In Computer Supported Cooperative Work: The Journal of Collaborative Computing 6: 327-368, 1997.

Hypermedia Use in Group Work:
Changing the Product, Process, and Strategy

Gloria Mark
FIT-Institute for Applied Information Technology

Jörg M. Haake, Norbert A. Streitz
IPSI - Integrated Publication and Information Systems Institute

GMD - German National Research Center for Information Technology
gloria.mark@gmd.de, {haake, streitz}@darmstadt.gmd.de

Abstract. Hypermedia structures have been integrated with CSCW functionality to develop the DOLPHIN system, an electronic meeting room environment. In this paper, a study is reported investigating how the DOLPHIN environment affects group work. Different aspects of group problem solving were examined to understand the effects of working with hypermedia: the group's product, cognitive factors, and the group process. The results showed that groups can easily work with hypermedia structures, and that these structures influence groups to produce a different product, to use a different strategy, and to use a different collaborative style, namely of dividing up their labor. The experimental results are explained in a model which suggests the involvement of both procedural and semantic components in hypermedia use. We discuss wider implications of hypermedia for CSCW and group work.

Key words. Hypermedia, cooperation support, electronic meeting room, electronic whiteboards, group process, division of labor, collaborative style, evaluation, empirical study

Introduction

As new technologies provide more options for ways that groups can collaborate, the social and cognitive effects on group work need more attention. In the same way that technology has introduced major changes in document production for single users, such as in speed of production, ease of modifying, production of multiple copies, and archiving, to name just a few, we can expect that comprehensive meeting room technologies designed to aid document production for groups will also cause substantial changes. Advantages of group products using electronic tools have already been shown in tasks ranging from brainstorming, where technology has helped to increase the number of ideas that groups generate (Hymes and Olson, 1992) to writing, where a collaborative tool has aided groups in producing documents of a higher quality (Olson et al., 1993). However, because of the variability of groups, the results from using technology are not always predictable from one situation to the next.

With respect to social aspects in the group process as a result of working with technology, important changes are also evident. For example, working with technology affects the work style of a group (e.g. Stefik et al., 1987), the amount of time a group assigns to task components (Olson et al., 1993). and the frequency of group interaction, generally by decreasing it (Watson et al., 1988).

Yet in comparison to the reports of design innovations of electronic collaboration systems, the number of empirical studies addressing the results of group work, and the effects of such systems on social and cognitive processes of groups have remained scarce. This is a deficit since an understanding of how technology affects group work is extremely important in guiding design. In this paper, we report on a study which addresses exactly this question: how a specific aspect of a collaborative technology affects the product, strategy, and process of group work. Although we understand the importance of examining work in a real-world context, we decided instead to examine the effects by taking an experimental approach. By using an experimental approach, we are able to control certain factors, and isolate others, enabling us to focus on specific changes which might occur as a result of introducing a manipulation: in this case, hypermedia structures which are provided in a CSCW system as in the case of the DOLPHIN system (Streitz et al., 1994).

This paper is organized as follows: In the first section we present our view of how hypermedia can be successfully integrated with CSCW, as realized in the DOLPHIN system, situated in an electronic meeting room. In section 2, we discuss our hypotheses of how hypermedia can influence group work, focusing on synchronous CSCW environments. We describe how we would expect hypermedia to affect the group’s products, the group’s strategy for structuring information, and the group process, summarizing our hypotheses into some basic experimental questions. In section 3, we describe the experimental method used, including the setting, the evaluation infrastructure, the design, task, and measures used to evaluate the data. In section 4, we present our results, in terms of the four experimental questions that we posed. In section 5, we discuss our results, and in section 6, we propose a model which integrates the experimental evidence to explain how the group process is affected through hypermedia use. In our conclusions in section 7, we present an overview of how hypermedia might have broader implications for group work in other task domains.

1 Hypermedia-based CSCW systems

1.1 The DOLPHIN system

The design of DOLPHIN (Streitz et al., 1994) is grounded in a conceptual framework which brings together the research areas of CSCW and hypermedia in combination with innovative user-interfaces for the corresponding functionalities. While a wide range of collaborative work can benefit from the approach taken, one can characterize our starting point as support for document-based activities, i.e., work where documents are either produced or are used as a means during group work. In our approach, we use the notion of "documents" in a very general and comprehensive way. Documents are not only traditional memos, letters, articles, or books, but they also include sets of scribbles and drawings on a whiteboard or an overhead transparency, and collections of information items which can contain multimedia elements (pictures, audio, video) which are usually available in electronic formats. In general, we consider documents to be hypermedia documents.

The design of DOLPHIN is guided by recognizing two central roles of hypermedia for CSCW (Streitz, 1994): 1) to provide a representation of the content and subject matter, e.g. in terms of a hyperdocument, and 2) to also provide a medium for cooperation and coordination in cooperative work.

Considering documents to be hypermedia documents implies the possibility of employing nonlinear network structures of complex relationships (links) between information components (nodes) as it is defined by the concept of hypertext (e.g., Nielsen, 1995). Defining relationships between different tasks, coordinating and assigning tasks to different group members, commenting and communicating on the progress of parts of the overall work activities, and much more, can be represented as hypermedia structures. At the same time, this information can be linked to the (hyper)documents which are created and used in the course of group activities. DOLPHIN now combines this hypermedia functionality with a graphical presentation of network structures and a pen-based user-interface including gesture recognition. Being also a cooperative system, DOLPHIN allows the separation and / or sharing of information and workspaces via its CSCW functionality. We propose DOLPHIN as an example of our general statement about the two beneficial roles of hypermedia structures in cooperative work environments because not only can they provide multiple ways and parallel views of a problem representation (see Gick and Holyoak, 1983), but they can also provide an efficient means for groups to work collaboratively.

From a CSCW perspective, we have focused on the support for face-to-face meetings and their equivalents facilitated by telecooperation techniques resulting in distributed "virtual meetings" (Johannsen et al., 1996). Observing the central role of public displays as they are provided by overhead projectors and whiteboards in face-to-face meetings, we investigate what kind of added value could be provided by using an electronic, i.e. a computer controlled, whiteboard either as a passive or an interactive display. Another line of research for computer support in face-to-face meetings focuses on providing a computer to every participant in the room. Using common application software, each participant can actively enter, edit, and use information while sitting in the meeting room. Our approach is based on recognizing that a combination of two scenarios is needed, i.e., an interactive electronic whiteboard *and* computers for each participant. As group activities are dynamic, rather than static, their support requires a high degree of flexibility. Flexibility can be provided in many ways. It is our view that hypermedia systems are able to provide this flexibility and serve as an ideal basis for building the next generation of cooperative information systems.

From a hypermedia perspective, flexibility translates into investigating which class of hypermedia structures should be provided for supporting group work. A prominent aspect is the modularity and inherent annotation capability of hypermedia structures. Hypermedia allows flexible decomposition, restructuring, and reuse of components in a dynamic fashion. Decomposition can be used for domain structuring as well as for structuring the division of labor in the group. Another aspect of flexibility refers to the degree to which a system is able to provide a wide range of structures in accordance with the requirements of the tasks, their coexistence and means for transformation (Haake et al., 1994). Combining these requirements, we propose to provide a basic node-link hypermedia model for structure representation. In addition to and coexistent with this model, we support more informal ways of communication by "free form" scribbling, drawing, gesturing, etc. as is the case on (electronic) whiteboards. Since DOLPHIN is a cooperative hypermedia system, this range of flexible structures can be used by each group member simultaneously at the workstation and/or at the public electronic whiteboard. DOLPHIN's collaboration functionality provides shared workspaces between all group members with an additional distinction between private and public workspaces (Streitz et al., 1994).

We describe now in more detail the functionality of DOLPHIN, organized in two sections: its hypermedia functionality and its cooperation support functionality.

