 gray scales since most systems available can only display 256 colors or 64 shades of gray. Thus the gray scale reduction did not reduce image quality but reduced compressed image size by about 10K.

6.7. Robustness and soft resets

All robot motions are monitored by server *C*. Each command sent to the robot is verified to be within the robot's workspace. Acknowledgments from the robot are monitored to detect errors. When an error is detected, server *C* automatically resets the robot controller, recalibrates, and returns the robot to its previous position.

7. Performance

7.1. History and statistics

Daily statistics are available and may be correlated with **project milestones**⁴. The site received over 2 million hits from over 50000 unique sites during its active period from September 1994 to March 1995.

7.2. Refresh rates via ethernet

System response time seems to be mostly dependent on network link speeds. Locally, we get screen refreshes at rates of 5–10 seconds per page. Similar response times have been reported from Europe. Obviously, a slow local link or SLIP connection will drastically affect the update speed, since the robot control image is essential to the system. Updates are also strongly affected by the speed of the WWW client application.

7.3. Uptime

The system is designed for 24 hour use. The WWW server scripts are generally modified, tested and then loaded into the running system. Background programs monitor the system and notify the team members if there are problems.

7.4. Operators' logs

When an operator has finished driving, he or she is prompted to make a textual entry into an "Operator's log". The **Operator's log**⁵ provides

⁴ <http://www.usc.edu/dept/raiders/>

⁵ <http://www.usc.edu/dept/raiders/>

an ongoing forum for discussion of the system and record of artifacts discovered in the sand.

For example, several skeptics have claimed that the system is an elaborate hoax where all images are taken from a prestored library (much like the celebrated Apollo Moonwalk hoax of 25 years ago). We have had encouraging comments from the robotics community, including several researchers at NASA.

8. Discussion and future applications:

This project is an initial step in an ongoing educational and research project at the University of Southern California. It brings together faculty and students of different backgrounds to collaborate in the design and implementation of a networked system that combines robotics with archaeology and interactive art.

This system exemplifies RISC Robotics, which advocates Reduced Intricacy in Sensing and Control. The SCARA-type robot requires only 4 axes, is relatively inexpensive and robust, and it is easy to avoid singularities. The end effector we've used here is also about the minimum. For more on RISC as applied to industrial robotics, please see [2].

We see the project leading in several directions. For Mosaic and the WWW, the required interface design prompted new developments related to several issues, including user authentication, user queuing and interface security (as discussed above).

For this project we chose a very simple application. The server can be extended to a variety of platforms that permit remote inspection and manipulation of objects—for example, providing unique and unedited access and views of priceless and otherwise inaccessible resources (a Grecian urn, a Gutenberg Bible, etc.), thus providing an alternative to pre-stored libraries which are limited in terms of perspective, depth of resolution, etc. For example, see <http://vive.cs.berkeley.edu/capeK/>

Further extensions for this project might include: the robot could be placed out in the field, in a remote anthropological site or on the moon;

the camera could be replaced with a scanning electron microscope; or the remote operator could be a doctor examining a patient or a specialist performing remote inspection or manufacturing. All of these areas also have significant implications for education, as they present the opportunity for virtual “field trips” to a live site while permitting remote manipulation from the classroom.

Anthropologists have conventionally recorded the diverse cultural heritage of humankind by means of varied media: written text, graphics, film, sound and still images. The advantage of a system like the one described in this paper lies in the fact that you do not have to rely on pre-recorded media. It enables the user to view and possibly record her or his own “slice of reality”. We see the Web as a perfect medium for updating pre-recorded media as described in [5]. Furthermore, we now have the possibility to combine updateable prerecorded media of all sorts with live recordings and live remote interactions. The possibilities of a system that combines global access to up-datable prerecorded media and combines it with the possibility of live remote interactions are just beginning to unfold.

