The Tele-Reporter: A Collaborative Interface for Goal-Oriented Teledirection

by

Derek J. Tang

Submitted to the Department of Electrical Engineering and Computer Science

In Partial Fulfillment of the Requirements for the Degree of

Master of Engineering in Electrical Engineering and Computer Science

at the Massachusetts Institute of Technology

May 24, 2002

Author_____
Department of Electrical Engineering and Computer Science
May 24, 2002

Certified by_____
Judith Donath
Thesis Supervisor

Accepted by_____
Arthur C. Smith
Chairman, Department Committee on Graduate Theses

The Tele-Reporter: A Collaborative Interface for Goal-Oriented Teledirection

by

Derek J. Tang

Submitted to the Department of Electrical Engineering and Computer Science

May 24, 2002

In Partial Fulfillment of the Requirements for the Degree of

Master of Engineering in Electrical Engineering and Computer Science

ABSTRACT

The Tele-Reporter is an extension of the Tele-Actor project, a Web interface where users share control over a human actor agent. The Tele-Reporter extends the premise of the Tele-Actor by allowing users to view a live news broadcast and collaboratively "teledirect" the reporter.

The Tele-Reporter addresses many of the concerns in designing an effective Web-based collaborative spaces. These include establishing a system goal, implementing an effective client reputation system, and deciding on a practical vote timing scheme.

Thesis Supervisor: Judith Donath
Title: Assistant Professor of Media Arts and Sciences

Acknowledgments

Thanks to the following people:

Judith Donath, for her guidance, supervision, straight good advice, and patience.

Kelly Dobson, for coming up with the plans for the Tele-Reporter hardware.

Matt Lee, Dana Spiegel, and the Tele-Actors Guild for starting the Tele-Actor project and laying the foundation for my work.

Karrie Karahalios, for hiring me as a UROP two years ago, and for being a bottomless source of knowledge and common sense.

The rest of the Sociable Media Group, for all their help and for being such sunny people.

Patrick Winston, for telling me how to write a thesis.

And my family, for all the right and obvious reasons.

Contents

1 Introduction

The Internet is often viewed as a connector between people in a one-to-one conversation, as is the case in email correspondence. However, when a group of people comes together to interact in a single online location, as in a message board, the Internet can take on definite shape. The information transmitted through the interactions of the group creates a virtual shared space that people can congregate and communicate in, then leave and return to over time.

Like a physical space, an online group space can develop its own timbre and mood and atmosphere. Changing the design, the rules, the layout, or the form of this space can profoundly affect the nature of the communication within its walls. We can see this clearly in real life: a room filled with couches will inspire more relaxed, casual communication than a lecture hall with stadium seating and a podium at the front.

A large amount of current research is dedicated to examining this behavior, and designing new and effective online group spaces. In this paper I will focus on an important subset of this research: designing virtual spaces to maximize collaboration.

Why is this important? Quite obviously, in the real world, groups of people come together to do more than just talk to each other. They combine forces to work together, the power of the whole greater than the sum of the parts. As a group, they make decisions and then effect action based on those decisions. Collaboration is a very powerful and very common facet of group behavior.

Yet, online spaces that are truly collaborative are few and far between. Often it is just too impractical because the participants do not share a physical space and are limited in the amount they can accomplish while still operating within the space.

In a collaborative situation where participants are geographically distributed, teleoperation, the remote control of an agent object, can be a useful concept. An online crowd can be given teleoperative control over an object. Then when the group comes to a decision, it can immediately put the results into action by sending remote directives to the agent object, and continue. The agent object can also be a human, in which case I will use the term teledirection. Teledirecting a human agent allows more flexibility in what can be accomplished remotely, and also introduces some new factors into the social interaction of the space, since a human agent, unlike a robot, is able to respond to its board of directors.

The Tele-Reporter is a collaborative teledirecting system that is an extension of the Tele-Actor project begun in 2001. Users of the Tele-Reporter system view a live broadcast from a reporter on location and participate in the report by issuing directions to the reporter as to where to go and who to talk to, and what to say. Whereas the Tele-Actor was useful primarily as an entertaining performance piece, the Tele-Reporter attempts to direct user collaboration toward more serious or practical purposes. The system interface and design takes into account many factors in an attempt to create a truly productive, collaborative online social space.

In this thesis I will discuss the background and work leading up to the Tele-Reporter project, describe the operation of the system, and discuss the reasoning behind the various design decisions that were made in its construction. I will also discuss the implementation details of and problems faced building the system and reflect on possible future directions for the project. I will close with a comment on the importance of the system as a model for designing collaborative teleoperative systems.

2 Previous Work

2.1 Mercury Project

One of the earliest installations in Web-based teleoperation, the Mercury Project allowed users to participate in a "tele-excavation" [01]. The target was the Mercury Site, actually a sandbox installed in the basement a laboratory at the University of Southern California and helpfully filled with various artifacts.

From 1994 to 1995, the Mercury Project website allowed users to log in and wait for their turn to take control of a telerobotic arm positioned over the site. Through a simple interface consisting of a schematic of the robot arm and a photo image of the local area of the site , a user could direct the arm to any point in the site, lower it, and finally command the robot to search for artifacts by shooting a blast of compressed air at the sand.

Although teleoperation is far from a new concept, the Mercury Project was one of the first attempts at building a teleoperated system that was accessible from any computer connected to the Internet. In doing so, it added a new dimension to the popular perception of the Web as a completely virtual world where data could be exchanged but nothing tangible could be affected. The Mercury Project let users use the Web to actually alter the physical world.

In designing a Web interface for the Mercury Project, the researchers made a conscious decision to sacrifice "telepresence" – the feeling of actually being present at the remote operation site – to accommodate the limitations of Web browsers of the time and to ensure consistency for all users. The result was a non-immersive, yet simple and intuitive interface. In later chapters, I will detail how designing such an interface is also a major component of the Tele-Reporter project.

