Supporting Collaborative Interpretation in Distributed Groupware

Donald Cox
IBM Software Solutions
Toronto Laboratory
Toronto, Ontario, M3C 1H7 Canada
dcox@ca.ibm.com

Saul Greenberg
Department of Computer Science
University of Calgary
Calgary, Alberta T2N 1N4 Canada
saul@cpsc.ucalgary.ca

ABSTRACT

Collaborative interpretation occurs when a group interprets and transforms a diverse set of information fragments into a coherent set of meaningful descriptions. This activity is characterized by *emergence*, where the participants' shared understanding develops gradually as they interact with each other and the source material. Our goal is to support collaborative interpretation by small, distributed groups. To achieve this, we first observed how face-to-face groups perform collaborative interpretation in a particular work context. We then synthesized design principles from two relevant areas: the key behaviors of people engaged in activities where emergence occurs, and how distributed groups work together over visual surfaces. We built and evaluated a system that supports a specific collaborative interpretation task. This system provides a large workspace and several objects that encourages emergence in interpretation. People manipulate *cards* that contain the raw information fragments. They reduce complexity by placing duplicate cards into *piles*. They suggest groupings as they manipulate the spatial layout of cards and piles. They enrich spatial layouts through *notes, text* and *freehand annotations*. They record their understanding of their final groupings as *reports* containing coherent descriptions.

Keywords

Collaborative interpretation, emergence, meeting support tools, real-time distributed groupware.

INTRODUCTION

Everyday working meetings include many types of activities: sharing information, brainstorming, ongoing design, argumentation, decision-making, and so on. In this paper we are interested in one type of meeting activity, which we call *collaborative interpretation,* and how it can be supported in distributed real-time groupware. We define collaborative interpretation as:

> a process where a group interprets and transforms a diverse set of information fragments into a smaller, coherent set of meaningful descriptions.

Many collaborative activities involve group members, either independently or as a part of a meeting, in collecting and generating information fragments. Depending on the group's goals, these fragments could come from many sources. These include the scattered set of ideas produced in brainstorming sessions [20]; a list of issues arising from early design deliberations [3]; a candidate list of product or system design requirements [1]; a collection of unranked design scenarios [17]; or a set of usability problems produced by independent usability evaluators after inspecting an interface [19, 4].

Information fragments by themselves are of only limited value. To be truly useful, they must be interpreted and recast in light of the goal for the final product. For example, within the context of heuristic evaluation [19, 4], people write usability problem descriptions (the fragments) as phrases and short sentences. Yet others may find these difficult to decipher, or may miss the nuances implied by them [10]. Fragments may be overly specific or low-level—they may indicate particular instances of ideas, issues or problems, rather than a more powerful high-level abstraction with broader coverage. As a set, they may contain irrelevant or duplicate fragments. The set may be unstructured, or may be organized into tentative and perhaps ineffective groupings.

In collaborative interpretation, the group deliberates over these information fragments. The basic goal is sense-making, where fragments are transformed into something that is meaningful. Typically, participants in this process remove unneeded fragments, structure the fragments into tentatively meaningful groups, reconsider fragments and groups at differing levels of abstraction, reorganize the fragments based on their evolving understanding of what makes the most sense, and then solidify the final groups as meaningful, high-level and coherent descriptions [e.g. 4].

Our focus in this paper is on how distributed groups, supported by groupware, can perform collaborative interpretation. We begin by considering how emergence occurs in collaborative interpretation, and on how it can be encouraged. Because we want to support distributed groups, we then describe design principles that capture the essence of the processes that encourage emergence to happen as well as

ACM Proc CSCW'00, p289-298, December 2-6, 2000, Philadelphia, PA.

the basic requirements necessary for effective groupware. We next present our system and discuss how it embodies these principles. We close by summarizing our evaluation of the system and discuss how our observations compare to similar face-to-face activities.

EMERGENCE IN COLLABORATIVE INTERPRETATION

The inputs to collaborative interpretation are the raw information fragments. As mentioned previously, these are often very brief, sometimes to the point of being indecipherable to all but their authors [10]. The output from collaborative interpretation is a set of coherent descriptions. These contain germane explanations, descriptions and details, and are expressed at a level of abstraction appropriate to the task being addressed by the group.

Emergence is a characteristic of the process by which the group interprets and transforms these initial raw fragments into the rich final descriptions. Moran, Chiu, and van Melle [15] characterize emergence as the observation that:

> Ideas do not arise well formed. At first there are expressions of fragments of thoughts. Once there is some rough material to work with, interpretations gradually begin to emerge as they are discussed. [15, p46]

Edmonds, Moran, and Do [6] elaborate.

> Emergence happens usually after a long period of evolution in a design process. It is sometimes a background process, a change of attitude or getting information about a new aspect or attitude for perception and interpolation or shifting context and paradigm. It sometimes involves revisiting old ideas and seeing the breakthrough … to see new possibilities emerge.

Essentially, the understanding of the fragments and what should be done emerges out of an extended consideration and manipulation of the 'rough material'. This sense-making is not a simple aggregation or union of the data. Rather, it is the result of reflecting on the fragments, of making additional connections between individual fragments, and of creating new abstractions that better describe what is going on. This often goes beyond what is in the raw data. It adds the tacit knowledge of the people who produced these fragments, the grounded perspectives of all participants, and the new insights formed as a by-product of discussing and manipulating the items.

While we know that emergence happens, the question remains on how best to support it. Most researchers concur that the method should rely on manipulation of artifacts in a spatial medium that affords the following activities [6].

Arranging and spatial reasoning. People lay items out in a visual workspace and express developing relationships amongst items—even ambiguous or partial ones—via spatial proximity and visual cues [4, 12, 15, 17, 22].

Informal structuring. Because people pursuing emergent outcomes often develop new (and perhaps uncertain) concepts and structures or rearrange existing ones, the process should support informal and lightweight re-structuring of objects rather than a formal structure that may be difficult or heavyweight to alter [6].

Sketching. Some researchers emphasize the importance of quick and easy freehand sketching. They argue that sketching lets people: experiment and explore options; reduce information to abstractions that emphasize structure instead of detail; and create and alter items in spite of uncertainty and ambiguity [6].

