Improving Interpretation of Remote Gestures with Telepointer Traces

Carl Gutwin and Reagan Penner
Department of Computer Science, University of Saskatchewan
57 Campus Drive, Saskatoon, SK, S7N 5A9, Canada
carl.gutwin@usask.ca, rpenner@cs.usask.ca

ABSTRACT

Gestural communication is an important part of shared work, both in face-to-face settings and distributed environments. However, gestures in groupware are often difficult to see and interpret because of disruptions to their motion caused by network jitter. One way to improve the visibility of remote gestures is by using traces—visualizations of the last few moments' of a remote pointer's motion. We carried out an experiment to test the effectiveness of traces in helping people interpret gestures. We found that telepointer traces dramatically improved people's accuracy and confidence in their decisions as jitter delays grew larger. Our results suggest that telepointer traces and other visualizations of interaction history can be used to enrich communication among remote collaborators.

Keywords

Gesture, consequential communication, network delay, jitter, telepointer traces, groupware usability

INTRODUCTION

Gesture is a common and essential part of communication in shared spaces (e.g. [1,22]). Explicit gestures help people convey things that are difficult to put into words, such as sizes or locations, and gestural motion that occurs as a natural consequence of activity allows people to gather awareness about each others' actions in a workspace. For example, people may explicitly indicate a path through a space by 'drawing it' over the work surface, or may implicitly show their activity through characteristic actions such as erasing pencil marks from paper.

When collaboration happens in shared-workspace groupware, gesture remains a valuable communication mechanism. Gestures in groupware systems are often conveyed using telepointers, and even though telepointers have limited degrees of freedom, people are still able to use them in powerful ways. However, telepointer gestures are often difficult to see and difficult to interpret because of network jitter. Jitter is variance in the delay of a stream of messages (such as the stream of telepointer positions that determines a gesture). Jitter alters the pacing of the stream, and with telepointer movement the result for the viewer is halting and jerky motion that is difficult to follow.

Although network jitter can be reduced through various network techniques, the problem cannot be removed entirely. An alternate and complementary way to address the problem of jitter is to use *traces* to enhance the visual representation of the motion. Traces are visualizations of the past motion of a workspace embodiment [8], and with telepointers, they can be displayed as a fading line that connects each telepointer position (see Figure 4). The telepointer tail makes visible some of the motion information that is lost due to jitter.

We carried out an experiment to determine whether telepointer traces can improve the visibility and interpretability of gestures during jitter conditions. Twenty-four participants watched and identified gestures that were subject to different amounts of jitter delay. Without traces, we found that as delay magnitude increased, people's accuracy and confidence became significantly worse. With telepointer traces, however, neither of these measures dropped substantially. Performance with traces was significantly better than performance without traces, and the difference between the two conditions increased as the jitter amount grew larger.

This study shows that embodiment enhancements such as telepointer traces can dramatically improve support for gestural communication over real-world networks. In our experience with traces in realistic applications, as long as the visual effect is carefully designed, people find the traces easy to understand and not overly distracting. Since telepointer traces can be easily added to groupware systems, designers should consider including them as a means for enriching distributed interaction and improving the usability of distributed tasks such as shared editing and group design, where gestures are particularly common.

In this paper we first outline three foundational areas: the use of gesture in collaboration, the problems that are caused by network jitter, and the idea of traces as a way to deal with these problems. We then report on the experiment, outline reasons why telepointer traces were effective, and discuss the problems and prospects of using traces in both shared workspace groupware and other kinds of collaborative environments.

GESTURES

In face-to-face work, gestures are frequent and are used for a variety of purposes. For example, in face-to-face design tasks, Tang [22] reports that 35% of hand actions were communicative gestures, and Bekker and colleagues [8] recorded an average of nine gestures per minute in groups of three people. In general, information may be communicated through gestures in two ways, either explicitly or as a consequence of activity.

Explicit gestures are intentional actions designed to convey a particular message to another person. There are several types of explicit gesture: for example, *pointing* to indicate objects, areas, and directions [24], *drawing* to show paths, shapes, or abstract figures [1,22], *describing* to show orientations, distances, or sizes [1], or *demonstrating* to act out the use or operation of an artifact [21,22]. There are other more specialized types as well, such as the emblem, where a gesture stands for a particular word or phrase (e.g. thumbs-up for "OK") [21].

The second type of gesture is that which communicates implicitly and unintentionally – where others pick up information simply from watching another person's movements and actions. This is consequential communication [20], so called because the information is communicated as a consequence of activity rather than as an intentional act. These consequential gestures are also important for the smooth operation and coordination of a group, since they provide valuable awareness information about others' actions and activities [15,20].

Many kinds of activity have characteristic and recognizable motions that, although not intended to convey information, can easily be seen and interpreted by another person. For example, the back-and-forth motion of erasing a drawing with a pencil eraser can be understood even from a distance. In some computer applications, these characteristic motions have even become formalized as commands in gestural interfaces, such as editing gestures in PenPoint [4] and control gestures in marking menus [12].

Gestures in groupware

In groupware systems, gestures are conveyed through the motion of a visible *embodiment* – a visible representation that stands for a person in the workspace. Embodiments vary widely, from simple pointers [7] to video images [18,23] to fully-rendered human avatars in collaborative virtual environments (CVEs) [2]. Although current embodiments are often limited by the capabilities of standard input devices, future systems will allow gesturing that is more like what is possible in real-world situations.