1.1.1 Hypermedia functionality

DOLPHIN provides the ability to create and edit scribbles, text, nodes and links. A scribble is a freehand writing or drawing, such as exclamation marks, or figures like boxes, circles, arrows, tables, etc. Text is a string of ASCII characters.

DOLPHIN documents begin with a top node corresponding to the highest level of the document. Nodes consist of a title and content. The content of a node can contain scribbles, text, and links to other nodes. The content of a node is displayed in a DOLPHIN window. DOLPHIN provides operations for creating, editing, selecting, moving, copying, pasting and deleting scribbles, text, nodes and links.

Links connect different nodes. DOLPHIN supports two kinds of links: inter-node links for navigational purposes and intra-node links as a kind of graphical representation of relationships between nodes. Inter-node links start in the content of a node and end in the content of another node. They are displayed in the source content by a node representative for the destination node (thus, functioning as anchors). Following the inter-node link leads to the content of the destination node and displays it in a new DOLPHIN window. Users can create such a link either by creating a new node and including its representative in the currently displayed content, or, they can copy an already existing node representative and paste it into a different content. In hypermedia systems, these links are usually called embedded links. With inter-node links users can create hierarchical as well as non-hierarchical structures between nodes. Intra-node links connect two nodes contained in the same content. Users can create these links between any two node representatives in the same content. These links present themselves as arrows. In hypermedia systems, they are usually called node-to-node links.

Using the above types of objects and operations, users may create different structures, ranging from hierarchically nested structures, i.e. each node at a higher level of the hierarchy contains the nodes of the next lower level (thus forming tree-like structures) to nonlinear structures where nodes are included in the content of several other nodes (thus constituting nonlinear graph structures). Users can also create graphical arrangements of objects in each node's content. Figure 1 shows an example screen of a DOLPHIN document as it was actually produced during the experiment described in this paper. The DOLPHIN interface has been redesigned and now looks different (see Bapat et al, 1996, Streitz et al, 1997).

Figure 1: User Interface of DOLPHIN at the time of the experiment

The pen-based user interface of DOLPHIN provides gestures for creating, deleting, moving, and selecting objects as well as for opening a node's content. An always visible menu with buttons for cut-

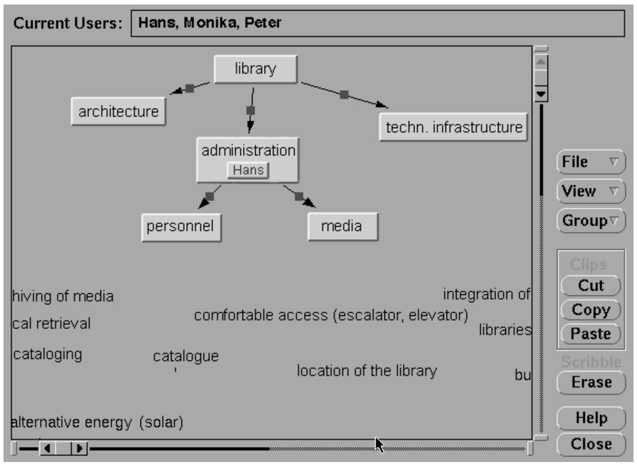

copy-paste operations, closing of DOLPHIN windows, and erasing of scribbles is provided at the edge of each DOLPHIN window. The DOLPHIN objects are mapped onto a general hypermedia data model provided by an underlying cooperative hypermedia server. Each object includes its content and additional presentation attributes (e.g., size, color, position).

1.1.2 Cooperation support functionality

When working with DOLPHIN, users can share the content of a node, thus using it as a public workspace which displays all nodes, links, text and scribbles at that level of the hypermedia structure. Users can also decide to work in different nodes in parallel, thus each using a private workspace. Within a shared workspace, DOLPHIN supports concurrent operations performed by different users. Shared access and active update/synchronization of concurrent DOLPHIN windows displaying the same node's content are provided by DOLPHIN's cooperative hypermedia server. All changes to the DOLPHIN hypermedia document are reflected in the hypermedia server and made persistent. DOLPHIN also provides different modes of coupling in shared workspaces. This allows users to scroll to different areas of the content of a shared node independently. This mode can be considered as relaxed WYSIWIS (What You See Is What I See). Another cooperative feature is that users are provided with awareness information by a user list and so called "activity markers". The user list provides the names of all users who are currently sharing the content of a node displayed in this window (see the three names “Hans, Monika, Peter” at the top of Fig. 1). Furthermore, one can see who is currently active in a node which is one level down in the structure displayed in the current window (see the name “Hans” on the node “administration” in Fig. 1).

Further technical information on DOLPHIN can be found in Streitz et al. (1994) and Haake et al. (1994). It was our goal to conduct an experiment to address the two roles of hypermedia, namely providing a representation of content, and serving as a medium for cooperation, and to investigate how these might affect the collaborative process. To this end, we tested different versions of the DOLPHIN functionality which will be described in section 3.

1.2 Related work

Since the evaluation of using DOLPHIN in our OCEAN Lab environment (see section 3.1 Setting) is concerned with (at least to our knowledge) a unique combination of physical setting and cooperative hypermedia functionality, it is difficult to directly compare it with other approaches. But there are, of course, forerunners in the development of technologies relevant for DOLPHIN and the OCEAN Lab. Due to space considerations, we do not compare DOLPHIN to the full range of hypermedia systems or CSCW systems in general. For an overview of hypertext / media systems see, for example, Nielsen (1995). For an overview of CSCW systems see, for example, Baecker (1993).

We rather concentrate on systems more closely related to our setting and compare our approach to these by categorizing them according to the following list of features provided by our physical setting and software (see also section 3): (a) a large public display, (b) used as an interactive pen-based electronic whiteboard operated by users standing in front of it, (c) connected to networked computers provided to everybody. (d) The software provides cooperative editing, (e) including shared documents and workspaces as well as hypermedia functionality.

The CoLab environment (Stefik et al., 1987) and the Capture Lab (Mantei, 1988) are examples of the combination of (a), (c), and (d). We do not consider the use of a public display as an equivalent of a traditional whiteboard / blackboard when it is only remotely controlled by a workstation, e.g., operated by a scribe or by a participant at the table. In the NICK project (Rein & Ellis, 1989), there was one experimental condition of networked computers only (c) while another condition of the NICK experiment realized an example of (b). Here, the public display offered the possibility of direct drawing on its surface as well as being operated by a scribe. A more recent example meeting our criteria of an electronic whiteboard is the LiveBoard (Elrod et al., 1992) in combination with appropriate software, e.g., Tivoli (Pedersen et al., 1993) which employs a pen-based user-interface providing support for free form scribbling and drawing. Another hardware platform is a SMARTBoard in combination with the SMARTIdeas software which supports concept maps (SMART, 1997). A prominent example of software used in a setting like (c) with or without (a) is GroupSystems (Nunamaker et al., 1991) supporting a variety of group activities. Of course, the scenario (c) requires software which provides appropriate cooperation functionality (d), e.g., ShrEdit (McGuffin & Olson, 1992). There are a number of approaches differing in their goals and functionality. An early example is Cognoter as part of CoLab providing WYSIWIS and the distinction between public and private workspaces. Due to space limitations, we cannot discuss systems which are primarily used in telecooperation situations, e.g., ClearBoard (Ishii et al., 1993), cooperative hypermedia systems, e.g., DHM (Grønbaek et al., 1994), SEPIA (Streitz et al., 1992, 1996), or collaboration in distributed design (Marmolin et al., 1991). A special case is the provision of cooperative hypermedia functionality (e) in meeting rooms in combination with all of the above features. DOLPHIN provides this and also allows extensions for distributed "virtual meetings" (Johannsen et al., 1996). For a further summary of empirical evaluations of electronic meeting rooms, see McLeod (1992). In addition, there will be more references to related work when discussing our results and conclusions in relationship to other findings in sections 5 and 7.