2.2 Telegarden

The Telegarden, the direct successor of the Mercury Project, gave Web users the ability to cultivate and plant seeds in a remote garden [02]. It was quite different in concept from the Mercury Project's tele-excavation robot as it was clearly created to allow users more than just the ability to execute a task, but also the opportunity to enter a shared, organic space available to any Web user.

In contrast to the Mercury Project, the Telegarden allowed multiple users to be "in" the garden simultaneously. Garden users controlled the seed-planting robot one at a time: orders were placed in a queue and executed in order of placement. However, users were highly aware of what moves were being made by other users, since each user's movements affected the aesthetic appeal and health of the garden as a whole. The community space was enhanced by the implementation of a "village square" where people in the garden could directly interact with each other. Thus, the Telegarden was a simplified collaborative environment in which users worked together to tend to the garden.

The long-term consequences of planting a seed made the garden a persistent environment which users visited repeatedly to watch the results of their efforts bud and flower. As users collaborated and bonded with each other over time, the Telegarden grew into an evolving, developing social community that was both physical and virtual.

2.3 Tele-Actor

The Tele-Actor project begun as a joint effort between U.C. Berkeley and the MIT Media Laboratory and is still under independent development at both ends. It allowed multiple users to simultaneously share control of a teleoperated actor agent [03]. To date, there have been several implementations of the Tele-Actor, but throughout this thesis, I will reference

Figure 1: Tele-Actor in action.

the system built at MIT and demonstrated at a Media Lab sponsor function in the spring of 2001.

The goal of the Tele-Actor project was to design an interface that allowed and encouraged collaboration among the audience [04]. In particular, it was necessary to decide on a method of aggregating the input of a large group of users into a single unified output command to send to the actor. Majority vote was deemed the simplest and most direct option – users would make suggestions and other users would indicate their support by placing votes for that suggestion.

It was decided that the actor would be human. This solved many problems that would have been encountered with a robot agent. A human actor, in addition to coming pre-built, was more capable and flexible than a robot in both interpreting complicated directions as well as executing them.

Figure 2: Tele-Actor interface.

Another important part of the project was the actor's costume [Fig. 1], which had the dual function of both communicating the actor's role and purpose to bystanders and also masking the equipment required for it to operate. The actor wore an elaborate headpiece with a camera mounted on her forehead. She also wore two LCD terminals where audience commands were displayed. One terminal was on the wrist, where the actor could easily see it, and another was mounted across the chest of the actor's padded vest, where bystanders could see it. A laptop connected to the camera and LCDs was hidden in a backpack.

The actor's gear was purposely flashy and space-age in design in order to draw attention to the actor and signal that her behavior was part of a performance. The terminal on the chest was instrumental in making strangers aware of the actor's goals and allowing them to decide whether to help or hinder her on her quest, if so desired.

The Tele-Actor client interface [Fig. 2] was designed to maximize telepresence. Toward that end, the screen was dominated by a real-time audio/video stream broadcasting the

actor's point of view. Users in the system were assigned square icons which they could move around the space, as well as on top of the video screen. They could type text messages, which were interpreted as directions for the actor if the user was on top of the video screen, or chat messages if the user was on the black space outside the video screen. When a suggestion was made, other users could indicate their approval by placing a vote – moving their square on top of the circle surrounding the suggestion. Spatial locations were important, as suggestions were intended to be contextual – a user might move over a doorway in the video screen, for instance, and suggest *Go downstairs*.

Voting was divided into rounds, with the winner chosen at the end of each round sent to the actor. If more than 50% of the users in the system voted for one suggestion, the voting round ended immediately and the suggestion was automatically sent to the actor. Otherwise, the system waited for a fixed interval of time, then sent the suggestion with the most votes to the actor at the end of the round.

The evolution of the Tele-Actor project into the Tele-Reporter is discussed at length in the next two chapters.

2.3 Slashdot Message Boards

Tele-Reporter makes extensive use of a client reputation system in order to maintain orderly behavior among its clients. The reputation system in place on the message boards at Slashdot.org is notable for its complexity as well as its success.

Slashdot.org is a technology website with a huge and very Web-savvy audience distributed across the world. The message boards on the site experience extremely high levels of traffic, and so face the enemy of all large public group spaces: disruptive user behavior. In order to keep its discussions civil and relevant, Slashdot has implemented one of the more comprehensive rating systems on the Web to filter and monitor its boards [05]. This system

centers around two concepts: moderation and karma. Moderation reflects the ratings of individual messages, while karma reflects the overall reputation of a poster.

User posts appear with a default rating (-1 to 5) based on the karma of the poster. Rating numbers are accompanied by comments that describe the value of the post, from "informative" to "flamebait." Users with particularly high karma are appointed by Slashdot to act as temporary moderators, and awarded a limited number of moderation points. These points can be used within three days to raise or lower the ratings of messages as the moderators see fit. The moderation up or down of messages, in turn, affects the karma of its poster.

Slashdot readers, when viewing posts, are given the option to select a threshold value. Then they will only be shown the text of messages with rankings above that threshold. This allows readers to skim through long discussions and be assured of catching all of the valuable posts.

The rating system has a powerful effect on posters; generally, it inspires people to attempt to post better or more insightful messages in hopes of gaining karma. The success of the system is solid evidence for the effectiveness of reputation systems, proving that while anonymous troublemaking is always attractive, prestige and power can be even more so.

3 Goals

The Media Lab Tele-Actor system had its first complete trial run with a human actor at the Digital Life sponsor meeting in the spring of 2001. As sponsors sat in the Media Lab atrium eating lunch, the actor made her way around the floor with a camera on her head and a terminal on her chest bearing her mission. Across the floor, several kiosks were set up for users to log in and provide direction to the actor. As with all demonstrations, valuable lessons were learned from the experience and new goals were set for the next stage of the project, the Tele-Reporter.