Everyday technology easily supports these activities. In our earlier research [4], we observed how face-to-face groups performed collaborative interpretation by recording information on Post-it™ notes, sticking them onto a large whiteboard, and quickly rearranging those notes into spatial groupings by moving them about as emergent understanding unfolds. They used marking pen to draw lines, identifying labels, and other text fragments near these groupings. They sketched new ideas on both the whiteboard and the notes.

DESIGN PRINCIPLES

We now consider basic principles that should be adhered to by designers of systems supporting distributed collaborative interpretation. Specifically, we describe how systems should:

- support the key behaviors observed in people engaged in activities where emergence occurs, and
- facilitate distributed groups working together over visual surfaces.

Supporting Emergence

We begin by providing four principles for supporting emergence in groupware; most are related to the points raised previously. These principles are drawn and collected from earlier research. This includes our own observations of face to face collaborative interpretation using Post-it™ notes, marker pens and a whiteboard within the contextual setting of *results synthesis* [4], described later.

Provide a spatial visual workspace. People have highly developed spatial and visual abilities [12], and they apply them to see and manipulate items as a way of expressing their evolving understanding. More specifically, the first inclination of many people when faced with the task of making sense of a large number of related elements is to rearrange the elements spatially [14]. When the entire data set is displayed in the space, people can simultaneously consider all the data elements and can easily switch their focus from individual elements to larger groupings. Thus a system supporting activities encouraging emergence should provide a spatial visual workspace that allows participants to see and manipulate raw fragmentary data directly [4].

Let people express relationships amongst the data using spatial proximity. People use spatial proximity to express relationships by arranging related items close together. That is, the distance between any two items on the space becomes a rough measure of the relation between them. Relations can be strong or weak: items with strong relations are very close (e.g., they may touch or overlap), while items with ambiguous, tentative or tacit relations are expressed by having them somewhat further apart [4]. In many cases, seeing the spatial proximity between items suffices as a visual classification and explanation of the relationship.

Allow free-form annotation of the underlying space. Spatial proximity is limited in that it cannot add explanatory detail, and it can only signify a single dimension of relatedness. We can mitigate this limitation by letting people draw or write directly on the space [6, 4, 23]. For example, people can alter or extend existing structures by drawing boxes around items; this partitions the space and strengthens groupings. Similarly, they can indicate secondary relationships by using lines and arrows to join more distant items. With text labels, people can both explain and abstract the meaning behind groupings. Labels can also become meaningful anchors for restructuring a workspace by placing related items close by [4]. People can elaborate further by placing text descriptions (explanations, questions, etc.) close to the label and surrounding items.

Allow the free creation and movement of data in the space. Because working with emergent ideas and concepts involves much exploration, people must face minimal overhead or impediment when trying things out in the space. We believe that people should be able to easily add new elements to the workspace, where these capture parts of their evolving understanding e.g., generalizations, summaries of existing elements, or new points. We also believe that people should be able to freely move all elements in the space, so they can express whatever relationships or structures they perceive in the data. Free movement allows people to express their current understanding of relationships even though they may be tentative, ill formed, or momentarily inexplicable. Moran, Chiu, and van Melle [15] have summarized this requirement as the design principle of "agility" for their work on supporting generic meetings of small groups. They argue that agility is important because it allows people to concentrate on expressing their ideas rather than on using the system or making their ideas fit the preconceived notions that have been implemented in the system.

Supporting Distributed Activities within a Visual Workspace

Because we want to support the way *distributed* teams perform collaborative interpretation, we need to adhere to relevant distributed groupware design principles. Fortunately, there exists a body of knowledge on effective groupware design for real-time collaboration over a spatial and visual workspace [23, 7, 8, 5], which is exactly the kind of system required for collaborative interpretation.

Provide a common and visually similar environment to all participants. In order to collaborate, participants must be able to establish a common grounding for their communication [2]. In a visual workspace, this implies that people should see the same things in similar ways. Specifically, a system supporting interpretation activities should ensure that items, spatial relationships and groupings appear similarly to all participants. While people are remarkably resilient at adapting to systems that show things differently, we believe that such behavior is counter-productive as it would inhibit the rapid, exploratory and lightweight nature of activities that produces emergence.

Provide timely feedback of all actions within the workspace. Collaborative interpretation is at its best when people can see what others are doing to objects as they are doing it. People need to see visual objects as others create, move and delete them, text as it is typed, and so on. This allows them to understand what is going on, provides opportunities for interaction and discussion, and helps them better coordinate their actions [5, 7, 8].

Support gesture and diectic references. Gestures (such as pointing to objects) and diectic reference (verbal references that accompanies gestures such as 'this one') are natural parts of communication [2]. Tang [23], for example, noted that almost half of a design group's actions over a visual workspace involved some type of gesture.

Support workspace awareness. Workspace awareness is the up-to-the-moment understanding of another person's interaction with the shared space [7, 8]. It goes beyond the notion of feedback and gestures by insisting that all actions are seen within the context of the entire workspace even when people are working in different parts of it. For example, a person in a group performing collaborative interpretation over a whiteboard may be concentrating on her own area of interest. Yet through peripheral vision and by hearing accompanying sounds and speech, she can track other people's activities. She can glance up and look around to see where others are working and what they are doing. By refocusing her gaze, she can quickly engage into other areas of the workspace if called upon. These awareness acts can be difficult to do in computer systems that are limited by small displays, unless some mechanism is available to support workspace awareness.

SYSTEM DESIGN

Using these principles, we built—using the GroupKit groupware toolkit [21]—several versions of a distributed system supporting collaborative interpretation. We evaluated each version through usability tests, modifying the design where necessary [4]. In this paper, we describe the final version, which we called PReSS[1]. PReSS was designed to support a particular exemplar of collaborative interpretation called *results synthesis* [4]. In this contextualized setting, people—usually members of an interface evaluation team—take a set of usability problems (the information fragments) discovered through heuristic evaluation [19] and transform them into detailed problem reports (the coherent descriptions) that would be given to system designers.

---

[1] To encourage replication and extension of our ideas, PReSS and Groupkit are available from www.cpsc.ucalgary.ca/grouplab/. Tcl/Tk can be found at dev.scriptics.com/.