In this paper, we will concentrate on telepointers – cursors that track the location and movement of each person's mouse pointer. Telepointers are small and unobtrusive, but they can convey considerable information. Even though they are much less realistic an embodiment than avatars in CVEs, telepointers often allow a wider range of expression and communication, because they concentrate all of the input capability of the pointing device into one place.

However, one of the problems with telepointer gestures is that they can be difficult to see – a situation that is made considerably worse when network delays disrupt the stream of telepointer messages. In the next section we discuss the ways in which one kind of delay called jitter affects the visibility and interpretability of gestures.

JITTER

There are two main types of delay in groupware systems: latency and jitter. Latency is the lag time between the occurrence of an event on a local machine (e.g. movement of the mouse) and display of that event on a remote machine (e.g. movement of the telepointer). Latency has been shown to cause problems in a variety of group interactions (e.g. [9,16,25]). However, simple latency is less of a problem for the interpretation of gestures (if all movement messages are delayed by the same amount) than it is for synchronization of the gesture with other streams such as verbal conversation.

The second type of delay – jitter – is much more problematic for gesture interpretation. Jitter is variance in the latency of a stream of messages, causing some messages to arrive too far apart, others too close together (see Figure 1). Jitter is only an issue for streams of information that have a meaningful spacing to begin with, such as voice data or cursor movement. Jitter occurs because network traffic changes from moment to moment, causing variable slowdowns, loss, and bottlenecks. To a viewer, the characteristic effect of jitter is halting, jerky motion. This effect can be broken down as follows:

1. The sender produces a regular stream of messages (e.g. pointer positions every 20 ms);
2. Something in the network (e.g. a traffic surge) causes a delay in the stream of position messages;
3. As a result of the delay, several messages 'pile up' and arrive at the same time at the receiver;
4. Processing this group of messages results in several coincident requests to draw the telepointer (assuming that no application-level buffering is done);
5. The display system draws each position, but the screen is refreshed with the next position too quickly for motion to be discerned by the user.
6. The viewer sees the telepointer jump from its initial position, where it has been frozen since the delay began, to the final position after the delay, without seeing any of the intermediate positions.

Jitter has primarily been studied in streaming media, where is has been shown that people notice even very small variations in playback of sound and video files [19]. Jitter has also been shown to cause problems for certain types of group interaction in groupware [9]. For example, jitter delays of more than about 300ms make it more difficult for people to predict where another person's telepointer is moving, or when it has stopped moving. In the next section, we propose how augmenting the representation of a telepointer with a trail or trace can help to smooth the halting motion of a telepointer in a jittery network.

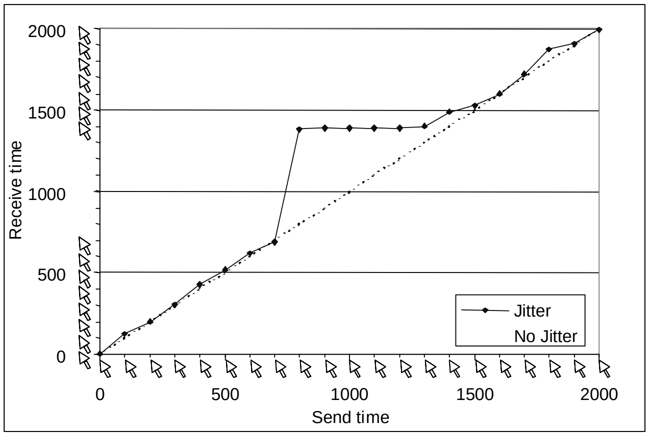

Figure 1. Example of pileup caused by network bottleneck in a stream of UDP messages sent at 100ms intervals. Messages 1-8 are all basically on time and are spaced appropriately. Messages 9-14 are all delayed by a network bottleneck, and arrive at the receiver at the same time. The line of cursors at bottom shows the original spacing of pointer positions; at left shows which positions are drawn.

TELEPOINTER TRACES

Hill and colleagues [13] introduced the idea of interaction histories for computational artifacts, where the system "records on computational objects…the events that comprise their use…and displays useful graphical abstractions of the accrued histories as part of the objects themselves." ([13], p. 3). Researchers have used interaction histories in several subsequent projects but the idea has not been applied to embodiments (although Dourish and Bellotti have suggested the notion of "slime trails" [6]). Our goal is to make use of interaction histories to make gestures and consequential communication more visible and more understandable.

Traces are interaction histories for embodiments, and visualize a person's recent movement in the shared workspace [8]. The inspiration for the visualizations comes from cartoon and comic art [14]. Artists working in these media have long had to address the problem of showing movement convincingly and comprehensibly, even with low frame rates and static pictures. Comic illustrators built on experiments by artists such as Duchamp and Marey, and developed three distinct techniques for emphasizing movement: motion lines, motion blur, and stutter blur (see Figure 2). Motion lines are the simplest, with one or more lines tracing the path of the moving object; motion blur adds the optical effect of streaks along the path; and stutter blur shows several representations of the object on the path.

Cartoonists regularly use these techniques to make movement seem smooth and understandable—not always to show the path of a moving object, but more to emphasize certain movements and to make objects and characters more convincing and real to the viewer. As Chang and Ungar state, the techniques work extremely well, allowing even impossible motions and events to be easily understood [5]. In computer interfaces, these techniques have been used to a limited extent: for example, in cursor trails for passive-matrix LCD screens that help the user find and track the moving cursor, and in some games to show special types of movement (e.g. Mortal Kombat 2).