2 Hypermedia and Group Work

In this section we discuss how, in our view, we see hypermedia benefiting group collaboration. We understand the benefits of hypermedia for groups in terms of the two roles of hypermedia for CSCW (Streitz, 1994). First, as a unique way of representing content, we discuss how hypermedia can facilitate problem solving by enabling a more comprehensive problem representation. Second, as an aid in group communication and coordination, we discuss how the group process can benefit from hypermedia. We discuss the importance of considering the flexibility of group work and the role that hypermedia can play in supporting flexible work arrangements. We then complete this section by presenting our experimental questions.

2.1 Hypermedia and the group's product: an advantage for multiple views?

Using hypermedia structures, one can construct multiple and parallel views of a problem representation. One means for doing this is via multiple links, which suggest their relationship to several concepts, and not just to a single concept, which occurs when ideas are structured in a linear outline form. The different relationships to the original node may highlight different aspects of the node, allowing different representations to be created. Psychological evidence suggests that the quality or accuracy of the problem solution is dependent on the appropriateness of the external problem representation (Newell, 1980, Schwartz, 1971, Streitz, 1985). Opportunities for multiple views and representations of the problem structure can facilitate the problem solving process (Gick & Holyoak, 1983, Mayer & Greeno, 1975). For example, Gick and Holyoak showed that when a verbal problem was presented in a visual analogical form, more subjects were able to solve the problem.

Structuring information into nodes and links requires that people specify and separate the concepts and define their relationships. On the one hand, this could be limiting, since people may tend to censor ideas which they do not deem important enough to convert into nodes. On the other hand, the capability of separating concepts could lead people to expand and deepen their information structure by clarifying differences between ideas. It may be easier for ideas to be developed and relationships to be established if they are first clearly separated.

This possibility of benefits in an expanded problem representation led us to investigate whether DOLPHIN's hypermedia functionality compared to DOLPHIN without hypermedia functionality would cause people to structure their information differently.

2.2 Hypermedia and the group strategy: new strategies for new information structures?

According to the depth of processing view of memory (Craik & Lockhart, 1972), information that is more deeply processed should be reflected by a better memory of the information. Therefore, we expect that the decision-making process of turning information into nodes should result in deeper processing of the information, and would be reflected in a better memory of the ideas compared to groups not using hypermedia structures. In addition, we expect that the higher the level of the hyperdocument structure, the better would be the memory of the information at that level due to repeated exposure, importance of information, and strength of association (Wingfield & Byrnes, 1981).

The information structures that one can build with hypermedia may also lead to the use of new strategies in planning and organizing information. Hypermedia may facilitate the clustering of information, since users must reason in terms of nodes, i.e. separate pieces of information, and the relationships between them. The separation and clustering of ideas may involve new strategies to create an overall structure.

2.3 Hypermedia and the group process: flexibility for collaborative style?

In addition to the possible benefits of providing an enhanced problem representation, shared hypermedia may also provide advantages for the group process. Hypermedia provides the potential for creating modular structures which, in a collaborative environment, can be beneficial for the distribution of tasks among group members. Nodes and links provide a method for separating and relating information which we expect would facilitate parallel work.

Groups may choose to work in various types of work arrangements, dividing up their labor to work in parallel, working together collectively, or varying their work arrangement throughout a meeting. By "collective work", we refer especially to synchronous work as in meetings, when a group works jointly; this is generally in a serial, turn-taking fashion. We introduce the notion of collaborative style as a way of expressing that the type of work arrangement groups adopt may not be a fixed process, but may rather be a style for a group influenced by various factors. A group’s work style may vary with the situation, the task, the group, and the technology (Mark et al., 1996). Situational factors, such as the structure and culture of an organization could influence a group to work in parallel or collectively, as could the group’s leadership style (Schein, 1990). Time pressure on a group could also influence a group to change its work style, particularly if dividing up the task could save the group time (Posner and Baecker, 1992). Characteristics of the group also can influence the work arrangement, such as the roles individuals adopt in a group (Zigurs and Kozar, 1994), maturity of the group (Posner and Baecker, 1992), individual group member characteristics, expertise in different areas, or even the general group atmosphere.

2.3.1 Collaborative style and task phase

Collaborative style can change with the task phase, but unfortunately, the results of groups do not provide a consistent picture. In a study of real-world collaboration in research, Kraut et al. (1988) report that collective work is generally done in the planning stages. In the case of collaborative writing, Posner and Baecker (1992) confirm that collective work functions best in the early stages of planning an outline, before individual ideas get a chance to mature. However, in an empirical study of collaborative writing in the Capture Lab, a meeting room where workstations are networked to a public screen, groups using technology engaged in less group discussion in planning than groups without technology (Horton et al., 1991). It is unclear why technology has such an effect.

In later stages of the task, where ideas are developed, some studies suggest that the task is divided up. Real-world collaborators generally work in parallel in the execution stages of the plans (Kraut et al., 1988). Groups using the Capture Lab also worked in parallel in the writing stage (Horton et al., 1991). In an empirical study using ShrEdit, a shared editor which enables users to open multiple windows, Olson et al. (1993) found parallel work to occur during the brainstorming phase, and a "divide and conquer" phase where the members divide the task into sections and develop the sections. However, a theory is still lacking as to how technology might support parallel work in these task stages.

One possible explanation for collective work during planning is that groups prefer to discuss their plans with each other and that speech constrains the collaborative style. Speech is more easily processed serially. For this reason, Whittaker et al. (1991) argue that speech imposes a turn-taking style on the group. This idea is supported by experiences with Cantata, a conferencing system using networked personal workstations. Although parallel conversations were possible, users generally communicated serially. The designers claim that turn-taking helps maintain the continuity of a topic and propose integrating turn-taking mechanisms for the group that emulate round-robin, first-in-line, and chairman mediation (Chang, 1989).

In the brainstorming phase of a task, some results now suggest that when members have the opportunity to brainstorm in parallel using technology, they will produce more ideas than nominal groups. This occurred with groups of more than nine members using the Arizona Brainstorming tool (Dennis et al., 1990). In a parallel work condition, groups brainstorming using ShrEdit produced more ideas than groups in a serial work condition (Hymes and Olson, 1992). The authors propose that parallel work could counteract production blocking, when ideas are censured, repressed, or forgotten, due to turn-taking.

2.3.2 The effect of technology on collaborative style

Introducing technology to group work seems to have a general effect of increasing the amount of parallel work. Parallel work has also been observed to occur more when users work with shared datastructures such as calendars, matrices and lists as opposed to text (Whittaker et al., 1991). In other cases as well, the use of computer-based media to keep track and to have a permanent record has been shown to increase the tendency for groups to work in parallel (Stefik et al., 1987; Whittaker et al., 1991). Observing some differences between non-technological and technological communication media could provide some clues for this change in behavior. We can differentiate non-technological media, such as speaking or writing on a chalkboard, from technologies such as an electronic whiteboard or workstations where one can store documents. We can discuss these differences along two dimensions. First, the (intermediary) products of the non-technology media are more temporary or fleeting. Second, with these non-technology media a group must use more of its own resources to store the information, such as relying on the individuals' memory, personal note-taking, or audio taping. Perhaps the reduced load on memory enables group members using technological media to work independently since they don't need to rely on the group memory so much to reconstruct the information.

The physical distribution of users can also interact with the technological environment to influence the collaborative style. For example, Ellis et al. (1991) report observing more parallel work in distributed settings than in face-to-face sessions using GROVE, a cooperative text editor. The authors propose that this may occur due to distributed tasks usually having more group members, or perhaps there is less social pressure to work collectively. In a face-to-face environment one cannot easily excuse himself and leave the meeting or simply begin working alone. On the other hand, the close proximity of members in a shared drawing space also had an effect on collaboration influencing the group to take turns in order to prevent overlapping as well as to perform parallel work (Tang, 1991).