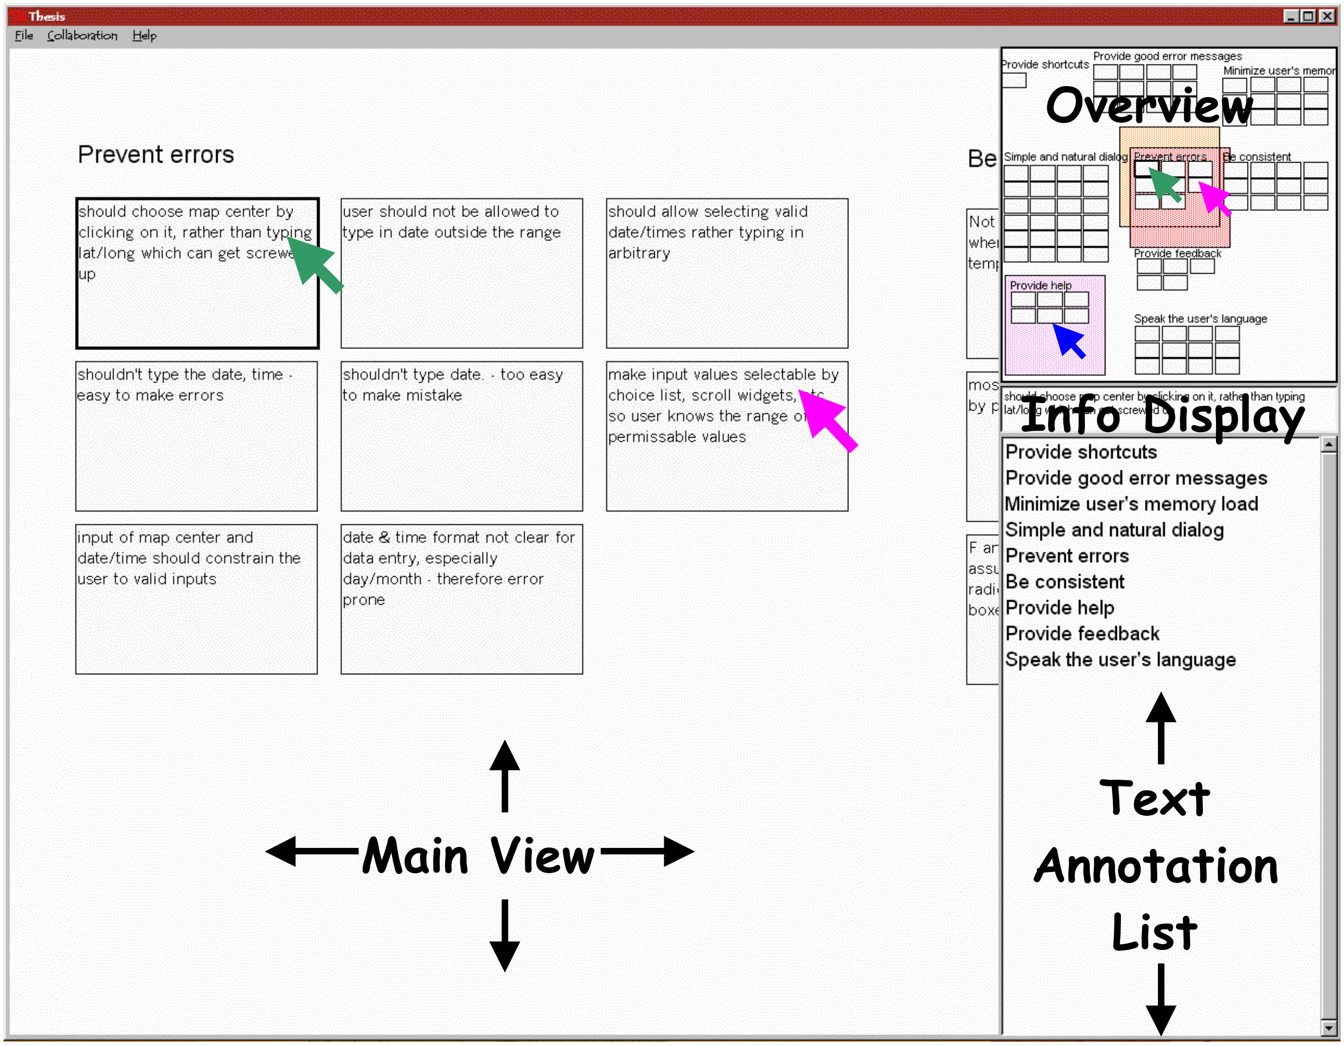

Figure 1. The initial workspace, laid out automatically into categories by PReSS. Its overview is shown larger in Figure 2a.

PReSS, illustrated in Figures 1–6, is loosely designed around the metaphor of Post-it notes and index cards that can be manipulated on a drawable surface similar to a whiteboard. As distributed groupware, particular design challenges include visualizing information on a small display (as compared to a real whiteboard), and the need to support workspace awareness. PReSS consists of:

- several primitive object types we believe are relevant to collaborative interpretation that can be created and manipulated within the workspace; and
- the visual workspace, with three different views for seeing and manipulating objects (Figure 1).

Primitive Object Types

Cards are the means by which raw information fragments are represented in the workspace (Figure 1, left pane). In the context of results synthesis, these are semi-structured fragments, where information includes a brief problem description, its severity, the name of the evaluator, the applicable heuristic, and so on. This customization of cards to results synthesis is minor, and the reader should be able to see how the cards could be adapted to other contexts of collaborative interpretation. PReSS displays only the problem description on the standard card (Figure 2 left side); secondary information appears only when a person double-clicks or edits the card (Figure 2, right side). Our reason for hiding this information is that we believe it important to fit as many cards as possible in the visual workspace while still having the fixed-sized card with its contents easily readable by using a reasonable font size. This decision is an example of a fundamental trade-off that we faced repeatedly in the design of PReSS: legibly showing detailed information while still showing as many data points as possible. Finally, cards can be cloned (duplicated) so that copies can be placed around the workspace when required.

Notes are similar to cards, except they are yellow (like Post-it™ notes). People can create, edit and place notes anywhere (Figure 2, top). We make notes visually different from cards to distinguish content added to the workspace during the interpretation from the original set of fragments. This is important in results synthesis, but may not be significant in other collaborative interpretation activities.