In view of other interactive WWW sites, we propose a three-tiered system describing interactively on the WWW. Under Level I, interaction is solely between digital information stored on computers or created by scripts running on such machines, and connected or communicating with the WWW and Mosaic clients. In Level II sites, the clients are able to observe the “real world” by means of a camera observing and digitizing visual and, hopefully soon, audio-visual information. The camera acts as an “eye” for the Web, providing multiple “real world images” from a global theater. A number of Web sites fall into this category, such as the **Coffee Pot**⁶ and the **Fish-tank**⁷ sites. All have the same characteristic of passively observing the real world. We also know of one restricted site that allows the user to alter

⁶ <http://www.cl.cam.ac.uk/coffee/coffee.html>

⁷ <http://elmendorf.camb-lads.loral.com:8080/~rg/aquarium.html>

the user's point of view (see **LEGO pan and tilt site**⁸).

The Mercury Project introduces a third level. Level III sites reach beyond the digital domain to allow users to alter a remote physical environment. We envision this project as a first glimpse into the possibilities available at Level III. We might also speculate about other levels, which might allow remote users to control a mobile robot and thus “tele-ambulate”.

Acknowledgements

Rick Lacy, Mark Brown, and the Center for Scholarly Technology, Departments of CS and Anthro, Prof. Peter Danzig (CS), The alpha and beta testers, The Laika Foundation, The Los Angeles Museum of Miniatures.

References

- [1] D.L. Akin, M. Minsky et al., Space applications of automation, robotics, and machine intelligence systems (ARAMIS), Phase II, Vol. 3, M.I.T. Contract NASA 8-34381, NASA Marshall Space Flight Center.
- [2] K. Goldberg and J. Canny, **RISC for Industrial Robots**⁹: Recent results and open problems, *1994 IEEE Conf. on Robotics and Automation*.
- [3] E. Kan, J. Tower et al., The JPL telerobot operator control station: part 1 – hardware, *Proc. NASA Conf. on Space Telerobotics*, 1989.
- [4] B.J. Makinson, Research and development prototype for machine augmentation of human strength and endurance, Report S-71-1065, General Electric Company, Schenectady, NY, 1971.
- [5] M. Mascha and G. Seaman, **Interactive education**¹⁰: transitioning CD-ROMs to the WEB, *Comput. Networks ISDN Systems* 27 (1994) 267–272.
- [6] R.S. Mosher, Force reflecting electrohydraulic servomanipulator, *Electro-Technology* (1960) 138.
- [7] N.J. Mizen, Preliminary design for the shoulders and arms of a powered exoskeletal structure, Cornell Aeronautical Laboratory Report VO-1692-V-4, 1965.

- [8] E. Freund, J. Rossmann, J. Uthoff and U. van der Valk, Towards realistic simulation of industrial robots, *Proc. IEEE/RSJ/GI Intelligent Robots and Systems*, IROS, Muenchen, 1994.

Kenneth Y. Goldberg received dual degrees (B.S. in Electrical Engineering and B.S. in Management) from the University of Pennsylvania in 1984. Working with Matt Mason at the School of Computer Science at Carnegie Mellon University, he completed the M.S. and Ph.D. degrees in 1987 and 1990 respectively. He also studied at Edinburgh University and the Technion. In January of 1991, he joined the faculty at USC with a joint appointment in Computer Science and EE-Systems. Currently he is Director of the Modular Robotics Lab and Associate Director of the Institute for Robotics and Intelligent Systems, where he leads a research group at USC that studies algorithms for feeding, sorting, and fixturing industrial parts. In September 1995, he will join UC Berkeley's Department of Industrial Engineering and Operations Research. Prof. Goldberg serves on the Advisory Board of the IEEE Society of Robotics and Automation and is a National Science Foundation Young Investigator. Prof. Goldberg's dissertation on geometric algorithms for manipulation was nominated for the 1991 ACM Dissertation Award and resulted in two U.S. patents. He leads a research group at USC that has published over forty papers on algorithms for feeding, sorting, and fixturing industrial parts. He emphasizes mathematically rigorous solutions that require a minimum of sensing and actuation so as to reduce costs and increase reliability on the factory floor. This research is sponsored by the National Science Foundation and Adept Technology. He is currently working closely with Adept, Silma, Inc, and a consortium of companies organized by the National Center for Manufacturing Sciences to combine modular tooling with simulation software that can dramatically reduce the time it takes to design and install assembly lines. Goldberg co-directed two projects combining tele-robotics with the WWW (Internet). The “Mercury Project” was a finalist for the 1995 National Information Infrastructure Award, and the “Tele-Garden” won the Kobe Prize at the Interactive Media Festival in 1995. Prof. Goldberg was elected a Senior Member of the Society of Manufacturing Engineers. He has served as reviewer for a number of journals and conferences and organized workshops at the past two IEEE Robotics and Automation Conferences. At the 1994 Conference, a paper written with his students on Modular Fixturing won the Best Paper Award.