3.1 Goal-Oriented Collaboration

It was highly amusing, though somewhat disappointing, to observe the directions received by the actor during the demonstration. Within a few minutes of opening, goals had degenerated from innocuous suggestions like *Talk to Brian Smith* into *Stand on a chair and bark like a dog*. Taking full advantage of its power, the audience took distinct pleasure in commanding the actor to perform silly stunts like singing songs to sponsors or switching peoples' plates of food.

Of course, there is nothing inherently wrong with this kind of behavior. However, it does beg the question of what can be expected from an audience that has been given free rein to control the mouth and limbs of a human being. In the Telegarden, members were restricted in their actions by the inherent limitations of the robot. The Tele-Actor, as a human, has a large range of physical movement and, for better or worse, an unlimited vocabulary of trickery.

The chaotic and unstructured behavior witnessed in the Tele-Actor demonstration can be traced to two key elements. First, the lack of an overriding goal state. The actor was not out to grow a garden. She was not interested in finding buried treasure. She existed purely for the purpose of doing whatever she was told. As such it can hardly be considered surprising that in searching for a goal and finding none, users directed the actor to run in circles, make animal noises and pester sponsors in the restroom.

The second element was a lack of a voting economy in the system. Although a basic economy had been planned, none was actually implemented for this demonstration. All users were given equal weight regardless of their actions, and so there was nothing to keep a few disruptive users from swaying the entire behavior of the group. In a nutshell, the lack of an overall goal meant that there was no forward or backward, and the lack of a user judgment system meant that there was no right or wrong.

It was decided that the next phase of the Tele-Actor would have an interface designed with goal-oriented collaboration as a key priority, and would feature both an overriding goal and a voting economy.

3.2 Intuitive Interface

The Tele-Actor interface, while highly immersive, was also somewhat difficult to use effectively. The decision to center the interface around the video screen and to turn the video screen an interactive bulletin board enhanced the feeling of telepresence, but complicated the process of user interaction. In particular, the central concept that messages placed while the user icon was on the video stream were suggestions, while messages placed outside of the video space were only chat comments was eventually impractical.

Icon placement on the video screen, implemented when the system was designed to direct a user playing a video game, became somewhat irrelevant as it was difficult to pass location-based context in a text message to the actor. A command “Go in there!” lost its

meaning without the context of a doorway framing it. As the actor was constantly moving, objects in the video screen also rarely stayed in one place and users were forced to make constant adjustment. Furthermore, the video screen was often obscured by the mass of user icons.

The Tele-Reporter project seeks to resolve some of these issues by relying on a simple, streamlined interface.

3.3 Practical Timing

Timing is always an issue in a system powered by real-time voting decisions. It is difficult to nail down a discrete fixed interval of time that is just long enough to accurately gauge the opinion of the group as a whole. This is complicated by the fact that in the Tele-Actor, the processes of suggesting goals and voting on them happened simultaneously and continuously. This meant that different suggestions were available for voting for different lengths of time, skewing the weight of votes. Finally, any fixed voting time was often inadequate because different missions inevitably took the actor a different and unpredictable amount of time to complete successfully.

Voting rounds in Tele-Actor lasted for five minutes. If in that time a suggestion gained the votes of more than half of the audience, the round ended immediately and the goal was sent to the actor. More often than not this led to a deadlock situation where voters were split among multiple decisions and ended up standing off, waiting for the clock to expire. This situation was helpful neither to the audience nor the actor.

With the Tele-Reporter, timed rounds became impractical for reasons that will be discussed. A new timing system was implemented to better suit the framework of the interaction and collaboration.

4 Overview

Like the Tele-Actor, the Tele-Reporter allows users across the Web to collaboratively provide direction to a human agent in real time. However, the agent is no longer an actor. In the Tele-Reporter, users control, as one might guess, a reporter.

This shift from actor to reporter changes the relationship between audience and agent. In previous collaborative interfaces, the agent was viewed as a tool – an object to be manipulated in service of the audience. The Tele-Reporter system is designed so that the reporter is viewed not as an instrument but as a *representative*.

The reporter can go where ordinary people are not allowed. He can attend exclusive press conferences and meetings, go backstage at concerts and shows, enter locker rooms at ballgames, and bring the audience along for the entire ride. On television, the audience is only able to observe. In the Tele-Reporter, the audience is given the ability to speak and move through the reporter. Given this chance, audience members should be inspired not to toy with their power and send the reporter crashing into walls, but to use the reporter to do what they would want to do in the same situation. It's the viewer's chance to cut through the pleasantries and ask the questions the public really wants to know. Forget talking to the Secretary, go straight for the President! Don't ask the celebrity about his kids, ask about his girlfriend! Don't take "manager's decision" for an answer – why was the first baseman really taken out of the lineup? Once and for all, ask the Lab Director all those questions no one dared to before – it's okay, you're speaking for all of us!

As the reporter makes his way through the room, the audience communicates and collaborates through a very different interface than in the Tele-Actor. The basic dynamics of interaction remain the same. The audience is expected to carry on a running conversation

amongst itself, commenting on the action, trading thoughts, getting to know each other better, and brainstorming for possible suggestions.

When users make up their mind, they can then proceed to make a suggestion, as in Tele-Actor. Other users can choose to vote for a suggestion or come up with a suggestion of their own. Unlike Tele-Actor, however, users are now given the freedom to vote, chat and suggest continuously – they do not have to choose only one state to be in at a time. Audience members are also allowed to make as many suggestions as they want. This allows more facilitatory users to pose several suggestions and allow others to vote on them. Audience members now have multiple votes, which they can spread across various suggestions. They are able to rearrange and redistribute their votes at will, as the changing real-time flow of the broadcast renders what was interesting a few seconds ago completely irrelevant now. These design decisions are meant to encourage increased interactivity and cut down on deadlock situations.

A major new development is that the reporter is now allowed to respond to the audience in a limited fashion. One reason for this is that unlike the Tele-Actor, where anything was fair game, the reporter must abide by some general rules of conduct in order to stay useful. For instance, if the reporter is interviewing the CEO of a large company, it would be significantly damaging and counterproductive for the reporter to jump on the table and bark like a dog. The reporter cannot stop viewers from putting forth such suggestions. The interface is designed to discourage such disruptive behavior and to prevent subversive suggestions from winning the majority vote. However, should such a suggestion end up winning anyway, the reporter now has the option to veto it. This in turn will punish the audience member responsible for posting the suggestion, reducing the likelihood of further disruption.