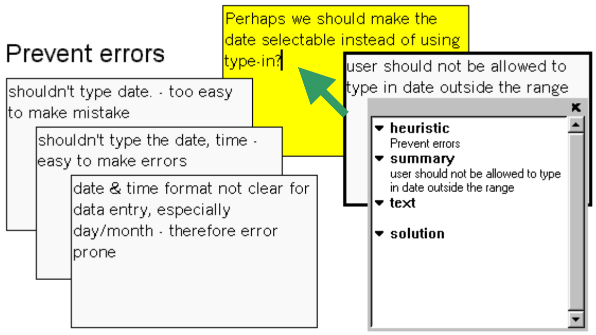

Figure 2. Four cards and one note (top). The rightmost card is displaying its semi-structured content in a pop-up.

Piles are a way of collapsing a collection of cards into a single entity. Within results synthesis, this is supposed to occur when people construe particular cards as duplicates. Figure 3, for example, illustrates a single pile comprising the duplicate cards previously seen on the left side of Figure 2. Creating piles is easy: a person just places one card directly atop another card or another pile. The person can also select which card best represents the pile by *shuffling* through its members until the desired representative is on the top of the pile: this is done by clicking the dog-ear on the pile's lower right corner, or by selecting from its context menu (Figure 3). Piles can be further manipulated: *Remove* takes the top card off the pile and places it in the workspace as an individual card, and *Unpile* separates all cards in the pile into a selected group of side-by-side individual cards. As with cards and notes, piles can be freely moved throughout the workspace.

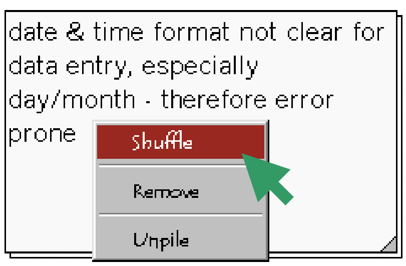

Figure 3. A pile and its context menu. Clicking the dog-ear shuffles through its cards.

Textual annotations provide another means for adding text marks to the workspace. These large text objects can be used for many purposes: for annotating existing structures, for listing possibilities [e.g., 23], and for labeling clusters. This last use is important to collaborative interpretation, as labels can be used to stabilize the emergent meaning behind the spatial arrangements of cards, notes and piles. That is, people can add labels and other bits of text to the workspace that identify, categorize, or abstract the essence of the objects around them: examples are illustrated in Figure 4 and in most other figures. People further manipulate the cards around the label to refine their emerging understanding. As in most drawing packages, text is easy to add, edit, reposition, or delete.

PReSS leverages text annotations as labels in two ways. First, the annotations provide a visible naming of surrounding organization of space. Fragments in the workspace are almost always being organized into some sort of categorization or grouping. When there are a reasonable number of fragments in the workspace, the groupings of fragments need to have their meanings stabilized as it is not

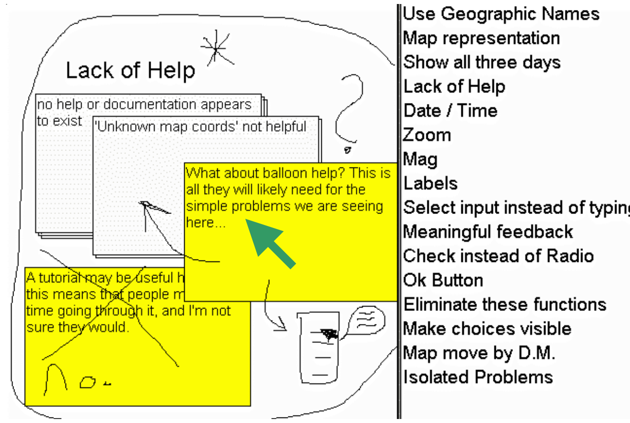

Figure 4. Text and graphical annotations. Note that the text annotation appears in the text annotation list on the right.

possible for any one person to remember the entire state of the workspace. For example, the layout in Figure 1 and 6a was automatically produced by PreSS where the raw problem statements that serve as input to results synthesis are pre-categorized by heuristic (this will be described further in a later section). While we expect these clusters and labels to change as people reflect on the emerging outcomes, they serve as good starting points. Second, text annotations, given that they name space, are leveraged for navigation. All text phrases are automatically added to a 'text annotation' list (bottom right of the screen in Figure 1). When the user clicks any phrase in the list, that user's window into the workspace will immediately scroll to bring that label into the center of it. Figure 4 illustrates a label 'Lack of Help' around several cards, as well as how that label appears on the text annotation list.

Freehand annotations are drawing marks that can be created in a manner similar to conventional paint programs. Marks are graphical objects: they can be both moved and deleted. For example, Figure 4 illustrates several freehand annotations in this workspace: arrows relate cards and sketches; a set of lines reenforces this grouping; a star-shaped glyph indicates that the nearby label is somehow important; an X tentatively discards an idea in a note; a sketch associated with a note illustrates an idea; and a question mark emphasizes an idea in a nearby note.

Reports, created via a workspace pop-up menu, contain the final coherent descriptions. Reports are usually created when the group is satisfied with a particular cluster of cards, piles and annotations. For example, Figure 5 illustrates a report that summarizes a cluster containing two piles, one card, and a text label. These surrounding objects serve as context and detailed reminders of what belongs in particular reports when the group is ready to write it. The report illustrated in Figure 5 is specifically for results synthesis. Through its tabs, people fill in sections that describe the usability problem in detail, recommended actions, and problem severity. Other report formats can be created for other collaborative interpretation contexts. As elements in the workspace, reports also act as progress indicators. They help participants track what groupings they have dealt with and what remains to be considered.

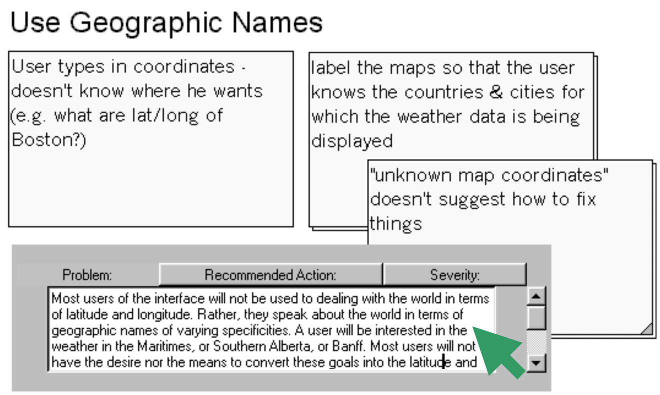

Figure 5. A report that describes in detail the cards around it.