2.3.3 Costs and benefits of collaborative styles

Both parallel and collective work styles offer costs and benefits. One main advantage in working collectively in an electronic meeting room environment is that the continual reference to common objects maintains a shared focus and common context for the group (Stefik et al.,1987). However, users of Cognoter in the Colab often had difficulty keeping track of and locating the shared references, especially if the windows were different sizes and had different views (Tatar et al., 1991). With collective work, members must process the content of discussion as well as locate items on the screen which can involve a cost of information processing. It can also hinder the group to wait for its slowest member.

In parallel work, overhead costs include negotiating subtasks, keeping track of and being informed of what others are doing, coordinating finished work, the difficulty of making mid-course corrections (Tatar et al., 1991), losing information on how the task is accomplished, as well as increasing the potential for conflict (Stefik et al., 1987). However, despite these overhead costs, parallel work may offer benefits in performance, such as work groups experiencing less delays (Sproull and Kiesler, 1991), taking advantage of members' different expertise, and increasing productivity during brainstorming (Dennis et al. 1990; Hymes and Olson, 1992).

In collective work, the group needs protocols for coordinating its actions. In face-to-face groups, social information, such as gestures or nonverbal back channel responses, serve as a coordinating device for negotiating actions and maintaining collaborative behavioral norms. With remote collaboration using technology, this rich visual information can be lost or dramatically reduced. Technological protocols can be introduced to reduce coordination overhead for the group but they generally impose a turn-taking process (Ellis et al., 1991) and can limit the flexibility of the group in adopting different collaborative styles.

The overhead cost of communication and coordination may actually not be so high for parallel work groups since it may not be necessary for group members to communicate about all details of their work process when they work in parallel. In real-world collaboration without technology, dividing a task actually reduced the communication and coordination overhead because the members did not communicate the details of their individual work processes, only the result. Individuals were able to concentrate on their task components and only communicate what was necessary (Kraut et al.,1990).

Both collaborative styles complement each other and seem to be appropriate for different task phases (e.g. Kraut et al., 1988; Posner and Baecker, 1992; Olson et al., 1993). Therefore, it is important to design systems to provide adequate support for groups whose work styles may change with the situation. A collaborative system must be able to offer users a smooth transition to switch styles when needed. For example, groups must be able to easily access the results of parallel work in order to perform collective work. Although parallel work seems to provide task benefits (work completed faster, advantages in brainstorming), it also requires an increased effort for users to maintain awareness of other group members’ activities.

2.4 The experimental questions

The goal in the experiment was to investigate how a group’s product and process is affected by using hypermedia structures. We therefore investigated the effects of adding hypermedia functionality to what can be considered to be a "standard" electronic whiteboard functionality. As a first step, we chose to add a "simple" hypermedia node-link model and to use an electronic meeting room scenario for the experiment. This was motivated by our focus on supporting processes especially in the early phases of group problem solving, and we wanted to observe the effects on problem exploration, idea generation, information structuring, and adding/reducing information elements. This resulted in two conditions of the experiment: our implementation of a more or less standard electronic whiteboard functionality vs. this functionality in combination with additional hypermedia structure functionality.

According to the distinction made by Horton et al. (1991), the DOLPHIN system functions within a low structure environment which does not prescribe a particular collaborative style to groups. By testing the group within such an environment where constraints of choice are minimized we could observe what strategies groups might naturally adopt when working with shared hypermedia structures. The possibilities discussed for how hypermedia may affect the group product and process led us to test these effects, being guided by the following questions.

1. *Hypermedia document creation: for ad hoc groups as well as for experienced individuals?*

We began with the question of whether hypermedia use would be limited to people having previous expertise. While there is evidence that users with much experience can create very extensive hypermedia documents over time (Landow, 1989; Schuler et al., 1995), it is not yet clear whether users without much experience and situated in a face-to-face meeting can also portray their ideas in a hypermedia format within a reasonable training time. This might become even more critical if the composition of the group is more or less ad hoc as it might happen, e.g., in a business environment. We therefore asked, to what extent can a group of users, within a reasonable training time, actually use hypermedia structures of nodes, links, and nonlinearity to create logical structures in a problem solving situation?

2. *The effect on the group's product: does hypermedia facilitate the formation and elaboration of ideas and relationships?*

Our second question focused on how hypermedia might change the documents that a group creates. The hypermedia functionality of creating nodes and links enables users to form multiple connections between concepts, i.e., many-to-many mappings. We would expect that this functionality would facilitate the creation of a network structure which we define as a structure containing at least one concept with multiple connections to other concepts including cross-hierarchical links. We would expect to see this structure as opposed to a pure hierarchical structure with superordinates at the highest level and subordinate concepts at lower levels (no cross-hierarchical connections). On the other hand, we expect that using a standard electronic whiteboard where links are not provided as a system feature (although subjects can draw arrows by hand to express relationships) would not necessarily lead one to create a network structure.

Based on the property that hypermedia node structures can be used to elaborate ideas by using nested nodes (each containing content), we expect that this functionality would lead groups using hypermedia structures to elaborate their ideas more, creating subsequently deeper levels containing new information compared to groups using only a flat electronic drawing surface. There might also be a tradeoff between depth and breadth. The nested nodes, along with the link functionality should result in more ideas associated with a superordinate concept compared to a non-hypermedia information structure.

3. *The effect on the group strategy and cognitive process: how does hypermedia influence strategy and memory?*

We expect that the role of hypermedia in group problem solving will also be reflected in the strategies used and in the retention of information. Working with hypermedia may lead groups to fundamentally change their planning strategy when organizing ideas.

4. *The effect on the group process: how does hypermedia structuring influence a group's collaborative style?*
The possibility of decomposition by using hypermedia structures can be used for the distribution of tasks among group members and for parallel work and may even affect group participation. We expect that different decision making occurs when organizing information into a hypermedia structure and it may be reflected in differences in group participation, such as speaking or system use. We examined whether hypermedia influenced groups to adopt a particular collaborative style.

3 Method of the Experiment

3.1 Setting

In our electronic meeting room (the OCEAN Lab), groups were seated around a rectangular table with a large interactive, electronic whiteboard facing them at one end of the table. While other scenarios are possible (cf. Streitz et al., 1994, Streitz et al., 1997), in this experiment, DOLPHIN was used as a multi-user application shared between an electronic whiteboard and networked computers. Fig. 2 shows the scenario at the time of the experiment in 1994. The current version (1997) of the OCEAN Lab uses a different setting for DOLPHIN. It includes four computers with flat LCD displays mounted into the table. One can operate DOLPHIN using a keyboard and a mouse at the workstations and also with a cordless pen at the whiteboard.

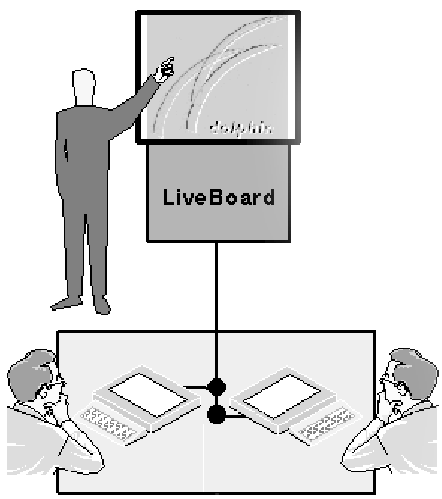

Figure 2. Physical setting of the electronic meeting room in the experiment

In this particular experiment, we used a SUN-based Xerox Liveboard (Elrod et al., 1992) as the hardware basis for the electronic whiteboard and two SUN Sparc-2 workstations with 17" color monitors mounted into the meeting table. One subject worked on the Liveboard, visible to everyone. Two subjects used the workstations while still being able to see and talk to the others. The public workspace on the Liveboard could also be presented on the workstations. For the public displays, we provide a relaxed WYSIWIS functionality (What You See Is What I See; Stefik et al., 1987), which means that users could scroll and resize their windows independently. The provision of workstations to two subjects allowed them to work in parallel in a "private workspace" (similar to taking individual notes on paper or looking at separate documents, etc.).