As mentioned before, the Tele-Reporter features the implementation of a simple reputation system, put into place to discourage that very behavior above. Under this system, audience members' clout within the system would be dependent on the constructiveness of

their prior behavior. In particular, audience members are rewarded for suggestions that the reporter decides to follow, and punished for suggestions that are vetoed.

Finally, a new timing system is implemented to fit the quick and unpredictable pace of journalistic reporting. Because the reporter controls the flow of conversation, it is now up to the reporter to pull up the winning suggestions whenever he is ready for a new prompt from the audience. This adds an element of unpredictability to the voting, as audience members can never be exactly sure when their decisions will be made final. However, it greatly benefits the process of interaction and keeps things flowing smoothly.

5 The Tele-Reporter Client

The Tele-Reporter client-side package is the applet through which users interact with each other and the reporter.

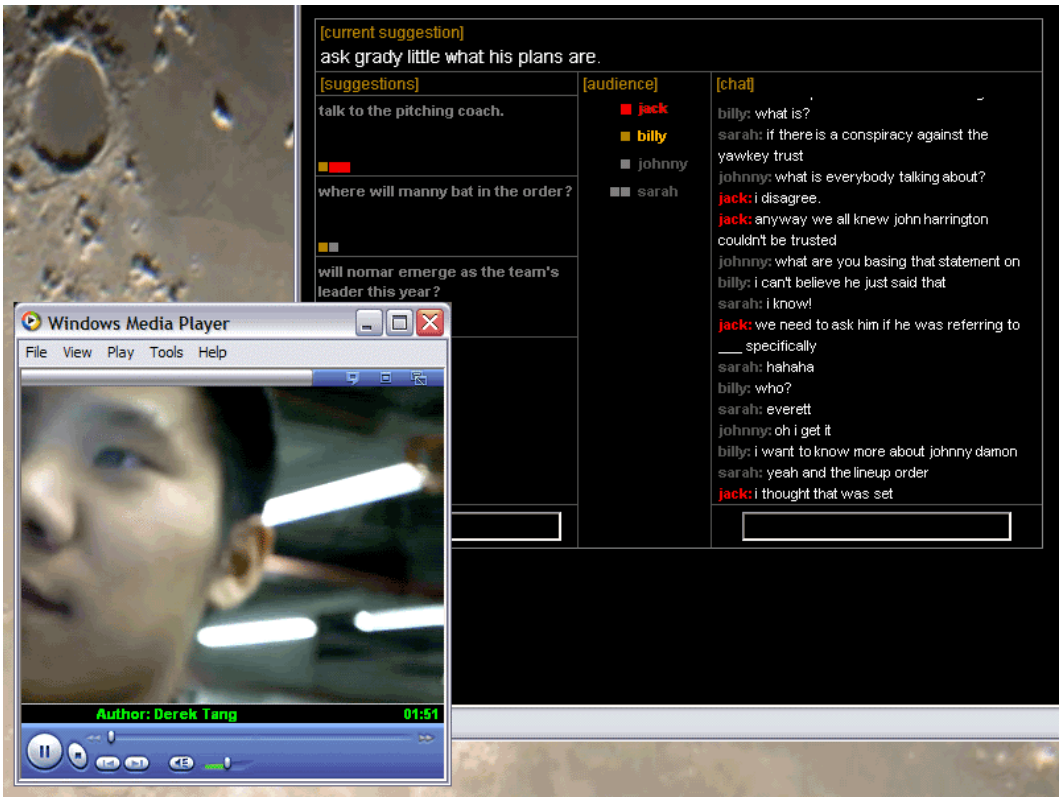

Figure 3. Tele-Reporter client interface (large view).

The main interface [Fig. 3] features several elements. The top panel displays the last winning suggestion that was chosen and accepted by the reporter. The suggestion panel, on the left, displays the current suggestions for the reporter and the votes each suggestion has received. The audience panel, in the center, shows all members of the audience and their voting status. Finally, the chat panel on the right allows users to participate in an ongoing text chat. The video panel is separate from the interface, in its own popup window.

5.1 Watching

Clients watch the reporter's progress in the video panel, which displays a live a/v stream directly from the reporter's camera.

The largest and most obvious difference between the Tele-Reporter and Tele-Actor interfaces is the separation of the video screen and the interface. This changes the audience's perception of its involvement in the actual report significantly. It reduces the sense that the user is actually present and physically participating in the remote action. This decision was made for two reasons.

First, context-based suggestion, based on positioning user icons on top of the video feed, is no longer necessary. As mentioned before, the Tele-Actor was designed to allow users to deliver contextual suggestions such as *Go there* and *Look at that* to the reporter. These context-based suggestions are difficult to transmit accurately to the reporter without replicating the video image on the reporter's end and significantly complicating the reporter interface. Furthermore, the fact that the reporter is constantly moving and the video frame is constantly updating and changing makes static placement of votes at different points on the video unnecessary and unenlightening. Therefore, there is no more reason to superimpose the interface on top of the video.

Furthermore, the separation of a/v feed and interface reflects the changing relationship between the audience and reporter. The reporter, unlike the actor, is not a mindless tool to be controlled robotically. It is important that the audience retain some respect for the reporter's autonomy and his ability to function independently as a reporter. Therefore a sense of telepresence is reduced in order to give the impression that the user is not so much taking control of the broadcast, but watching the broadcast as if it was on television and then providing guidance and suggestions in a secondary window.