The Visual Workspace

The central feature of PReSS is the visual workspace. Its area is large enough to fit approximately 240 non-overlapping information fragments (cards) that serve as the input to collaborative interpretation. The workspace is an undifferentiated space, and all actions can be performed anywhere within it. The workspace does not provide any explicit structuring mechanism. Of course, the workspace is completely group-aware: all participants share and see it, and anyone can do anything at anytime within it. People have three different views of the workspace on their display, as described below and illustrated in Figure 1.

The main view. The main view is a fixed-size scrollable view-port into a portion of the entire workspace (Figure 1, left side). About $1/10^{th}$ of the entire area is visible in full size, which is room to fit about 24 non-overlapping cards. They also see other people's telepointers (used for gesturing) if their viewports happen to coincide. Users typically spend most of their time in the main view, as they see all objects (cards, notes, reports, annotations etc.) in sufficient detail for easy reading and manipulation.

The overview. The entire workspace is visible as a miniature overview (Figure 1, top right). The overview provides the global context, and is especially good for helping people establish the spatial relationships between the various groupings. In addition, participants see not only their own actions reflected immediately in the overview, but all actions of others. This includes view rectangles (one per person) that show what others see in their main views, as well as all telepointers. The intent of the overview is to provide workspace awareness [7,8], where individuals can stay aware of where their collaborators are working and what they are doing. Scaling of objects within the workspace is proportional, with two exceptions. First, we do not show the text content of the cards, as the text is too small to read. Rather, an *information display*, positioned below the overview (Figure 1 right middle), displays the card contents as soon as a person passes the mouse over a particular card in the overview. This enables a person to rapidly search or browse the cards in the overview without resorting to scrolling their main view. Second, text annotations are made larger to make them easily readable. We do this because we believe they act as landmarks and thus should be visible in the overview.

The overview is active, allowing most actions available in the main view excepting those that are compromised by working at a small scale. For example, any object (cards, reports, text and free-hand annotations) can be moved around, and most object-specific actions are permitted e.g., cards can be cloned and unpiled. However, text editing is disabled. People can also scroll their main view quickly by dragging their view rectangle in the overview to the new location. For example, for close collaboration two people can overlap their view rectangles to align their main views so they are seeing the same thing.

The text annotation list, visible in the lower right-hand corner of Figure 1 and in Figure 4, contains a list of all the text annotations in the workspace and thus a summary of most groupings. As previously mentioned, it is designed to provide navigational and organizational shortcuts. Clicking an annotation in the list immediately scrolls the main view so that the text annotation is centered within it. The user can also drag objects from the main view and drop them on an entry in the list, which results in the object being 'teleported' to a location close to the annotation.

SCENARIO OF USE

To illustrate how all the above features work together, we briefly describe a scenario that captures the essence of how we observed various groups perform results synthesis—an exemplar of collaborative interpretation—using PreSS [4]. In particular, we describe how a group of evaluators transform the raw problem descriptions that are the output of heuristic evaluation into detailed problem reports that can be given to developers [10]. Our scenario begins after several geographically distributed evaluators have independently performed a heuristic evaluation of an interface. Specifically, they have already used a custom tool to capture electronically 92 raw problem descriptions, categorized by nine heuristics, which can be used as input to PReSS [4].

Preparation. When the evaluators are ready to begin collaborative interpretation, they set up a conference call and start PReSS. Through a menu option, they import the file containing the raw usability problems. PReSS then automatically organizes the 92 raw problems by the nine heuristics: a text annotation naming a particular heuristic is placed on the workspace and the relevant cards are organized beneath it. This initial layout is shown in Figure 1, with a close-up of the overview provided in Figure 6a.

Familiarization and collapsing duplicates. Participants familiarize themselves with the others' contributions by reviewing the contents of the entire workspace. This can be done directly in the main view, or by scanning through card contents by passing their cursor over cards in the overview and seeing their contents in the information display. When someone finds a card that is hard to decipher, they point to it with their telepointer and ask others what it means; the author, seeing the reference, explains it. As a part of this review, participants also notice duplicate or near duplicate problem descriptions. They collapse these into piles by placing one card atop the other. This action is usually announced briefly to the other participants by saying something like “Here’s a bunch of ‘No Help’ problems”. Figure 6b illustrates what the workspace looks like after this round; comparing it to Figure 6a shows that while the categorizations are still the same, some of the 92 cards have been relocated and moved into piles—only 56 objects (cards, piles) are now visible.

Round 1. The participants then begin reorganizing the

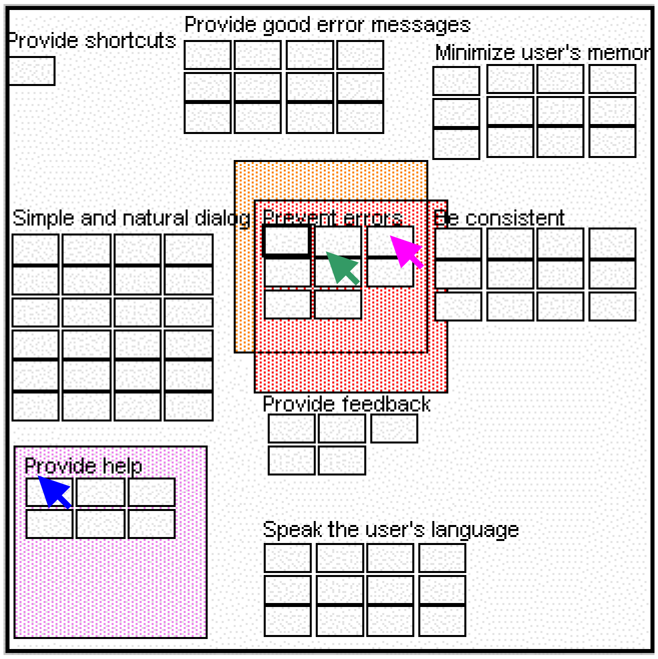

Figure 6a. Initial Organization. PReSS automatically lays out raw data by heuristic.