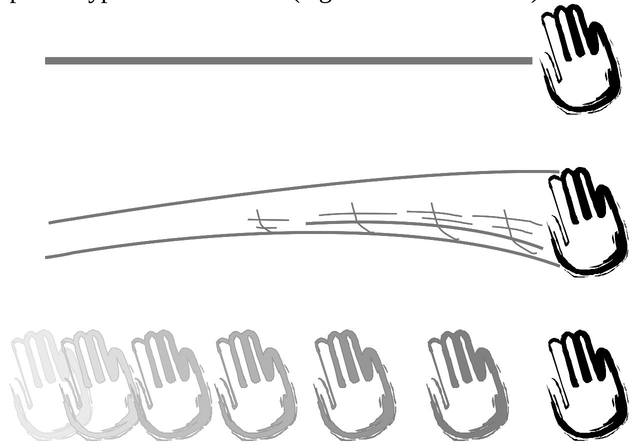

Figure 2. Motion lines, motion blur, and stutter blur in comic strip art (adapted from [14]).

These techniques for depicting and emphasizing motion appear to be a useful way of visualizing embodiment interaction histories, augmenting the basic representation to better convey motion-based information. We applied these ideas to telepointers, and compared several types of traces that vary the technique, length, contrast, and fade effect [8]. We determined that relatively short, low-contrast, fading motion lines showed motion well but did not add undue clutter to the display. Examples of telepointer traces are shown in Figures 3 and 4.

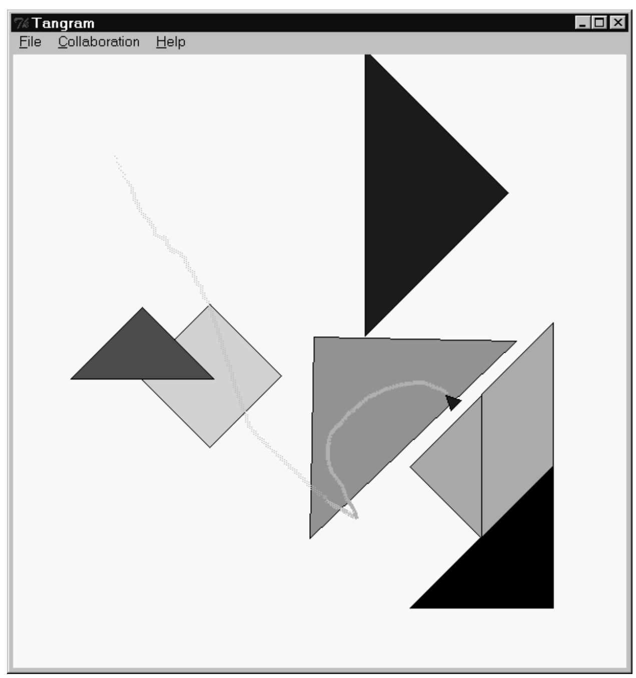

Figure 3. A telepointer trace in a puzzle game, where the remote user has just finished rotating a piece.

Our early experience with telepointer traces suggested that they could be effective in smoothing out gestures in jittery environments. To test this hypothesis, we carried out an experiment that examined people's ability to interpret three kinds of gestures with different levels of jitter, both with and without traces.

EXPERIMENT METHODS

Participants

Twenty-four participants (16 male, 8 female) were recruited from an upper-year HCI class at a local university. All participants were frequent users of mouse-and-windows based systems (at least 20 hours per week). Although all of the participants were familiar with networked applications such as web browsers and email, only about half had experience with real-time groupware in the form of multiplayer games. These participants were familiar with the problems of network delay, although only with latency and not with jitter.

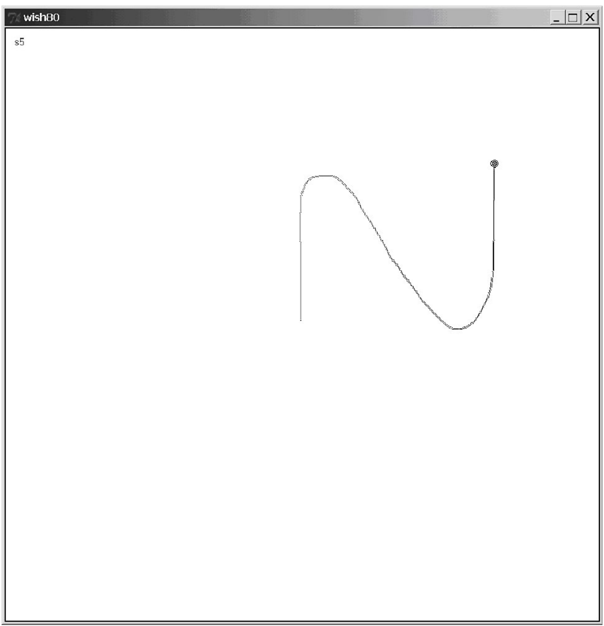

Figure 4. Telepointer trace as used in the gesture study (trace is lengthened in this figure for visibility; see [11] for a more accurate video representation of the effect).

Gesture types and tasks

From the types of gestures that have been seen in collaborative situations (described earlier), we decided to focus on drawing gestures where people trace lines and figures over the work surface. This type of gesture is common in many workspace situations, and incorporates many of the explicit and consequential communication events of shared workspace collaboration. We did not specifically test describing or demonstrating gestures because these make more use of hand and body orientation, two-handed input, and complex transitions such as twisting and rotating – all of which are difficult to compose with a single mouse cursor, implying that these gesture types are currently less likely to be common in groupware. For the study, we tested three types of drawing gesture:

- *Shapes* involve the tracing of a shape or symbol. To test the interpretability of shape gestures, we prerecorded a set of 33 shapes (see Figure 5) that included letters, numbers, and simple strokes. The shapes were all approximately 10 cm in height when drawn on the screen. The participant's task was to determine which of the 33 gestures was being shown on the screen; participants chose their answers from an answer sheet that showed all the gestures (similar to Figure 5).
- *Routes* indicate a route through a set of objects in the workspace. We recorded 35 different route gestures drawn through a workspace of simple numbered circles (see example in Figure 6). The participant's task was to write down the object numbers, in order, that were indicated by the gesture.
- *Areas* are gestures that outline a particular region of a workspace or a group of artifacts. To test area gestures, we recorded 35 instances of a person drawing a line to enclose several objects in a workspace (example shown in Figure 6). The participant's task was to write down the object numbers that were enclosed by the area gesture (in any order).