Michael Mascha. Born in Vienna, Austria (1958). Ph.D. in Anthropology and Communication Sciences at the University of Vienna (1985), followed by a post graduate course in audio-visual technologies at the Austrian Institute for Scientific Cinematography (1985–1986). Three years field work in Fiji for a series of films, in the position of director/scientist. Director of the First International Documentary Film Festival Vienna (1990–1992). Joined the Center for Visual Anthropology at USC as Research Associate for New Media (1992).

⁸ http://www.ludvigsen.dhhalden.no/webdoc/pt/pt_dummy.html

⁹ file:///palenque.usc.edu/pub/Tech_Rept/Iris/1993/IRIS_93_315.ps.Z

¹⁰ <http://elab-server.usc.edu/papers/www94.html>

Lecturer of Interactive New Media and appointment (1993) as the first Director of the E-Lab, a newly established teaching and research facility for New Media at the Center for Visual Anthropology. Promotion to Assistant Professor (1994) at the Department of Anthropology, teaching, developing and researching Interactive New Media in anthropology and communications. First Webmaster for the E-Lab World Wide Web server and developer of WWW based curriculums for visual anthropology and New Media. Founder and maintainer of the Virtual Library for Anthropology and Associate Editor for the Social Science Computer Review (Anthropology). Co-Director of the Mercury Project—the first robotic tele-excavation unit on the World Wide Web. Currently Creative Director and a founding partner in W3-design, Inc. (<http://www.w3-design.com>), a premiere Web site development and design company in Los Angeles.

Steven Gentner was born in Cardiff, England. He was raised in Nigeria, Africa, where he attended the American International School. He continued with his high-school education by attending TESIS (The American School In Switzerland) located in Lugano, Switzerland. With the goal of becoming an architect he started to acquire an extensive art background. However, midway through high-school, he decided to change his focus from art to that of computer science. He quickly became attached to the flexibility and challenge of working with computers and furthered this interest at the University of Southern California. He completed by Bachelor of Science in Computer Science and continued at USC for one more year to acquire a Masters in Computer Science. Upon graduating he is pursuing work involving the World Wide Web.

Nick Rothenberg graduated from Harvard University in 1987, where he studied ethnographic filmmaking with Robert Gardner, and received a Master's Degree in Film and Anthropology from The University of Southern California (USC). He was Assistant Director for the E-Lab, a multi-media lab out of USC's Center for Visual Anthropology, from 1992–1994. With Mercury Project colleague Dr. Michael Mascha, he formed a Web site design company in 1994, W3-design Inc. (<http://www.w3-design.com/>).

Carl F. Sutter is the Senior Programmer/Analyst at the University of Southern California's James Irvine Foundation Center for Scholarly Technology (CST). He holds a Bachelor's Degree in Electrical Engineering from Cornell University, and a Master's Degree in Electrical Engineering from the University of Southern California. At the CST, he analyzes new technologies and provides advanced training and advocacy on the scholarly use of technology.

Jeffrey Wiegley is a Ph.D. student in Computer Science at the University of Southern California. He has been working with Ken Goldberg now for about three and half years on various projects, the latest of which has been the TeleGarden Web Site. He has a B.S. in Computer Science from the State University of New York at Buffalo. His specialties are Unix System administration and programming as well as Mechanical design and fabrication. He is currently beginning work on his Ph.D. dissertation on parts feeding for industrial applications and hopes to complete his degree by the end of 1996.