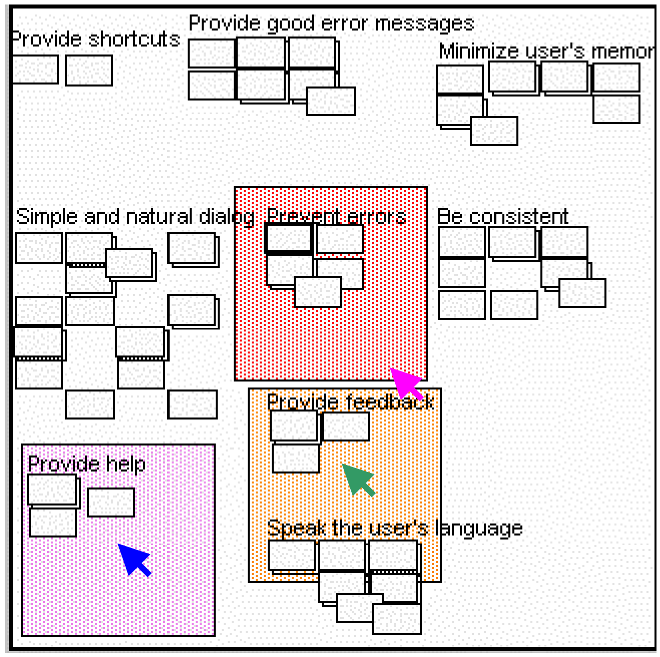

Figure 6b. Duplicates pruned. Categories are still the same, but some cards have been placed into piles.

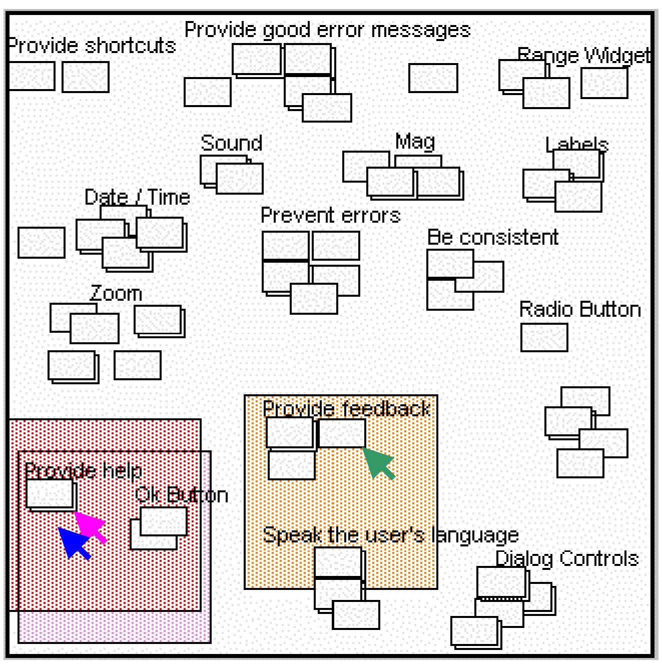

Figure 6c. Reorganization 1. Cards have been placed into different groupings, with most labeled by a text annotation. While some original categories still exist, others are gone.

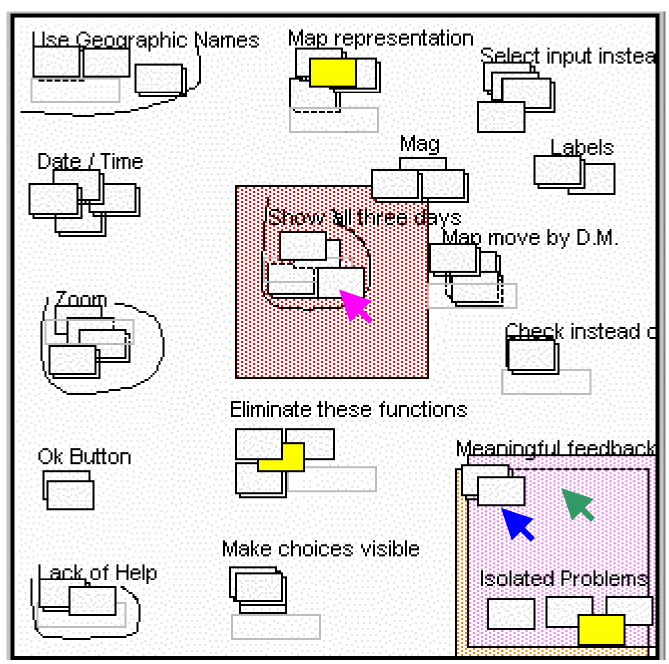

Figure 6d. Reorganization 2 representing final abstractions. Some of these have already been rewritten as reports that can be passed to developers.

Figure 6. Workspace overview snapshots taken over time, arranged as a storyboard to capture the essence of how we observed various groups perform results synthesis—an exemplar of collaborative interpretation—using PreSS.

fragments to reflect their initial understanding of what makes the most sense. In this scenario, they start moving some cards that seem related into a new part of the workspace. As this rearranging continues, they realize the groupings they are creating are organized by interface component; they then add text annotations that record this insight. This continues, with people moving many cards out of their original organization-by-heuristic into this new tentative structure. As this is done, people discuss this new organization, perhaps talking about how particular cards fit or don't fit within it. Figure 6c illustrates this new tentative reorganization: when compared to Figure 6b, we see quite a few new groupings and text annotations, and that some of the original heuristic categories have disappeared.

Round 2. As the process unfolds, participants are gaining a better insight into the emerging relationships that exists between the raw problem descriptions. The participants decide that the first go at interpretation is not satisfactory, perhaps because the organization does not account for all the problems contained by the cards, or because glimmers of better abstractions emerge. They negotiate the outline of a new structure: they discuss particular cards and groupings while gesturing over the workspace, and they move items around tentatively to 'sound out' the reaction of others.

They then move about the workspace, reorganizing the cards, notes and text annotations according to their new shared interpretive structure. They next go about refining the interpretation present in the workspace by adding more text annotations and free-hand annotations (and deleting old ones) to create a more nuanced interpretation of the spatial relationships displayed by the cards. Figure 6d illustrates how many of the categories and groupings in Figure 6c have changed. Freehand annotations around several groups and the tighter clusters also suggest that some participants now consider these new categories as fairly stable.

What has happened in these rounds is that participants are reorganizing, and as a consequence reconceptualizing, the entire data set: new meaning emerges. In doing so, they make use of the entire workspace, exploiting the spatial environment. Reorganization happens at not only at a macro level but also at the micro level. At the micro level, first the cards are arranged in a haphazard manner around a text annotation. Once all the cards are roughly grouped, the participants return to each grouping to impose a more nuance expression of the relationships between the cards grouped in each area. In doing so, they are using spatial proximity at a more local level to reflect a more sophisticated understanding of what is wrong with that particular element in the interface. The participants continue to organize at the micro level until they have straightened out the entire workspace.