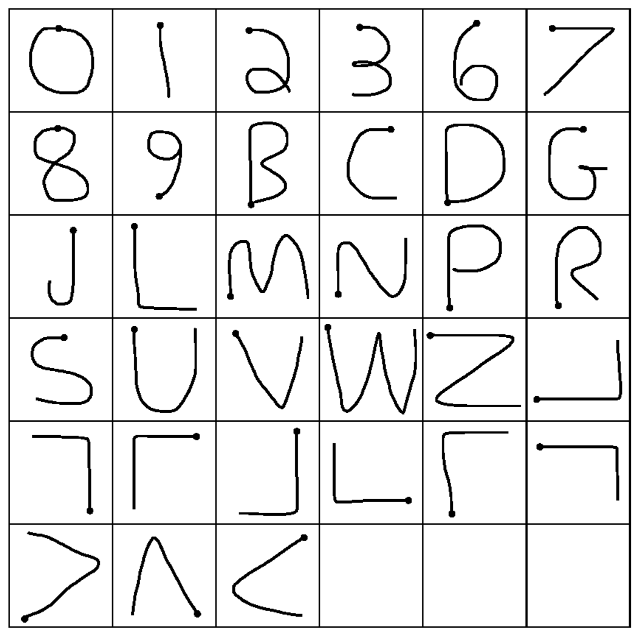

Figure 5. Shape gestures used in the study. The black dot indicates where the gesture begins.

Experiment apparatus

The experiment was conducted on a P3 Windows 2000 PC running a custom-built Tcl/Tk application. The study system was shown in an 800x800-pixel window in the center of a 1280x1024 21-inch monitor.

All gestures were prerecorded and stored as lists of time-stamped cursor positions, allowing us to control the speed and pacing of the gesture's motion during replay. Jitter delays of 200, 400, or 600 milliseconds were randomly imposed on 10% of the messages in the playback stream. These conditions mimic patterns observed in logs of real groupware use on the Internet [9,10].

Traces were implemented as line segments that connected cursor positions. The lines were 1 pixel wide, and faded completely over a period of 800 milliseconds. Although trace length is dependent on cursor velocity, the moderate cursor speed used to record the gestures meant that the traces were approximately 9 cm long, and showed about one-third or less of the gestures at any one time.

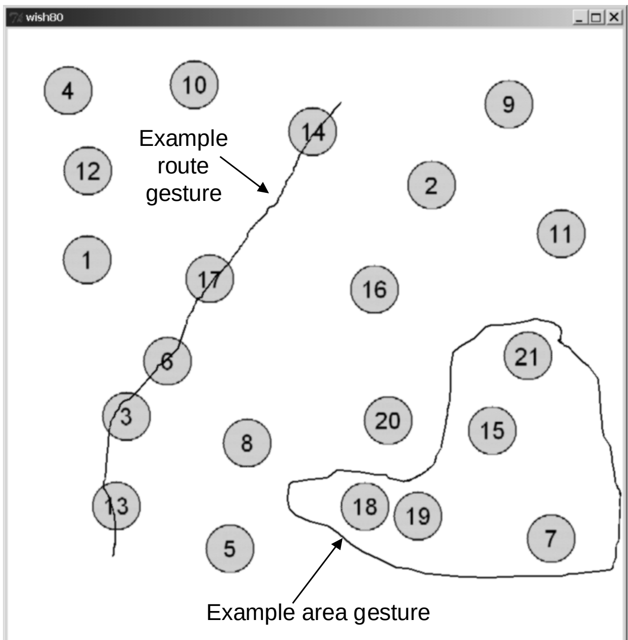

Figure 6. Map for route and area gestures, showing examples of route and area gestures (note: examples show full cursor paths, not traces; also, multiple gestures did not appear at the same time in the actual study).

Design

The study used a 3x4x2 within-participants factorial design. The factors were:

- Gesture type: Shapes, Routes, or Areas
- Jitter magnitude: 0 ms, 200 ms, 400 ms, or 600 ms
- Trace: On or Off.

With 24 participants and eight gestures per block, there was a total of 4608 gestures in the experiment. Data collection included decisions about what each gesture showed, as marked by the participant on an answer sheet, and confidence ratings on a five-point scale for each decision, also self reported on the answer sheet.

We tested two main hypotheses, one relating to the effects of jitter magnitude, and one relating to the effects of telepointer traces:

1. Participants will be less accurate and less confident in their decisions as jitter magnitude increases;
2. Participants will be more accurate and more confident when viewing gestures with telepointer traces than when viewing gestures without traces.

Procedure

Each participant completed 24 blocks of trials (2 trace conditions x 3 gesture types x 4 jitter levels). Order of the blocks was determined by a Latin square design, and each participant followed a different order. In each block, the participant completed eight trials. A trial consisted of four steps. First, the participant pressed the space bar to start the trial, and the gesture was played for them on the screen. Second, they marked their decision about the gesture on an answer sheet. Third, they marked their level of confidence using a five-point scale. Fourth, they pressed the space bar again to indicate that they were finished the trial.