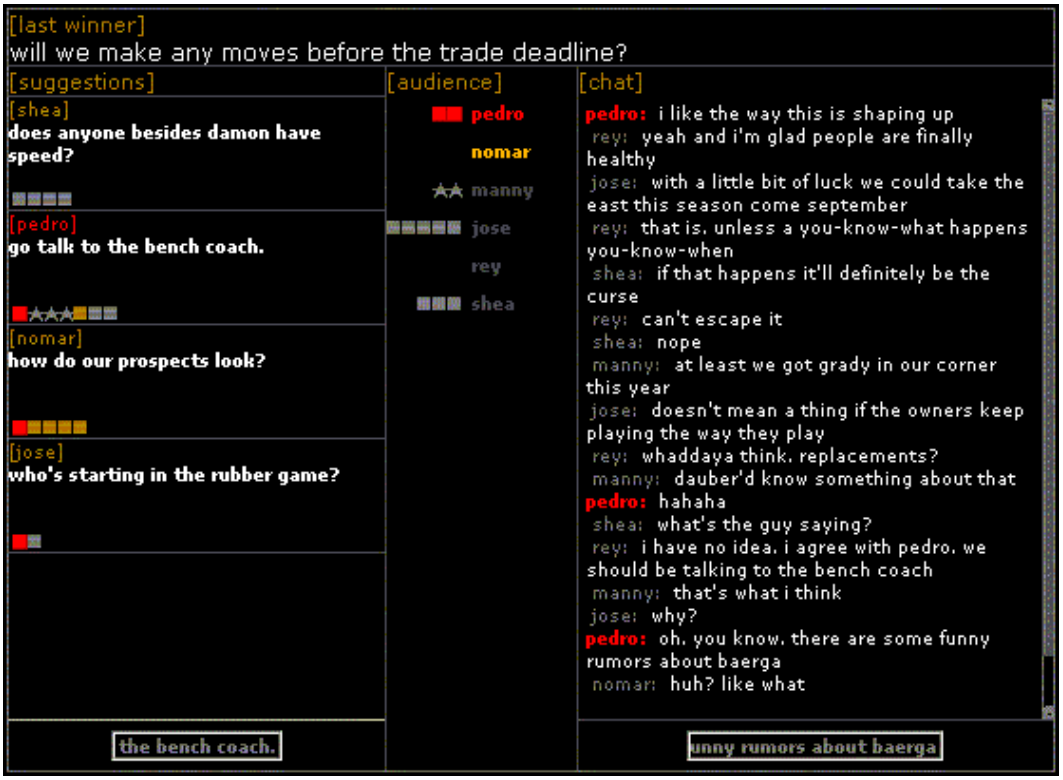

Figure 4. Tele-Reporter client interface (detail).

5.2 Chatting

Through the chat window on the right side of the applet, users can get together and interact in various ways. These include:

- Discussing the events unfolding in the video window. *I can't believe how he's dodging the question!* says one user.
- Tossing ideas back and forth. *I totally agree! I think we need to flat out force him to say if he killed her – yes or no*, says another. In this manner chatting accompanies the posting of suggestions and votes. *I'm going to post that right now*, says one. *You got my vote*, says another.
- Getting to know each other better. *You're smart! Are you from Texas?*

In any of these applications, the presence of chat adds immeasurably to the sense of community within the system.

The choice of a simple standard linear text-based chat window as opposed to the more advanced spatial chat system of Tele-Actor was made in order to allow a new user to ease into an ongoing conversation with as little adjustment as possible. The familiarity of the environment is intended to encourage users to use the chat capabilities casually and to allow the chat to supplement the voting and suggesting process, not supercede it.

5.3 Suggesting

The primary means by which audience members communicate with the reporter is through placing suggestions for the reporter's next course of action. Suggestions are intended to direct the reporter to do one of three things:

- Begin a course of action. When the reporter has finished completing any tasks at hand he will wander the room searching for a new course of action. At this time the audience provides directions suggesting places to go (*Head on into the ballroom*) and people to talk to (*Try looking for the manager*).
- Pursue a course of action. When a subject to interview is found, it is up to the reporter to control the flow of the conversation, and up to the audience to indicate which direction that flow should take. For instance, if a reporter is interviewing a politician and the politician mentions something very significant, say, "I'm absolutely opposed to abortion," the audience indicates its desire to pursue that end by supplying suggestions such as *Ask him if this will affect legislation!* or *Confront him about the contradiction!*
- Change a course of action. If a conversation is becoming bogged down in uninteresting details, it is up to the audience to signal the reporter to move on (*Change the subject*), move on completely *(Who's that standing over there?*) or just give up *(Forget this guy*) .

A user makes a suggestion by typing the suggestion text into the field below the suggestion pool and pressing Enter. This sends a suggestion command to the Tele-Reporter server, which sends the command out to the reporter and all users. The suggestion then appears in the pool with the name of the original poster and the text of the suggestion. Votes accrued are shown at the bottom at the suggestion box.

Under the current implementation, there is no limit on the amount of suggestions that users are allowed to post. This increases the flexibility of the system, as some users are more likely to formulate suggestions and others are more willing to just vote on already-posted suggestions. However, it can also increase the clutter. To make sure that only relevant suggestions stay in the suggestion pool, suggestions are allowed 30 seconds to receive a vote. If after 30 seconds there are no votes for the suggestion, it is automatically removed. Although user may be able to keep their suggestions alive by voting for it themselves, no user will be able to sustain more suggestions than he has votes.

Furthermore, if at any time after 30 seconds after posting the number of votes for a suggestion drops to zero (say, if a newer and better suggestion is posted that causes everyone to move their votes), the suggestion disappears as well, because it has become obsolete.

I should note that it is not the intention of the Tele-Reporter project to turn reporters into robots and render television reporting useless. It is instead to involve the audience in the process. Ideally, suggestions should be specific but allow the reporter room to maneuver within conversations. A suggestion should be more along the lines of *Find out more about his energy policy* rather than *Here's what you say*. This gives the reporter breathing room and allows him to put to use his own journalistic abilities.

5.4 Voting

The audience panel in the middle of the applet serves two purposes. It allows chatting clients to see who else is in the room that they can talk to. It also serves as the platform for voting and viewing the voting status and behavior of other clients.