Finalizing. As the participants come to agreement about the structure that has emerged, and as most of the raw problem descriptions are accounted for, they translate the knowledge encoded in spatial layouts and annotations into a form that is more easily communicated to people who have not participated in the session. The participants do this by creating a report object, usually placing it next to an associated grouping. They then write detailed problem descriptions for each identified grouping. Figure 6d illustrates a group finalizing their categories: some (but not yet all) categories have an associated final report (the longer rectangular box). As with all their activities, they can do this individually (where others are still aware of what a person is doing), and collaboratively (with several participants discussing and working on a single report at the same time). When they are done, they generate a report by using the PReSS export option, which generates a formatted html file containing all the final problem reports. This report is then passed onto the developers.

EVALUATION

Through a series of observational usability studies, we evaluated how results synthesis, as a contextualized exemplar of collaborative interpretation, is supported in both a face-to-face paper-based process [4] as well as electronically via PReSS. Evaluation of the system was formative rather than summative: we used it to verify design approaches, to detect and repair usability problems (see [4 Ch. 5-6] for the design evolution), and to see if the system supported a process similar to what we saw when groups used paper.

The system evaluation was divided into two stages: three single user studies and five multi-user studies. Studies with single users were done first to drive out software defects and establish a base level of usability. We then moved onto multi-user trials as soon as system operation was mostly defect free. In all studies, participants had previously performed a heuristic evaluation of the interface described in Exercise 8 in Nielsen's *Usability Engineering* [19]. After briefly demonstrating PReSS's capabilities, we gave them approximately 90 raw problems produced by all evaluators (including themselves) from the heuristic evaluation: the initial screen appeared similar to Figure 6a. We then observed them using PReSS to perform results synthesis for about an hour. We did not collect any specific quantitative measures; rather, we took notes whenever we saw or heard something we judged to be of interest i.e., what seemed to work and where breakdowns occurred. This was followed by a retrospective interview. Full details of the study and its results are reported in Chapter 6 and the Appendices of Cox's thesis [4] (an online copy is available in the 1998 section of http://www.cpsc.ucalgary.ca/grouplab/papers/).

Single users studies identified a number of defects in the system that were repaired as they were discovered. Aside from a number of routine low-level fixes, we added the information display seen in Figure 1 (to allow people to quickly search the workspace using the overview), as well as the addition of the note object (because people wanted to add more fragments to the workspace as they were working with its contents). We also noticed users behave differently from what we expected in their extensive and liberal use of piles.

While the original design intent of a pile was strictly for grouping duplicates, people used it as a more general grouping mechanism i.e., to collect all the problem fragments that they felt represented a single problem description. We believe this happened because people wanted to reduce the apparent complexity of the workspace: people mentioned that they initially felt overwhelmed by the large number of cards. On the balance, however, individuals used the system as anticipated, using a number of navigation mechanisms to move about the workspace, exploiting spatial layout in creating interpretive structures, and creating new groupings. All these support the validity of our design principles.

Multi-user studies. We then observed several cases of small groups (2 to 3 people who knew each other) performing collaborative interpretation in PReSS. The study environment simulated remote interaction: while participants were in the same room, they faced away from each other and could not see each other's screen directly. Participants were, of course, encouraged to talk with each other during the task.

In general, PReSS worked well as groupware for collaborative interpretation: people used it largely as expected and similar to how they did it on paper. Participants were able to do both individual and collaborative work, moving between the two extremes with little effort. Participants made and understood deictic references. They moved problems out of their initial groupings and into new categories of their own creation. Spatial layout was used extensively.

However, several aspects of the collaboration were somewhat inhibited due to factors that we believe are related to workspace awareness issues. First is the extra effort required to synchronize views when the participants needed to work together. Compared to the ease of glancing around and focusing on a common place in face to face settings, scrolling to a common place for sharing detailed work is effortful. Still, they were able to accomplish these transitions. Second, while participants worked both as individuals and as groups, the coordination did not seem to be as close as in the face to face situation—there was more working apart and, as mentioned previously, there was more effort expended re-synchronizing views. the three-person studies, we seldom saw all participants engaged simultaneously on the same task in the same part of the workspace after the familiarization stage. Third, understanding other people's references was more difficult. Especially in the three-person study, participants asked for clarification of references so they could figure out what was being referred to. A typical exchange might go:

P2 What do you think of this? *Clicks on card and jiggles it.*
P1 Where?
P2 This one here. *Jiggles card more demonstratively*
P1 Oh, yeah, that seems OK.

All these problems are not fatal, for in general participants did exploit the system features in much the same way as had been seen in the previous studies (although perhaps not as gracefully). We see these problems as indicating that we need to improve further our techniques for information visualization and for supporting workspace awareness [7, 8].

RELATED WORK

A number of other systems encourage emergence similar to how is seen in collaborative interpretation. Notecards [9, 14] is an early and excellent example of a (mostly) single user hypertext authoring system that allows people to create, manipulate and organize information fragments based upon a card metaphor. Our own system is most like VIKI, which was designed to promote emergence through spatial hypertext [12, 13]. VIKI does not provide any support for collaboration beyond allowing multiple people to access the same workspace asynchronously. Thus VIKI is focused more on longer-term loosely-coupled group work both by design and in the reported case studies of it use. VIKI has three kinds of entities in its workspace: objects, collections, and composites. VIKI objects are much more general versions of the cards in our system, allowing any number of named "slots" which can hold text. There is a lightweight user-extensible type system for the objects that allows objects to be created with a default appearance and set of slots. However, any individual object may be extended or modified in appearance or content. Collections and composites are structuring mechanisms for objects in the workspace. Collections are large sub-spaces that appear in the workspace as a window onto the subspace. Composites are "lightweight structures that consist of two or more objects or collections in particular visual/spatial configurations." [13 p17] These composites may be user specified, or the VIKI system has a background recognizer that will suggest the creation of composites when it notices recurring patterns of visual/spatial configuration.