RESULTS

Analysis followed the hypotheses stated above. The data was first analysed to consider the effects of jitter on accuracy and confidence. A second analysis was then carried out to determine whether the presence of a telepointer trace made a difference in these variables. A summary of the results for accuracy and confidence measures are shown in Table 1.

Effects of increasing jitter

In the trials without telepointer traces, it was clear that as jitter magnitude increased, people had more difficulty interpreting the gestures. For example, interpretation accuracy dropped from around 90% to around 50% for all gesture types as jitter delays increased from 0 ms to 600 ms (see Figure 7). Confidence in interpretations dropped in a similar fashion, and response time also increased for all types (although more for shapes than routes or areas).

Data from trials without the telepointer trace was analysed using ANOVA. Considering all gesture types together, there were clear main effects of jitter magnitude for both primary dependent variables:

- Accuracy: $F(3,69) = 79.38$, $p<0.001$
- Confidence: $F(3,69) = 219.57$, $p<0.001$

The results indicate that jitter delays of about one-third of a second or more cause serious problem for gesture interpretation. Figures 7 and 8 shows mean results for all dependent variables and all gesture types.

We also tested whether jitter had an effect on trials where traces were used. ANOVA showed main effects of jitter amount for confidence, but not for accuracy:

- Accuracy: $F(3,69) = 1.87$, $p = 0.143$
- Confidence: $F(3,69) = 6.44$, $p<0.01$

The lack of a main effect of jitter on accuracy suggests that traces make gesture interpretation much more resilient to the problems caused by this type of network delay. In addition, although there was an effect on confidence, the actual reduction was relatively small.

Comparing performance with and without traces

As can be seen in the following figures, the addition of telepointer traces had a marked effect on accuracy and confidence. When the telepointer had a trace, increasing jitter magnitude made very little difference to either of the two dependent variables. Analysis of our second hypotheses tested whether performance was better with telepointer trails than without. Clear main effects were found using ANOVA for both accuracy and confidence:

- Accuracy: F (1,23) = 148.48, p<0.001
- Confidence: F (1,23) = 264.67, p<0.001

The results show that performance with traces is substantially better than with unaugmented telepointers. Even at relatively low jitter levels (200 ms), having a trace implies about eight per cent fewer errors. This difference grows as jitter increases, and at 600 ms, the difference in accuracy is between 30% and 40%. The increasing difference suggests an interaction, and ANOVA did show significant interaction between Jitter and Trace for both accuracy and confidence (for accuracy, F 3,69 = 68.08, p<0.001; for confidence, F 3,69 = 165.26, p<0.001). This interaction was expected; since the aim of telepointer traces is to reduce the negative effects of jitter, we would expect that the difference would become larger as jitter grows.

There was one other significant interaction, between Gesture Type and Trace for confidence (F 2,46 = 10.83, p<0.001). This interaction indicates that although the beneficial effect of traces did increase for each gesture type, the degree of increase was different for different gestures. Inspection of Figure 8 shows that as jitter increased, traces led to slightly higher confidence ratings for shape and route gestures compared to area gestures.

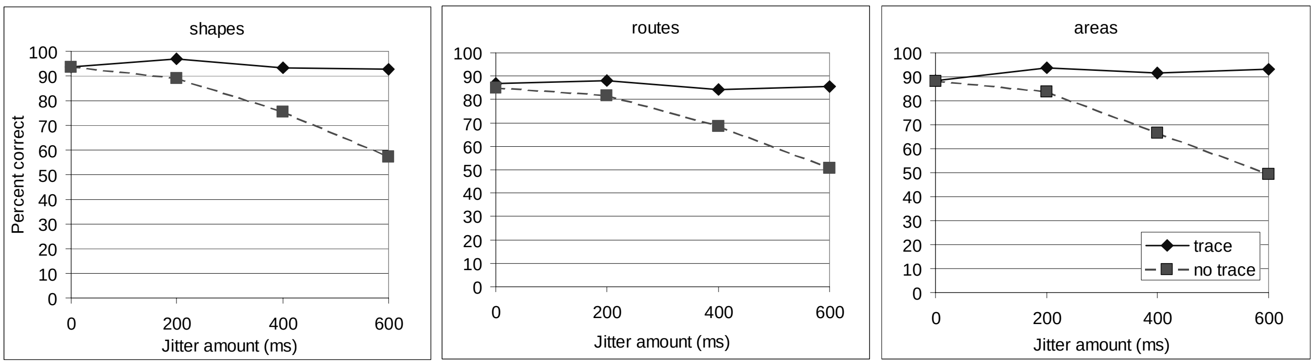

Figure 7. Mean interpretation accuracy for all three gesture types (higher is better)