When a client logs on, he immediately sees all current clients in the system listed in the audience panel, with his name in red at the top. The boxes next to his name reflect the number of votes he has; by default, when a new user logs in, this number is 5. The client can then place a vote by dragging the box into a suggestion in the suggestion panel.

Clients can choose to distribute their votes any way they wish. If a client feels strongly about a suggestion he himself has posted, he is welcome to go ahead and drop all 5 of his votes in it to try to boost the suggestion's popularity. Conversely, if a client sees two or three suggestions he agrees with, he may place a few votes in each one to indicate his support for them. There is no obligation to use up all votes, although statistically speaking the more votes placed, the better the majority vote will reflect the audience opinion.

Clients can also change their minds as long as suggestions are active. As mentioned above, if a client has placed his votes already but cannot make up his mind (or a new and more attractive suggestion is posted), he can transfer his votes back and forth between suggestions as much as he likes by dragging the box out of the suggestion box and into another one. He can also retract a vote by dragging it out of the suggestion pool, in which case it returns to the unused vote pool next to his name. Votes in transition are considered by the system to be unused.

Figure 5. Reporter a) accepts and b) vetoes.

Voting is intended to be neither anonymous nor confidential, as collaboration hinges on client awareness of each other's actions and decisions. To this end, the placement of unused votes in the audience panel is a useful feature that allows users to stay aware of how other users are using their votes and how many they have left to place. Clients can also see which suggestions another client has voted for by rolling their mouse over that client's name in the audience panel. This causes all the client's votes to be highlighted in the suggestion panel. Of course, clients can also chat with each other about their voting decisions and motivations in the chat panel.

5.5 Decisions and Reputation

When the reporter is ready to solicit audience input, he polls the system to return the current suggestion with the most votes. At this point, the system puts a freeze on suggesting and voting. All suggestions disappear except for the winning one, which is highlighted blue. In the

case of a tie, the most recent suggestion is chosen, under the assumption that newer items are more relevant.

At this point users can continue to chat but cannot suggest or vote until the reporter makes a decision regarding the winning suggestion. The reporter can do one of two things; he can either accept the suggestion or veto it.

In the case that the reporter accepts the suggestion [Fig. 4a], a check mark flashes across the suggestion to indicate the decision, and the suggestion text is moved up to the winner's panel. The system then rewards the audience member who posted the suggestion by boosting their internal reputation, as will be explained later.

If the reporter instead vetoes the suggestion [Fig. 4b], the suggestion is crossed out and the winner's panel is blanked out. The system then punishes the user responsible for posting the vetoed suggestion.

In either case, after a decision is made, suggestions are cleared and votes are restored to their users, with vote numbers to reflect their updated reputations. The process of suggesting and voting then begins anew.

Client reputation is implemented in order to establish a flexible social order within the group that maximizes collaboration and balances the relationship between the audience and the reporter. Reputation manifests itself very straightforwardly in the system. A client is assigned a reputation value between −4 and 10.

- For a reputation between 1 and 5, a client's reputation corresponds to the number of votes he is free to place. A client is never able to exceed the default assignment of 5 votes. However, by being punished it is possible to lose votes and thus voting power and reputation within the system.
- For a reputation between −4 and 0, a client has no votes at all. Furthermore, because he has no votes, he is not allowed to make suggestions either (since it was probably a poor suggestion that dropped his reputation in the first place). He is reduced to

chatting and observing others, and must wait for his reputation to rise again before regaining the ability to participate fully in the system.

- For a reputation between 6 and 10, a client is not rewarded with extra votes but is marked as being an outstanding contributor. This is designated by his votes, which are displayed as stars as opposed to boxes. These votes have the same value as any other vote within the system, but signal to other audience members that they are votes made with greater authority and weight. This encourages other audience members to look to these "prestige" clients for guidance in placing votes, and increases the rewarded clients' influence within the group. The hope of becoming a prestige client, as well, inspires other clients to make suggestions they feel will be popular and accepted by the reporter.

In its current implementation, an accept from a reporter raises the reputation of the winning suggestion's poster by 5 points, with 10 as a maximum reputation value. This is to prevent users from gaining permanent prestige status within the group. After each round initiated by the reporter, a prestige client's reputation lowers by one point until it returns to 5. Therefore, in the absence of any further activity, a user with a reputation of 5 who posts a winning suggestion will be granted prestige status for 5 rounds.

A veto, conversely, lowers the reputation of the suggestion's poster by 5 points. The minimum reputation value of –4 is naturally enforced because only voters with at least 1 vote are allowed to post suggestions. After each round initiated by the reporter, a punished client's prestige also rises by one point until it reaches 5.

6 The Field Reporter

The reporter makes up the mobile, remote portion of the Tele-Reporter system. There are three major components to the field reporter package.

6.1 The Clipboard

The clipboard is the means by which the reporter receives instruction from the audience and responds in kind.

The clipboard [Fig. 6a] is small and practical and can actually be used as a real notepad. At the top is a small directional microphone. This is the ideal spot for a microphone as it can pick up the audio of the reporter as well as the targets.

Fastened to the clip at the top of the board is the reporter's suggestion display, a LCD terminal (Matrix Orbital LK 402-13) which is backlit and can display up to 80 characters of text. Two the left of the terminal are two buttons, one black and one red.

The terminal is blank until the reporter decides to solicit direction from the audience, at which time he presses the black button. This informs the server of a poll call and causes the terminal to display the winning suggestion. At this point the suggestion appears flashing on the terminal to indicate that the system is waiting for the reporter's verification. The reporter then evaluates the suggestion. If the suggestion is acceptable, he presses the black button again to accept the suggestion. Otherwise he presses the red button to veto it. The confirmation decision is sent back to the server.