3.2 The evaluation infrastructure

During the experiment, various data was recorded. Screendumps were taken from every workstation and the Liveboard every 15 seconds by a snapshot program. The total view of the meeting room was videotaped. The final state of the meeting document was stored in DOLPHIN's cooperative hypermedia server. After the meeting, the document structure was reconstructed on a piece of paper and coded. An observer took notes during the meeting.

To be able to evaluate concurrent activities of multiple people using the computers provided in the meeting room, a specialized evaluation tool was developed. The screenshots from the workstations and the Liveboard plus the videotape were digitized into Quicktime movies on a Macintosh computer. The evaluation tool can present all four digital video streams concurrently on the Macintosh. One may choose to play one or all of the four streams available, resize the active video window, hear the sound in parallel to the video, and synchronize the screen content of all participants. By playing the video streams backwards or forwards, one can observe how the meetings developed. This tool was especially valuable for analyzing the groups' strategies during the meetings.

3.3 Design

In order to isolate the effects of using hypermedia structures, a between-subjects design was used. Subjects worked in groups of three, with eight groups per condition. Each group was assigned to one of two conditions which correspond to two different functionalities of the DOLPHIN system:

Non-hypermedia structures (N-Condition). Subjects were trained only in the non-hypermedia structures which provide the standard electronic whiteboard functionality. Here, DOLPHIN offers scribbles with the pen (and mouse) and typed text using a keyboard. These are objects which can be created, selected, moved, and deleted. Subjects could work only in one window, but could scroll to any area of the window and could scroll to different views independently from each other if they preferred.

Hypermedia structures (H-Condition). Subjects were trained both in the non-hypermedia structures of DOLPHIN (see above) as well as in the hypermedia structures of the system: for nodes and links the operations included creation, selection, opening (of nodes), deletion, and movement. Subjects were also shown how to create non-linear structures by copying and pasting nodes and links.

3.4 Subjects

The experiment took place during an eight-week period in the fall of 1994. A total of 48 subjects were recruited mainly from students of the Darmstadt University of Technology, while eight of them were recruited from the staff of GMD-IPSI, but were not part of the DOLPHIN development team. Most had received their university degree within the last three years. Subjects were assigned to conditions so that students and staff were randomly distributed among the groups. Subjects were not paid. No significant differences were found in subjects' computer experience, or age, between the conditions.

3.5 Procedure

Subjects first received a 40-minute training session on the system functionality. In both conditions subjects were shown examples of information structures on the Liveboard as possible information structures that one can create with the DOLPHIN system: e.g. a "network" structure. We tried to match the structures in both conditions as closely as possible so as not to create bias. In the N-Condition, the structure was created with handwritten words and drawn arrows; in the H-Condition, the same structure was created with nodes and links. Subjects were shown additional "matched" examples so that they could see that they could also use scribbles in the H-condition. At the end of the training, subjects worked on a 20-minute practice exercise. The practice exercise consisted of concepts associated with ordering food in a restaurant, and subjects were told to practice using the functionality to organize these ideas any way they like. In the H-condition, subjects practiced creating, opening, closing, and moving nodes, and using the pen to draw on the Liveboard. No systematic collaborative styles were observed during training.

After training, groups were instructed to spend the first 20 minutes in brainstorming and were read standard instructions for this task (Osborne, 1957). Subjects then were told to spend 40 minutes structuring and developing their ideas which they had generated during brainstorming. In both conditions, subjects were told that they could order their ideas into categories or graphical arrangements using the functionality that they had learned. It was emphasized that subjects in both conditions could structure their ideas any way they preferred, using their own strategy, and were not bound to any particular format. Subjects could choose to work in any style they wanted: together, separately, or varying their styles. Directly after the experiment, subjects were given a memory test and then filled out a questionnaire in order to survey users' work group experience, and opinions of the system, and to receive suggestions.

3.6 Task

It was decided that the task should meet the following requirements: 1) it should have the potential to allow subjects to generate a wide range of ideas, 2) its solution should be able to comfortably assume both the form of a non-hypermedia and hypermedia structure, 3) it should have a realistic purpose to increase motivation, and 4) it should be complex enough so that its solution could take advantage of DOLPHIN's capabilities in both experimental conditions.

Subjects were told to design a "library of the future" for the city of Darmstadt, Germany. They were to prepare a proposal for the city, which did not have to be in the form of a finished text document. The instructions stated that the citizens of Darmstadt would be the main users, cost was not to be a factor, financing should be left out, and ideas were to be generated without consideration of their feasibility, i.e. that current technology did not have to exist to realize their ideas.

3.7 Measures

The choice of our measures was guided by the hypotheses that we developed earlier addressing hypermedia use:

Relationship between ideas: The information structures produced by both conditions were categorized into network, pure hierarchical, and other.

Elaboration of superordinate concepts: We counted the number of levels (depth) in the documents in both conditions. Depth of the document is defined as the deepest pathway that one can reach beginning from the highest level of the document. In the H-condition it was the top node, and in the N-condition it was the highest level which in all cases was identifiable by some graphical marker such as an underline or enumeration. We also measured breadth, defined as the number of superordinate concepts at the highest level of the document.

Ideas connected to superordinate pathway: We counted the average number of ideas following intra- and inter-node links along each superordinate pathway.

Changes of ideas from brainstorming to information structuring. We counted how many ideas were generated during brainstorming and the net number of ideas the group ended up with after structuring the ideas, after additions and deletions. Each node was counted as one idea, and details in parentheses were included as part of the same idea.

Quality of documents. We asked two experts, a linguist, and a sociologist specializing in conversation patterns, to judge the documents along four dimensions. *Logic of local structure* was a measure of the logical relationship of elements within separate parts of the document (nodes in the hypermedia document, and clusters in the non-hypermedia document); *logic of global structure* was a measure of the logical relationship between separate parts of the document; *originality of solution* was a measure of how inventive/unusual the ideas were, the general approach to the solution, and the framework in which the ideas were presented; and *comprehensiveness* was a measure of how multifaceted the document was.

Group Strategy. Using the evaluation tool, groups' strategies were categorized as: primarily top-down (all/most of the superordinate concepts are set up initially; subordinate concepts are then categorized under them), depth-first (one superordinate concept was set up, developed with subordinate concepts, and then the process repeated for each superordinate), primarily bottom-up (subordinate concepts are first grouped, and then superordinate concepts assigned), and mixed (a mixture of the preceding strategies).

Memory of information. Directly after the experiment, subjects were given 15 minutes to try to reconstruct with paper and pencil the ideas and structure which they had just created.

The effect of hypermedia structuring on group processes was measured by:

Group process. The length of time in seconds of speaking and of system use by each group member was coded from the videotapes. In addition, the videotapes were observed to determine the collaborative style of each group.

Group satisfaction. Satisfaction was measured by the questionnaires.

3.8 Coding

For all dependent measures that we report, with the exception of group participation, coding was done by two separate coders, with percentage agreement exceeding 89% (this value occurred with quality of documents). For group participation we used one coder, since the coder had to simply observe start and end times of talking and typing, and we felt that this was an objective measure that could be done competently by one coder. Where there was any question about the data, the data was checked by a second coder.

4 Results

In this section we report the results in terms of the four questions that the experiment addressed: 1) use of hypermedia structures, 2) effect on the group's product, 3) effect on the group strategy, and 4) effect on the group process. A discussion of the results will be presented in section 5.