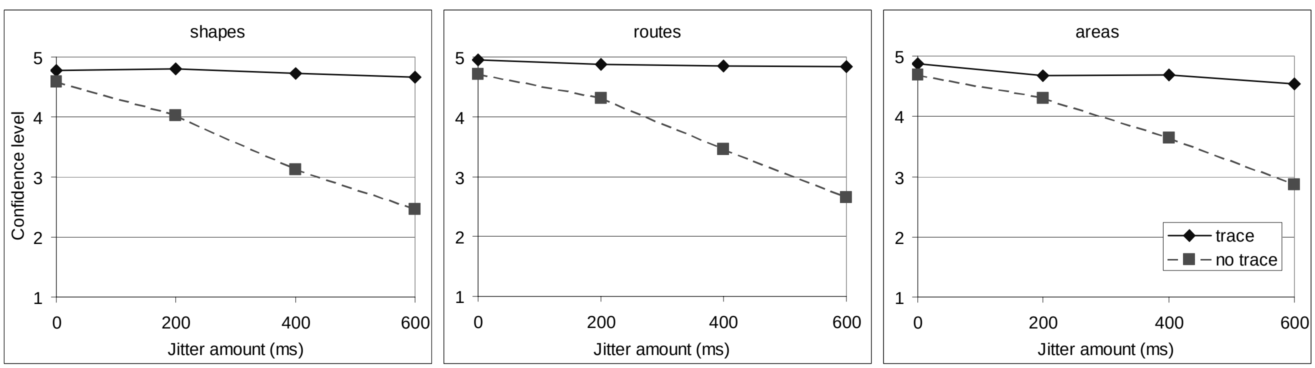

Figure 8. Mean self-reported confidence in interpretation for all three gesture types. (Numbers were anchored on the answer sheet as follows: 1:Just a guess, 2: Not confident 3: In between, 4: Confident, 5: Positive)

		Accuracy (%)				Confidence (1-5)			
		Trace		No Trace		Trace		No Trace	
Task	Jitter	Mean	SD	Mean	SD	mean	SD	mean	SD
Shapes	0 ms	93.75	8.24	93.75	9.03	4.78	0.46	4.59	0.54
	200 ms	96.88	5.53	89.06	9.26	4.80	0.39	4.03	0.71
	400 ms	93.23	7.35	75.52	15.85	4.73	0.39	3.12	0.75
	600 ms	92.71	12.18	57.29	23.86	4.67	0.48	2.46	0.78
Routes	0 ms	86.98	21.64	84.9	20.52	4.95	0.11	4.71	0.51
	200 ms	88.02	21.64	81.77	25.26	4.88	0.21	4.31	0.69
	400 ms	84.38	23.09	68.75	26.06	4.85	0.23	3.46	0.70
	600 ms	85.42	27.5	50.52	25.93	4.83	0.23	2.66	0.76
Areas	0 ms	88.54	9.69	88.02	11.35	4.87	0.29	4.68	0.53
	200 ms	93.75	6.38	83.85	12.49	4.68	0.49	4.31	0.80
	400 ms	91.67	10.21	66.67	12.04	4.68	0.52	3.65	0.77
	600 ms	93.23	9.01	49.48	20.35	4.54	0.54	2.87	0.72

Table 1. Summary of accuracy (%) and confidence (1-5) for all gesture types and all jitter delay levels.

DISCUSSION

The study clearly shows that increasing jitter magnitude has a negative effect on people's ability to interpret gestures, and that adding a trace to the telepointer greatly reduces the negative effects of jitter. In this section we discuss several issues arising from the study: why the traces work, differences between gesture types, whether the benefits will generalize to real-world situations, the problem of clutter, and how traces compare to other delay techniques.

Why telepointer traces work

There are two main reasons why traces are effective at improving gestures: first, traces change the experience of viewing the gesture to one that is more robust in the face of jitter; and second, traces are able to make use of position information that is not used effectively by the system.

When viewing the study gestures with and without traces, there are obvious differences in the way in which the gesture is experienced [11]. Without traces, the gesture is like a film clip or animation; with a trace, the gesture is like a combination of a film (the telepointer) and a drawing (the semi-persistent trace). As jitter increases, the effect on the film metaphor is to reduce the frame rate, making interpretation more and more difficult. The effect of increasing jitter on the drawing, however, is minimal; people are still easily able to 'read' the trace on the screen regardless of the jumpiness of the telepointer.

Second, traces allow the presentation of information that traditional telepointer implementations effectively discard. When several position messages arrive at the same time, traditional implementations imply that none but the most recent will be seen by the user. One way that this extra position information can be used is to change the representation from one that is position-based to one that is path-based – an alternative representation that makes different information visible. In the terms of Hill and colleagues [13], traces are an alternate informational physics for telepointers that make them appear in a form that is more appropriate to the task of seeing and interpreting gestures.

Differences between gesture types

Our observations of the ways that participants worked with the different types of gestures during the study suggest additional reasons why specific types of problems occurred.

In route and area gestures, jitter could cause entire sections of the gesture to be cut off. This likely led to errors of interpretation in which people simply 'connected the dots' of the telepointer positions, since there was no visual information to suggest where the gesture had gone between the points.

With shape gestures, however, corner-cutting from jitter had some other subtleties. When shapes are made up of straight lines and sharp corners, the person drawing the shape must slow down for the corners. This implies that there will be more pointer positions sent from the corners (assuming static frequency of mouse interrupts), increasing the likelihood that the corners will be shown during jitter conditions. Oddly enough, this seems to allow for the connect-the-dots interpretation strategy to work fairly well, making straight-line gestures more resilient to jitter than curved gestures.

Generalizing the results

While the results of the laboratory study strongly show the value of telepointer traces, we also need to consider how well the technique will work in different kinds of real-world groupware applications, whether traces will improve other types of gestures, and whether they can be generalized to other types of embodiments.