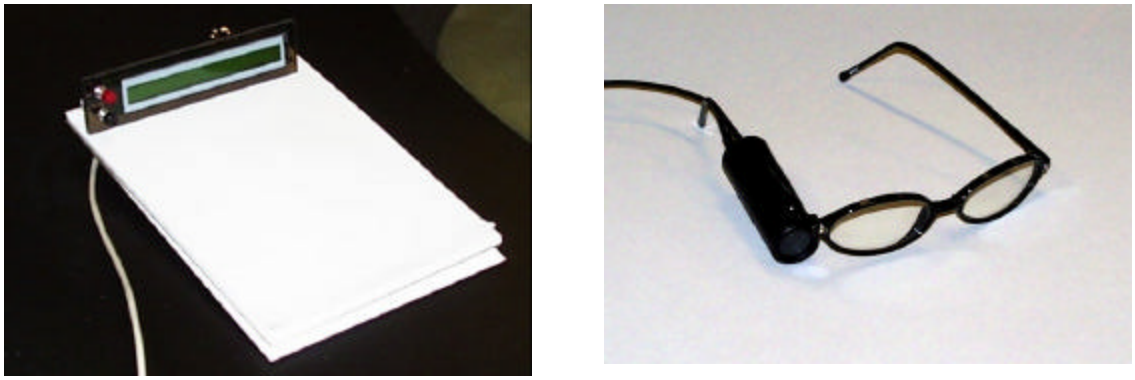

Figure 6. Tele-Reporter a) clipboard and b) camera.

If there are no suggestions that have received votes when the reporter polls the system, he receives a [none available] message, which must also be confirmed. In this particular situation an accept and a veto both have no effect, since there is no winning audience member to reward or punish.

If a reporter chooses to veto a decision, all suggestions are cleared and voting begins again. Therefore the reporter cannot poll again immediately following a veto and expect to get the second best answer. Instead, more than likely he will receive a [none available] instead. The reason for the automatic round clearing is to ensure that the reporter cannot just scroll through suggestions until hitting one he likes. In general, it is important for the reporter to respect the wishes of the audience and for the audience to feel that they are actually participating and directing the report. The reporter is only expected to veto a suggestion if it is absolutely out of the question.

This leads to a discussion of the timing in polling the system. Under the current implementation, the reporter is not aware of the progress of voting among the clients. It is usually not enough for the reporter just to poll the audience whenever he feels like it at random pauses in the conversation. He must anticipate audience reactions. There are three general points when the reporter should poll the system, corresponding to the three basic forms of suggestion behavior:

- When a conversation has ended. This provides the reporter with directions to begin a new course of action.
- When an interesting point has just been raised. This allows the reporter to get a sense of the audience response and to determine whether or not to pursue that point further.
- When a conversation has been going for a long time. This allows the reporter to see if the audience is bored or if he should continue with the interview.

In spite of the guesswork required on the reporter's part, it is still more practical to end voting rounds based on reporter decisions than to automatically forward winning votes to the reporter. This is because no matter what timing scheme is chosen, *someone* must decide when the turning points in a broadcast as defined above are reached, and it had might as well be the reporter. It would be much more difficult to have the audience collaboratively decide exactly when a round can be declared over and a suggestion the winner, especially when suggestions do not reach more than 50% of the votes and there is no obvious consensus.

The choice and design of the clipboard to house the terminal and other assorted pieces of equipment complements the general feel of the Tele-Reporter gear, which is designed to give the impression of a classic film-noir journalist. The idea is to, as with the Tele-Actor gear, allow bystanders to understand the purpose and goal of the reporter just by looking at him. However, in this situation, we do not want to make the gear so overstated and caricatured that people do not the reporter seriously and attempt to disrupt the broadcast.

6.2 The Camera

If the tele-reporter was actually appearing live on television, it would probably be necessary to have an actual cameraman follow him and provide the audio/video feed. However, our implementation is web-only and low-budget and designed for more low-key events (like staff meetings at the Media Lab) where the presence of a cameraman would

probably be intrusive. Therefore, it made more sense to house the camera on the reporter's person.

The camera itself is a relatively small lipstick-shaped bullet camera (Marshall V-2214) under an inch in diameter and 2.4 inches in length. It was not overly difficult to hide the camera on the reporter – the main decision was where to locate it.

It was decided that mounting the camera on the reporter's forehead as in the Tele-Actor project drew too much attention to the camera and gave the reporter too much of a performer's appearance. Instead, the camera is positioned behind the ear of the reporter. This position allows the audience to get a clear shot of the reporter's view, as well as a sense of the presence of the reporter, as the side of the reporter's head is often unobtrusively visible at the side of the video frame. This compensates somewhat for the lack of telepresence in the general interface, as the audience receives the impression that they are looking over the reporter's shoulder instead of actually being the reporter's eyes.

To fasten, the camera is mounted on the arm of a pair of glasses worn by the reporter [Fig. 5b]. The glasses, serious and scholarly, lend added gravity to the reporter's appearance without going over the top.

6.3 The Action Bag

The clunky hardware parts needed to allow the reporter to function correctly are stowed away in a shoulder bag carried by the reporter. The parts include:

- a 12V battery pack to power the camera.
- a Belkin USB camera frame-grabber to allow the computer to read the video input.
- a serial connector connecting the computer to the LCD terminal.
- the cord for the microphone.
- a Sony VAIO laptop running Win2K with wireless Internet connection, interfacing the various parts of the reporter gear and connecting to the Tele-Reporter server.

The laptop runs two programs:

- Sorenson Broadcaster, which captures video data from the camera and sends it to a streaming server.
- A Tele-Reporter reporter-side Java application, which accepts suggestion and voting data from the server, sends data via the serial port to the LCD terminal, reads button presses on the terminal and sends it back to the system server.

The bag sends wires running to the terminal, the microphone and the camera on the reporter, thus restricting the movement of the reporter somewhat. In any event, effort was taken to run the wires discreetly so they would not be noticeable and also to allowing enough wire to hang free so that the reporter would have room to maneuver the clipboard position and turn his head freely.

7 Implementation Details

All data passing aside from a/v data in the Tele-Reporter system between computers is handled by a three-part Java application, comprised of the Tele-Reporter client, server and reporter packages.