First we report how well subjects were able to work with hypermedia structures. With the exception of some minor clarification questions on tool use, no subjects had difficulty with using the system. In the H-Condition, during the information structuring phase nodes and inter-node links were used extensively. However, surprisingly, none of the H-groups used nodes or links during brainstorming. Despite the fact that a non-linear structure was not a requirement of the information structuring, 7 out of the 8 H-groups created a non-linear structure using inter-node links. One N-group also created a non-linear structure using hand-drawn arrows (see Fig. 3).

4.1 The groups' products

Relationship between ideas. It was expected that groups in the H-Condition would be more likely to create a network structure than a pure hierarchical structure. This was indeed the case: 6 out of 8 groups in the H-Condition created a network structure using intra-node links whereas 7 out of 8 groups in the N-Condition created a pure hierarchical structure (chi-square$_{(1)}$ = 6.36, p<.025).

Figure 3 shows a network structure created by an N-group. H-groups created a pure hierarchical structure in two cases. In the first case, they created a two-level (superordinate and subordinate) hierarchy, and in the second case, they created multiple levels (four). Figure 4 shows an example of a network structure created by an H-group. In most cases where N-groups created a pure hierarchical structure (6 cases), it was a two-level hierarchy. In only one case did they employ three levels .

Elaboration of ideas. It was expected that subjects in the H-Condition would be more likely to create structures having deeper levels. Indeed, H-groups did create structures having deeper levels and the result was highly significant (mean$_H$=3.5 levels, s.d.=.76, mean$_N$=2.25 levels, s.d.=.46 t(14)=3.99, p<.001). However, no significant difference was found in the breadth of the structures: (mean$_H$=4.5 concepts, s.d.=1.85, mean$_N$=5.4 concepts, s.d.=.74).

Ideas connected to superordinate pathway. It was also expected that more ideas would be associated with each superordinate concept in the H-Condition. Compared to the N-Condition, a strong trend showed that the information structures of groups in the H-Condition had, on the average, double the number of subordinate ideas connected to each superordinate concept, (mean$_H$=13.26 ideas, s.d.=10.08, mean$_N$=6.60 ideas, s.d.=2.70, t(14)=1.81, p<.09).

Changes of ideas from brainstorming to information structuring. More ideas were generated by the H-Condition than the N-Condition during brainstorming, but the difference was not significant. However, as a result of adding and deleting ideas over the duration of the experiment, the difference showed a trend of increasing in value (mean$_H$=48 ideas, s.d.=9.35, mean$_N$=38.25 ideas, s.d.=10.90, t(14)=1.92, p<.08).

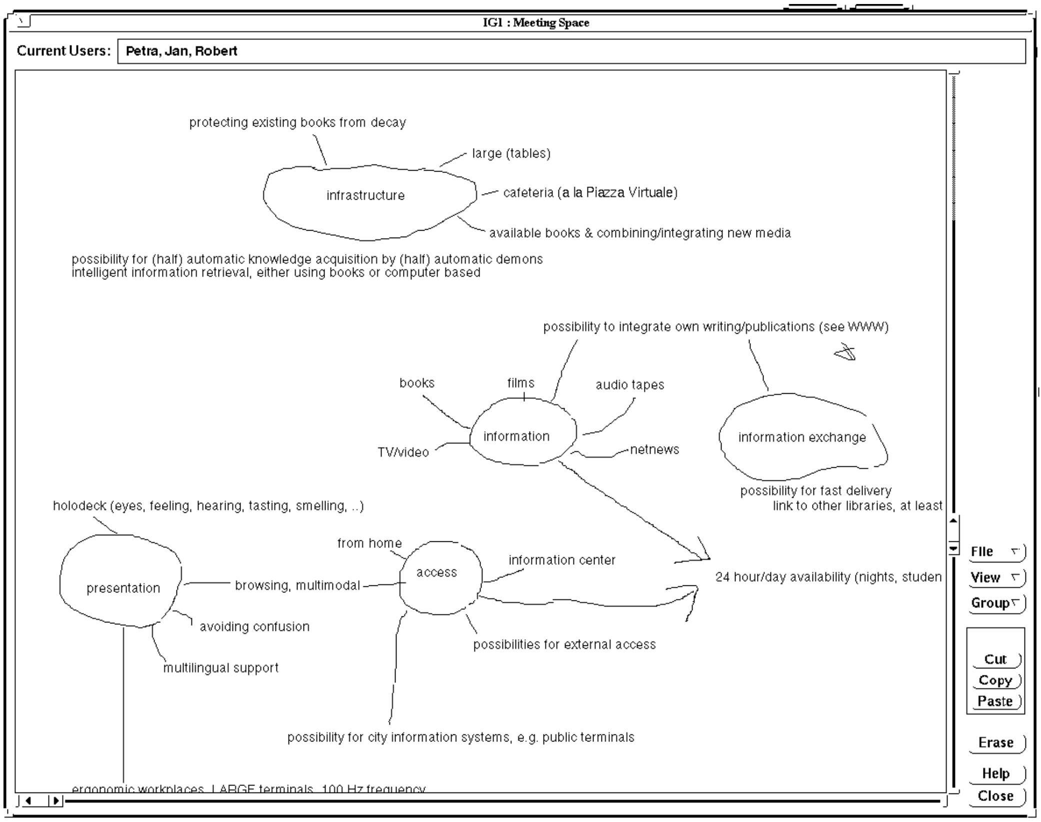

Figure 3: Example of a nonlinear structure created by an N-Group (translated from German)

Quality of documents. No difference was found by the expert coders in their judgment of the logic of local structures, logic of global structure, or comprehensiveness. However, although we provide a cautionary note since the results could be due to chance, the documents from the H-Condition were judged as being significantly more original in their solution ($t(14)=2.16$, $p<.05$).

In section 5.2, the results presented here of the groups' products will be discussed.

4.2 The group strategy

The type of strategy used in structuring the information was compared between conditions. Six out of eight groups in the H-Condition used a primarily top-down strategy. In contrast, groups in the N-Condition were found to be more likely to use a depth-first strategy (5 out of 8 groups). For the statistical test, a top-down vs. non top-down strategy was compared and the results show a significant difference in top-down strategy use ($\text{chi-square}_{(1)}=4$, $p<.05$).