Our experiences with three realistic applications in which we have implemented telepointer traces (a puzzle game, a drawing editor, and a file browser) suggest that the improvements to gesturing will translate to real-world applications, as long as participants use gestures in those tasks. In our informal tests with these applications, we considered two issues that were raised during the study: whether gesturing over a backdrop of workspace artifacts would make the traces difficult to see, and whether gesturing in a larger workspace would change the traces' effectiveness. Neither of these issues appeared to be a problem with our applications. Traces were easy to see over a variety of workspace artifacts including text, line drawings, and pictures. Also, gesturing in a larger space was easily accommodated, primarily since people usually called attention to their gesture with a verbal remark (e.g. "like this <gesture>"). However, we did observe several periods in our realistic applications where people simply did not gesture very often. Therefore, the value of telepointer traces will be limited to the amount of gestural communication that is expected in the task.

Types of gesture other than the drawing gestures used in the study are likely to see varying benefit from the addition of traces. An important factor here is whether the gesture's meaning derives more from path information or from other information such as location. For example, pointing is much more dependent on location than it is on path, and so traces will not benefit pointing gestures nearly as much as they did route, area, or shape gestures (although traces do help people track a telepointer to the location where the pointing takes place). In contrast, gesture types such as demonstrations and descriptions seem highly oriented towards path information. Therefore, although these gestures are not often not seen with telepointers because of the lack of expressiveness in the input device, when they are more common in groupware we expect that they will benefit from the addition of traces.

Another issue of generalization is whether traces can be applied to other kinds of embodiments, such as video images [18] or avatars in CVEs [2]. To begin with, the idea of visualizing the recent motion of the embodiment is certainly a possibility – and other types of embodiment will also be affected by jitter. However, since video images and avatars are much larger than telepointers and have more degrees of freedom (e.g. hands and arms as well as fingertips), it is unclear whether the visual effect of the trace would be difficult to calculate, or whether the larger visualization would cause too much screen clutter. It is certain that traces would not be applied in the same way to each part of a larger embodiment. For example, an avatar's torso movements convey different information than its fingertip movements do, and so should have a different informational physics; and since torsos are much bigger, they will require a trace that is more subtle than a fingertip trace in order to prevent clutter.

Distraction, screen clutter, and user control

A concern often mentioned in regards to traces is the amount of screen clutter they add to the workspace, which may occlude objects or distract users from their individual tasks. In our experience this is not the case if the visual effect of traces are carefully designed, and if the visual weight of the trace is controllable by the viewer.

The fading line trace used in this study was chosen after comparing several different types of visual effects [8]. Several of these, including the stutter-blur approach used with laptops, long trails, and trails that do not fade out were seen as distracting, but the more subtle representation used with the fading line was not seen as a problem.

Nevertheless, there will be situations where people will want the traces 'turned down' and this can be easily done, by allowing people to control the length, width, contrast, and transparency of others' traces. We believe that these types of controls could be tied to a focus and nimbus awareness model [3], where traces would become more or less apparent based on the focus of the viewer and the nimbus of the gesturer.

Comparison to other techniques

Traces can be compared to other means of dealing with network jitter, such as buffering and dead reckoning. Traces can also be compared with annotation facilities that appear in some groupware systems.

Client-side buffering of incoming messages is commonly used as the standard solution to network jitter for media such as audio and video streams. In this technique, a certain number of messages are stored at the receiver before 'playback' begins, and then messages are taken from the buffer at the appropriate rate for smooth display. As long as the message buffer holds enough messages to keep playing during jitter delays, this scheme will result in correct pacing of the stream. However, buffering presents a difficult tradeoff where jitter is reduced only by artificially increasing latency. In general, the receiver must delay the stream by the maximum expected jitter amount in order to prevent any halting in playback. This latency cost makes buffering a much less attractive solution for groupware, since the artificial delay will exacerbate the existing network latency, and because lag has definite negative effects on interaction (e.g. [9,25]). Traces have the advantage that they do not require the addition of any artificial delays, so that collaborators can see both the most up-to-date cursor position as well as the recent path information. Nevertheless, there may be situations where either of these techniques, or a combination of them, may be more appropriate for a given task type.

Another technique often used in delay situations is dead reckoning. This is a scheme where the motion of an object such as a telepointer is determined by extrapolation from past locations, when current positions are unavailable due to network delay. This technique has been used successfully to show object movement in CVE systems (e.g. [25]) and even to show avatar movement in some networked games (e.g. Unreal Tournament). However, dead reckoning is only effective when objects tend to move in relatively predictable ways, and gestures are not regular enough to make this technique a viable solution to the problems introduced by jitter.

Finally, the effect of telepointer traces appears similar in some ways to the annotation facilities that have been implemented in some groupware systems (e.g. [17]). In these systems, users can draw overtop the workspace by holding down a particular key or mouse button. Although annotation facilities are similar to traces, there are two important differences. First, traces do not require any effort on the part of the gesturer; annotations must be explicitly initiated by the user, a relatively heavyweight action for something as lightweight as a gesture. In addition, explicit annotations will only benefit explicit gestures, not consequential communication. Second, annotations are much more persistent than a trace, and for frequent and ephemeral actions such as gestures, an impermanent representation seems more appropriate. However, explicit annotations also have advantages (e.g. in making multi-stroke gestures), and it is again likely that annotations and traces can easily coexist in the same application.

Future directions

We plan to further investigate traces and network jitter in three ways. First, we are building adaptive traces that alter the length of the trace to match the amount of jitter currently in the network. For example, the number of telepointer positions included in the trace could be varied based on a regularly-measured jitter magnitude, or even based on the actual arrival times of the messages in the stream. This approach would guarantee that all telepointer information is made visible regardless of jitter, but would use only the minimum trace length required at any time. It remains to be seen, however, whether variable-length traces are appealing and understandable to users.