7.1 Tele-Reporter Server

The Tele-Reporter Server is essentially a large socket listener, waiting for remote users to log in, and then holding data connections to each user. The reporter is considered a special user who needs to log in before other clients.

In addition to setting up connections, the server passes information received from the reporter or any user to all users. These events are packaged in the form of Java objects called Deltas, which are marked with type (SUGGEST, VOTE, LOGIN, etc.) and data tags to indicate their source client and their pertinent information.

The server also keeps an internal state of the entire Tele-Reporter system – all suggestions, users and votes, so that when a new user logs in the middle of a voting round, the server easily passes on the current state of the system.

7.2 Tele-Reporter Client

The Tele-Reporter client system is a Java applet resident on the server computer and accessible via the Web to any computer in the MIT Media Lab. The applet is responsible for establishing a connection with the server, displaying the client interface, sending client events to the server and receiving server messages.

7.3 Tele-Reporter Reporter

The Tele-Reporter reporter application is responsible for maintaining a serial connection with the LCD interface (using the SerialIO Java package) as well as a data connection with the server. The reporter is updated of suggestions and votes, and sends POLL, ACCEPT, and VETO Deltas to the server.

7.4 Darwin Streaming Server

By far the most difficult piece to implement in this system was the streaming audio and video. This was because of the stringent requirements on streaming quality necessitated by the system, as well as the limited resources available in the laboratory.

It was necessary to capture video encoded on the reporter computer and multicast it reliably to a large number of remote users over the Media Lab local area network. One of the key concerns was buffering latency. Most live streaming video solutions introduce at least 10 seconds of delay between video capture and remote streaming to ensure reliability of the feed. However, this is infeasible in a system where meaningful communication between the reporter and the audience hinges on the audience being aware of what the reporter is doing at exactly that moment in time. It is impossible to make relevant suggestions to the reporter when the users are stuck 10 seconds behind the conversation.

Another problem was simply that there was not enough bandwidth for many users to stream video simultaneously over a single or few connections from an isolated video server. This led to choppy and poor-quality video feeds on the client end.

These problems were not solved in implementing the Tele-Reporter, although many attempts were made to determine a better solution. Instead, a best possible attempt was made. This involved unicasting the video feed using Sorenson Broadcaster from the

reporter's computer to a streaming server resident in the lab running Darwin Streaming Server. The video was then streamed multicast in QuickTime format to the client computers. Although video was captured at 15.0 frames per second, the output stream on the client computer usually dropped down to a frame per second or worse. Buffering delay was able to be reduced to about 4 seconds, but even this felt significant, and furthermore this reduction caused audio and video streams to sometimes come out of sync.

8 Future Research

8.1 Interface Elements

As the interface is a very important part of the Tele-Reporter, future work on the project should focus on streamlining the communication between the reporter and the audience. In particular the round-ending timing process that causes the reporter to receiving a suggestion must be refined. In Tele-Actor, suggestions were sent at fixed intervals, which proved to be inadequate. Sending only when a majority decision was reached was also insufficient because the required level of consensus was rarely reached. In the eventual implementation of Tele-Reporter, the reporter decides when to end rounds and poll the system, a decision which worked well but is still not perfect as the reporter is not always aware when the audience has actually made a decision.

One possible addition to the interface might be an instant interest meter. In addition to voting, clients could set a value indicating their current interest level – above average or below average – with the broadcast. The aggregate interest value of the audience would always be visible to the reporter. Sudden changes in interest value would signal the reporter that the audience was stirring and would allow him to get a sense of when would be a good time to poll the audience.

Eventually, a more advanced reputation system must be built as well. However, this step requires large user group testing of the system to determine the effectiveness of the current reputation system and where its shortcomings lie.

8.2 Applications

Because of the difficulties in streaming video, the Tele-Reporter has only received one small-scale test run, at a Media Lab sponsor function in the spring of 2002. For this demonstration, a reporter walked around the 4th floor of the laboratory, looking at projects and talking to other students. Although the interface was robust, interaction was, as expected, severely hampered by the poor video quality.

In the future I hope that the Tele-Reporter system can be tested on a larger scale. In order to really get a sense of how the audience interacts with the reporter, the Tele-Reporter system needs to be expanded to the point where it can be made publicly available to all users on the web. Public user testing is really the only way to test a system like this, and would provide valuable information as to how effective the system is in encouraging collaboration, and how it needs to be revised.

Before the system reaches this point, of course a much better streaming video solution is required. One possible solution would be to circumvent web streaming altogether, and work in conjunction with an actual local television broadcast. Then the system would be able to use a professional news crew to broadcast the event, perhaps eliminating the need to stream it directly through the Web. However, doing so might affect the balance of the interface significantly.

9 Conclusions

Collaborative teledirection on the Web has great potential for success and usefulness in the future. A successful system allows distributed users to collaborate efficiently and effectively, and alter the physical world via a shared social space.

The Tele-Reporter is a unique and valuable example for studying the design of collaborative teleoperative interfaces, as it reveals the elements necessary to make such a system work: an overriding goal state, a voter economy, and a practical and appropriate timing solution.

The implementation of the Tele-Reporter interface was highly limited by technological difficulties in streaming media. However, I believe that the project merits further development. Large-scale public testing it in various forums will reveal more about the best way to build collaborative teledirection systems and the most appropriate applications for such systems.

References

[01] Ken Goldberg, Michael Mascha, Steven Gentner, 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.

[02] Tele-Garden Info. <http://telegarden.aec.at/cgi-bin/knapsack/html/info.html>

[03] The Tele-Actors Guild. Tele-direction: a new framework for collaborative telepresence. In *Proceedings of CHI 2001,* Seattle, WA, April 2001.

[04] Judith Donath, Dana Spiegel, Matt Lee, Kelly Dobson and Ken Goldberg. Collaborative tele-directing. Unpublished.

[05] Slashdot FAQ. <http://slashdot.org/faq/com-mod.shtml>