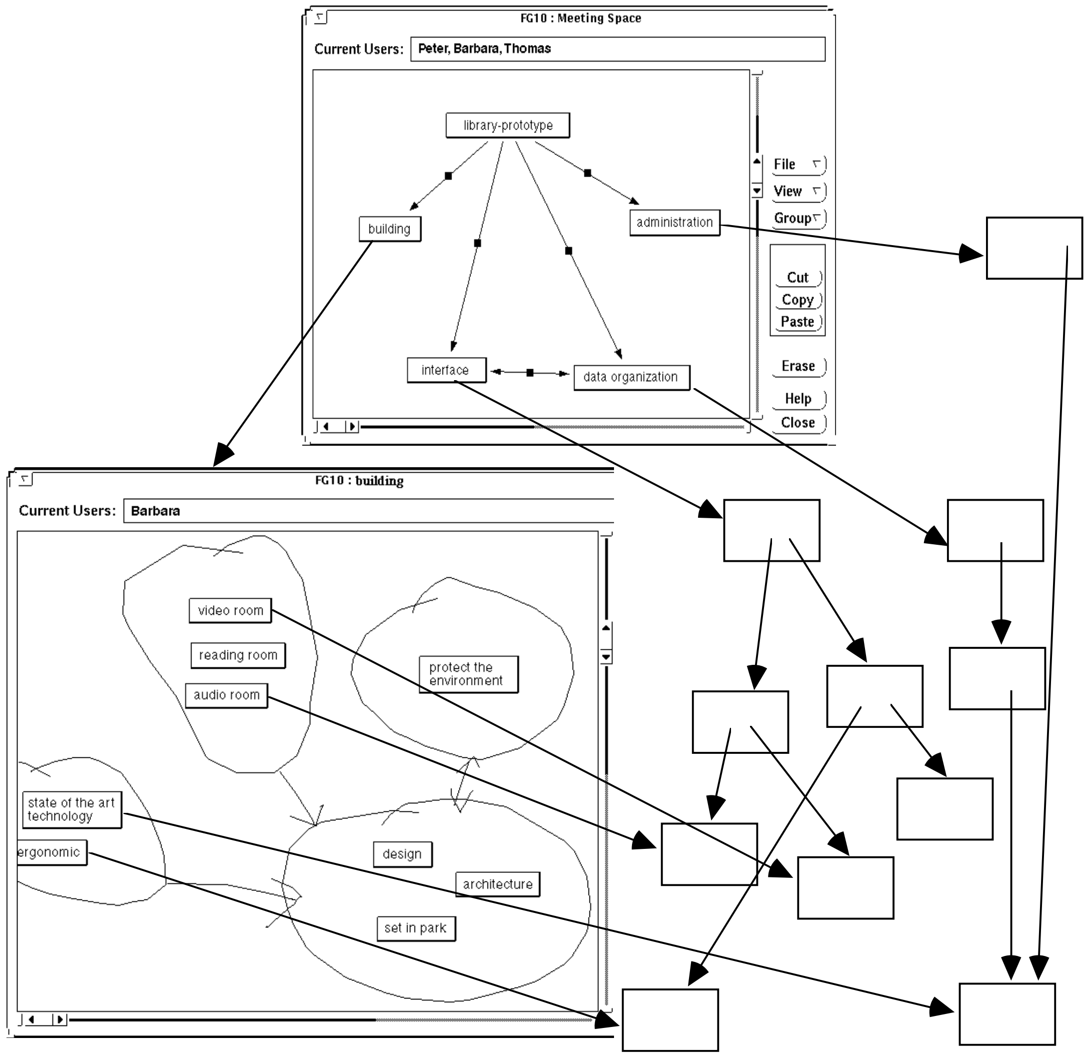

Figure 4: Results of an H-Group (translated from German). This Figure shows the content of the top node and one of the subordinate nodes. Due to lack of space we used a graphical representation of the rest of the document's structure. Here, empty rectangles denote nodes and fat arrows denote inter-node links.

Memory of information. We expected that a higher proportion of ideas in the H-Condition would be remembered than in the N-Condition. Surprisingly, the contrary turned out to be the case: subjects in the N-Condition actually remembered significantly more ideas ($mean_H$=49%, s.d.=7%, $mean_N$=64%, s.d.=7%, t(13)=4.11, p<.001). One outlying value from the H-Condition had to be eliminated. However, subjects in both conditions remembered the same proportion of superordinate concepts ($mean_H$=86%, s.d.=19%, $mean_N$=86%, s.d.=16%). We also expected that for the H-Condition, the memory of information would increase with the document level. The results confirmed this: subjects remembered 79% of the information on the first level of the document, 42% on all intermediate levels (combined for the analysis), and 31% on the bottom level of every branch in the structure. Some of the H-groups reconstructed the information structure as a network structure in the memory tests.

4.3 The group process: collaborative style

Temporal patterning and group operations

Since speaking and typing were coded in seconds, in order to be as precise as possible, we used only the first 30 minutes from all groups after brainstorming was finished and eliminated two groups due to minor system errors. Since we did not feel that the other results were affected by the errors, we elected to keep the results for the other analyses. No difference was found between conditions in the absolute amount of participation time for speaking or system use. However, we did find interesting results when speaking and system use were plotted in five minute intervals. Figure 5a shows the speaking distributions for the H-groups. Since some H-groups were faster or slower in their process, the values were standardized so that all dips are placed at the origin, in order to enable a comparison of the similarity of the forms of the distributions. Figure 5b shows the averages for both conditions, for speaking and system use. For the H-groups, speaking decreases and system use increases in one of the five minute intervals, generally occurring around 20 minutes into the structuring task. We found no such pattern with the N-groups. In contrast, the amount of speaking and system use in the N-groups generally follows a uniform distribution (Figure 5b). One sees the correspondence between the dip and peak in the group averages for the H-groups but no clear correspondence appears in the curves for the N-group averages.

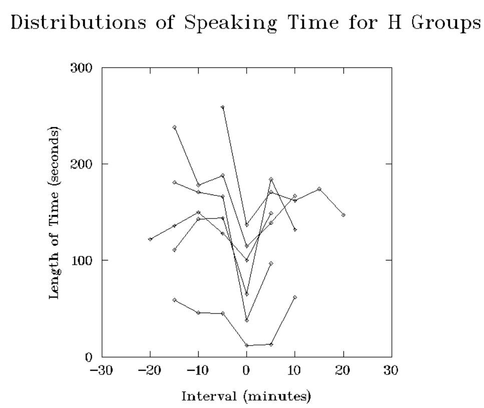

Figure 5a: Distribution of speaking times for H-Groups

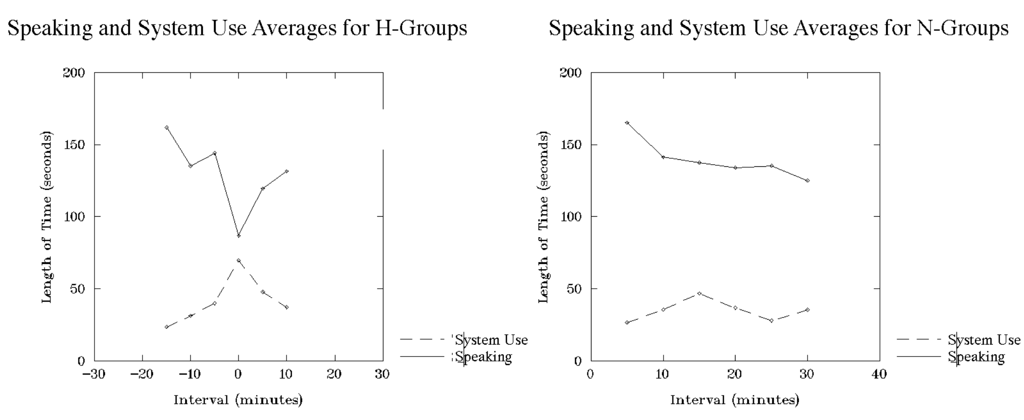

Figure 5b: Speaking and system use averages for both conditions

Group satisfaction.

N-groups reported higher satisfaction working in their groups than H-groups ($t(14)=2.22$, $p<.03$).

Collaborative style

General observations: hypermedia use, parallel work, and the division of labor. We observed that parallel work occurred in five of the H-groups (one of these groups worked only a short time in parallel, the rest of the time collectively) but not in any of the N-groups. By parallel work in our experiment, we refer to users working in separate areas of the document at the same time. In the H-Condition, this refers to opening up separate nodes and adding content to them. In the N-Condition, this refers to scrolling to a different area of the page and working there. This difference is highly significant (chi-square(1)=7.28, $p<.001$). Note that the scrolling by the N-groups could be detected by observing their screen views captured on the videotapes.

When the H-groups worked in parallel, they generally used a clear division of labor, in which subtasks were assigned to each member. We will discuss this process in more detail shortly. In contrast to the H-groups, only one N-group proposed a task division but none used a clear division of labor. With this exception, all N-groups began immediately working on the task together as a group without discussing other options of working.

The differences in temporal patterning were a clue to us that qualitative differences were occurring in the different group conditions. To explain our further observations, we now make a distinction between two types of proposals for group action which members can raise (Mark, 1991). A `content proposal refers to a proposal which concerns how a group can organize or generate content in terms of how to solve a task, such as We could enumerate the main points first, and then sort the ideas under these main points. In contrast, a process proposal refers to the group organization and how it can carry out operations. These proposals include, for example, (but are not limited to): assigning subtasks to members, suggesting that the members work in parallel, or proposing which roles members should adopt, such as leader.