Tivoli [15, 17] and the Insight Lab [11] are systems designed to support groups in real-time co-located intellectual activities that feature emergence and sense-making. Tivoli is based around an electronic whiteboard, and its focus is to provide a fluid interaction style using a pen as primary input device for both meeting-generated content and control input. The Insight Lab uses physical products (papers, conventional whiteboards that can be printed out), where barcodes link digitized versions of these products to a multimedia store. Similar ideas are found in Collaborage [16]. As with PReSS, Tivoli supports the use of pre-generated material as the basis for the meeting activity [17]. It also supports annotation of the workspace and grouping of the contents of the workspace. The largest difference between PReSS, Tivoli, Collaborage and the Insight Lab is that PReSS is aimed at distributed groups simultaneously accessing the workspace using conventional computer screens (in Tivoli, a single person manipulates a very large display whereas the Insight lab uses physical artifacts).

Gungen [18] is a system designed for real-time collaboration amongst distributed groups performing the KJ method, a technique for organizing ill-structured data. The KJ method is a very formal technique that, while a form of collaborative interpretation, may not be appropriate for all contexts. PReSS differs from Gungen in that it provides much better support for collaboration over the workspace e.g., for workspace awareness [7, 8].

CONCLUSION

Many actual work tasks can be described as collaborative interpretation of fragmentary source material. In these tasks, emergence is a key characteristic of how the work unfolds. The outcome of interpretive tasks cannot be predicted before hand, but rather must emerge out of the background of the interpreters interacting with the source material in the context of their particular goal.

Our goal was to support collaborative interpretation for small distributed teams. To that end, we contributed several design principles drawn from two relevant areas: how to support activities where emergence happens amongst distributed teams, and how to support collaboration over a shared workspace. We then applied these principles to our design of PReSS, a system for collaborative interpretation.

Our evaluation of PReSS is promising. We saw that small groups could perform results synthesis—a contextualized exemplar of collaborative interpretation—using a simple metaphor of cards and annotations on a drawable surface. Yet further refinements are needed. Perhaps the biggest problem is that the small workstation screen is a poor stand-in for a large visual workspace such as a whiteboard. While PReSS uses a variety of techniques to legibly display a reasonable number of cards and their spatial relationships, scanning the workspace is still more effortful when compared to glancing around at cards placed on a whiteboard. Another problem is that PReSS's techniques for workspace awareness support, while modeled on strong empirical research [7, 8] is still crude when compared to how workspace awareness is done in face to face settings. Finally, PReSS's card metaphor currently works only with text. In practice people may want to perform collaborative interpretation on other media forms e.g., drawings, photographs, video, spreadsheets, or even physical artifacts.

ACKNOWLEDGMENTS

This work was partially funded by NSERC, and by Microsoft Research. Carl Gutwin influenced our work.

REFERENCES

1. Byrne, J. and Barlow, T. Structured brainstorming: A method for collecting user requirements. *Proc. Human Factors and Ergonomics Society 37th Annual Meeting* 1, 427-431, 1993
2. Clark, H. *Using Language*. University of Cambridge Press, 1996.
3. Conklin, E. and Begeman, M. GIBIS: A hypertext tool for exploratory policy discussion. *Proc. ACM CSCW'88*, 140-152. 1988.
4. Cox, D. *Supporting Results Synthesis in Heuristic Evaluation.* M.Sc. thesis, Dept Computer Science, University of Calgary, Calgary, AB Canada, 1998.
5. Dix, A., Finlay, J., Abowd, G. and Beale, R *Human-Computer Interaction*, Prentice Hall, . 1993.
6. Edmonds, E., Moran, T. and Do, E. Interactive systems for supporting for the emergence of concepts and ideas. *SIGCHI Bulletin* 30(1), 62-76, 1998
7. Gutwin, C. and Greenberg, S. The effects of workspace awareness support on the usability of real-time distributed groupware. *ACM TOCHI* 6(3) 243-281, 1999.
8. Gutwin, C. *Workspace Awareness in Real-Time Distributed Groupwar*e. Ph.D. Thesis, University of Calgary, Calgary, AB Canada, 1997.
9. Halasz, F., Moran, T. and Trigg, R. Notecards in a nutshell. *Proc. ACM CHI+GI'87*, 45-52, 1987.
10. Jeffries, R. Usability problem reports: Helping evaluators communicate effectively with developers. In J. Nielsen & R. Mack (Ed) *Usability Inspection Methods.* Wiley 1994.
11. Lange, B., Jones, M. and Meyers, J Insight lab: an immersive team environment linking paper, displays, and data. *Proc ACM CHI'98*, 550 – 557, . 1998.
12. Marshall, C. and Shipman, F.M. Spatial hypertext: Designing for change. *Communications of the ACM*, 38(8), 88-97, 1995.
13. Marshall, C.C., Shipman, F.M., and Coombs, J.H. VIKI: Spatial hypertext supporting emergent structure. *Proc. of ECHT '94*. 13-23, 1994.
14. Monty, M. *Issues for Supporting Notetaking and Note Using in the Computer Environmen*t. Ph.D. Thesis, University of California, San Diego, USA, 1990.
15. Moran, T., Chiu, P. and van Melle, W. Pen based interaction techniques for organizing material on an electronic whiteboard. *Proc. ACM UIST'97*, 45-54, 1997.
16. Moran, T., Saund, E., van Melle, W., Gujar, A., Fishkin, K. and Harrison, B. Design and technology for collaborage: collaborative collages of information on physical walls. *Proc ACM UIST'99* 197-206, 1999.
17. Moran, T., van Melle, W. and Chiu, P. Tailorable domain objects as meeting tools for an electronic whiteboard. *Proc. ACM CSCW'98,* 295-304, 1998.
18. Munemori, J. and Nagasawa, Y. GUNGEN: Groupware for a new idea generation support system. *Information and Software Technology*. 38, 213-220, 1996.
19. Nielsen, J. *Usability Engineering.* Academic Press, 1993.
20. Nunamaker, J., Dennis, A., Valacich, J., Vogel, D. and George, J. Electronic meeting systems to support group work. *Communications of the ACM,* 34(7), 40-61, 1991.
21. Roseman, M. and Greenberg, S. Building real-time groupware with GroupKit, a groupware toolkit. *ACM TOCHI* 3(1), 66-106, 1996.
22. Shipman, F.M., and Marshall, C.C. Formality considered harmful: Experiences, emerging themes, and directions. Xerox PARC Report ISTL-CSA-94-08-02, 1994.
23. Tang J Findings from observational studies of collaborative work, *Int J Man-Machine Studies* 34, 143-160, 1991