Second, we are beginning to design traces that can work with different types of embodiments including video images and avatars. As discussed above, this work requires that several types of trace be designed for the different types of information that can be produced, both explicitly and consequentially, by the embodiment.

Finally, we are continuing to add telepointer traces to realistic groupware applications to gain more knowledge about the ways that both network jitter and augmentations like traces affect group work and change the way that collaborators interact in shared spaces.

CONCLUSIONS

Gestural communication in real-time groupware is made difficult by network jitter. One approach to overcoming the problems of jitter is to use traces to make the path of a workspace embodiment visible. We carried out a study to test the effects of jitter on gesture comprehension, and to test whether telepointer traces could reduce those negative effects. Our results showed that traces were remarkably effective at assisting people in interpreting certain types of gestures accurately. Traces can be implemented easily and efficiently, and it seems clear that groupware designers should consider adding them to groupware that will be used over wide area networks. Traces are one example of augmented representations and alternative informational physics that can help improve the richness and usability of real-time groupware systems.

ACKNOWLEDGMENTS

This work was supported by the Natural Sciences and Engineering Research Council of Canada. Thanks to Sheelagh Carpendale, Barry Brown, Jeff Dyck, and David Pinelle for discussions about earlier versions of the paper.

REFERENCES

1. Bekker, M., Olson, J., and Olson, G. Analysis of Gestures in Face-to-Face Design Teams Provides Guidance for How to Use Groupware in Design. *Proc. DIS'95 Symposium on Designing Interactive Systems*, 1995, 157-166.
2. Benford, S., Bowers, J., Lennart, E.F., Greenhalgh, C., Snowdon, D. User embodiment in collaborative virtual environments. *Proc. ACM CHI'95*, 1995, 242 – 249.
3. Benford, S., and Fahlen, L. A Spatial Model of Interaction in Large Virtual Environments. *Proc. Third European Conference on Computer-Supported Cooperative Work*, 1993, 109-124.
4. Carr, R., and Shafer, D., *The Power of PenPoint*. Addison Wesley, New York, 1991.
5. Chang, B., and Ungar, D., Animation: From Cartoons to the User Interface. *Proc. ACM Symposium on User Interface Software and Technology*, 1993, 45-55.
6. Dourish, P., and Bellotti, V. Awareness and Coordination in Shared Workspaces, *Proc. ACM CSCW'92*, 1992, 107-114.
7. Greenberg, S., Gutwin, C., and Roseman, M., Semantic Telepointers for Groupware. *Proc. OzCHI '96 Australian Conference on Computer-Human Interaction*, 1996, 54-61.
8. Gutwin, C. Traces: Visualizing the Immediate Past to Support Group Interaction. *Proc. Graphics Interface 2002*, 2002, 43-50.
9. Gutwin, C. The Effects of Network Delays on Group Work in Real-Time Groupware. *Proc. 7th European Conference on CSCW*, Bonn, 2001, 299-318.
10. Gutwin, C. *Groupware Latency and Jitter on the Internet*. Technical Report, Department of Computer Science, University of Saskatchewan, 2001.
11. Gutwin, C. *Telepointer Traces*. Video clip: hci.usask.ca/publications/2002/gestures-cscw/index.xml, 2002.
12. Kurtenbach, G., and Buxton, W. User Learning and Performance with Marking Menus. *Proc. ACM CHI'94*, 1994, 258-264.
13. Hill, W.C. & Hollan, J.D. Edit wear and read wear. *Proc. ACM CHI'92*, 1992, 3-9.
14. McCloud, S. *Understanding Comics*. Kitchen Sink Press, 1994.
15. Norman, D., *Things That Make Us Smart*, Addison-Wesley, 1993.
16. Park, K. and Kenyon, R. Effects of Network Characteristics on Human Performance in the Collaborative Virtual Environment. *Proc. IEEE Virtual Reality '99*, 1999, 104-111.
17. Roseman, M. and Greenberg, S. GroupKit: A groupware toolkit for building real-time conferencing applications. In *Proc ACM CSCW'92*, 43-50.
18. Roussel, N. Exploring new uses of video with videoSpace. In R. Little and L Nigay, eds., *Proc. 8th IFIP International Conference on Engineering for Human-Computer Interaction*, 2001, 73-90.
19. Scholl, F., *Planning for Multimedia*. White paper, Monarch Information Networks. Currently available at: http://www.monarch-info.com/pmm.html, 2000.
20. Segal, L. Designing Team Workstations: The Choreography of Teamwork, in *Local Applications of the Ecological Approach to Human-Machine Systems*, P. Hancock, J. Flach, J. Caird and K. Vicente ed., Lawrence Erlbaum, 1995, 392-415.
21. Short, J., Williams, E., and Christie, B., Communication Modes and Task Performance, in *Readings in Groupware and Computer Supported Cooperative Work*, R. M. Baecker ed., Morgan Kaufmann, 1976, 169-176.
22. Tang, J. Findings from Observational Studies of Collaborative Work, *International Journal of Man-Machine Studies*, 34(2), 1991, 143-160.
23. Tang, J., and Minneman, S. VideoWhiteboard: Video Shadows to Support Remote Collaboration. *Proc. ACM CHI'91 Conference on Human Factors in Computing Systems*, 1991, 315-322.
24. Tatar, D., G. Foster, and D. Bobrow. Design for Conversation: Lessons from Cognoter, *International Journal of Man-Machine Studies*, 34(2), 1991, 185-210.
25. Vaghi, I., Greenhalgh, C., and Benford, S., Coping with Inconsistency due to Network Delays in Collaborative Virtual Environments. *Proc. ACM VRST Workshop*, 1999, 42-49.