We found a correspondence between the peak in system use, as described in the results, and the introduction of a process proposal in many of the H-groups. In five out of the eight H-groups, one member introduced a process proposal outlining how a group could carry out its operations. These operations involved organizing the brainstormed ideas into the superordinate categories, and in one case, establishing the superordinate structure. The proposals appear to have been a catalyst for the groups to take action, i.e. to begin carrying out its operation. Speaking declined and system use increased as a result of the group carrying out this plan. In other words, the peak in system use corresponds to the initial period when groups began developing the superordinates. After working for a period (roughly five minutes), groups generally began communicating again.

There was some variability however, concerning the proposals, such as whether a proposal was acted upon immediately or discussed, or the number of proposals concerning the action. Whereas we generally found a correspondence between the introduction of a proposal for parallel work and actual parallel work, there were some exceptions (see Table 1). Although group H7 made a process proposal, the group ended up working collectively. Another group, H6, did not introduce proposals but ended up working in parallel, each member adding content to separate nodes. No clear process proposal could be identified from the videotape. However, group H6 appeared to work more haphazardly than those groups who worked on the basis of a clear process proposal.

Condition	Process Proposals	Parallel Work	Top-Down Strategy
H-groups			
H1	Yes-P	Yes	Yes
H2	No	No	No
H3	Yes-P	Yes	Yes
H4	Yes-P	Yes	Yes
H5	Yes-P	Yes	Yes
H6	No	Yes	Yes
H7	Yes-P	No	Yes
H8	No	No	No
N-groups			
N1	No	No	No
N2	No	No	No
N3	Yes-C	No	No
N4	No	No	No
N5	No	No	Yes
N6	No	No	No
N7	No	No	Yes
N8	Yes-P	No	No

Table 1. Relationship between parallel work, introduction of a process proposal for parallel/collective work, and use of top-down strategy in H- and N-groups. Yes-P=proposal for parallel work, Yes-C=proposal for collective work

Group examples

The way that H-groups divided up their labor was actually quite similar. For example, in group H3, the group divided the task into two subtasks: working on a conventional library, and working on an electronic library. Their decision to divide up the labor seemed to be reached by them quite naturally:

> *Subject 2* (Workstation 1): We do it the following way: we have two categories: conventional and electronic. Then, we can take a second page where one has very many ideas which are networked under one another e.g. we have media. When one clicks on that, then comes the 'offerings' of the conventional library....How can we do that the best?...How to divide it up? We can copy it and then continue writing ...no, that won't work.
>
> *Subject 3* (workstation 2): That's exactly what I thought the whole time. It's irrational, that one works, one looks on, and the other has nothing to do.
>
> *Subject 2:* How can we solve this? We're all now on the same page...we could in principle write out the ideas that we discussed.....Now it's clear how it can work. One can do the conventional and the other the electronic...If we would be a project group, then we would delegate it in any case. One would do that, and the other would do the other.

The group then assigned group members to work on these areas: two members working on the electronic library, and one member working on the conventional library. Other H-groups divided the labor up also according to the superordinate concepts. In group H4, three of the four superordinate concepts were first set up: Environment, Topic Outline, and Media (later Time Structure was added). One member then proposed that they divide the work up according to each group member's main interest. The group then distributed these three superordinate concepts among the three group members, and each person developed their superordinate concept respectively. Similarly, in group H5, the group established three work areas for the task: Building, Offerings, and Organization. The group then decided that they would assign each person to one work area to develop ideas on this topic. Each member then created a node labeled with their name, opened a new node and elaborated this work area with the ideas they generated during brainstorming. In group H1, one member proposed that now that the group had established the superordinates, they should then proceed to divide up the subpoints, assign them to members, and further develop them. During the peak in system use, one member worked on the Liveboard to turn the superordinates into nodes and created the network structure by linking the superordinates. After a five-minute period, speaking increased due to discussing details of how the subordinates should be categorized. In group H7, it is interesting to note that the group proposed dividing the task into the two subtasks of Technical Departments and Reorganization and proposed assigning these two subtasks to subgroups. However, one member began introducing content proposals, and the group did not continue to assign the work areas to group members, but ended up working collectively on these two superordinate concepts.

As mentioned, none of the N-groups used a clear division of labor. In group N3, the person working at the Liveboard asked the group at the outset, “should we work together?” and the members agreed. Without further discussion or consideration of other modes of working, the group worked together creating an outline for their ideas. Group N8 provides an interesting example of a group having a sense that they should divide up their labor, but not knowing how to go about doing it.

> *Subject 2* (Workstation 1): We could divide the screen (referring to the Liveboard) up into four parts (pause). No, that’s not good.
>
> *Subject 1* (Liveboard): How can we best do this? (pause) Could we erase my part and leave your parts in (referring to the workstation screens)? (pause) No, that’s also not good, because everyone sees it. We could make three parts, or four parts, and then divide it up.
>
> *Subject 3* (Workstation 2): I have everything on my screen that we want to see.
>
> *Subject 2:* One can dictate, and the other write. Everything is visible.

This group then proceeded to work collectively without any further consideration of how the brainstormed ideas could be divided up into parts, or who should be assigned which parts. It was our observation that neither of the N-groups’ proposals were at the level of detail comparable with those of the H-groups. In section 5.4, we discuss the combined results of group strategy and collaborative style.

5 Discussion

Through the group protocols and quantitative measures, we can begin to identify a common thread that emerges to bind together some of the results of the H-groups. The hypermedia functionality of creating nodes, adding content, and linking nodes changes the way that groups work and the documents that they produce. This capability is in contrast to working on a two-dimensional "flat" working surface. We explain these findings by addressing the four questions raised earlier in Section 2.

5.1 Hypermedia use

We begin by discussing our results of how the groups worked with the hypermedia structures. We found that with a reasonable training period, a length of what might be appropriate for ad hoc groups in a business setting, all the groups were able to use the hypermedia structures without difficulty. With the exception of minor clarification questions on tool use, subjects were able to create nodes, add content to them, and create links, although we felt that links were underutilized. Despite the fact that a non-linear structure was not a requirement of the information structuring, six out of the eight H-groups created a non-linear structure using inter-node links. During training, subjects were shown an example of a non-linear structure as one possible way that they can structure their data. The fact that six of the groups used non-linearity after being shown only one example demonstrates to us that groups were able to comprehend the concept of non-linearity fairly quickly. This result provides optimism for us that with only a small amount of training, inexperienced users can create hypermedia documents. One N-group also created a non-linear structure with their available means of using hand-drawn arrows (see Figure 3). We feel it is also important to note that unlike many studies which examine group use of technology and that use facilitators, we did not explicitly teach a process. Our experiment was specifically designed for the groups to create their own process.

5.2 The groups' products

In this section, we discuss whether hypermedia facilitates the formation and elaboration of relationships, based on our evidence. At this point we would like to add a comment about our measures selected. For some of the measures, it is possible that there is an inherent bias to our selection. The operations involved in creating a network structure, deeper levels of a document, and connecting ideas to a superordinate pathway may simply be easier with hypermedia structures than by performing operations that involve working in an "unlimited 2D space", as in the N-condition. On the other hand, even if a bias exists, we chose these measures because we feel that they nonetheless are essential characteristics of a document structure and it is important to understand how they might be affected by hypermedia use.

H-groups produced documents with more relationships and more deeply elaborated concepts, which confirmed our expectation. These documents also contained more levels of depth and had a greater number of ideas associated with the superordinate concepts. The functionality of creating nodes and adding content in our view explains the deeper levels. We argue that the capability of creating nodes makes information easier to sort and easier to access. The gesture in DOLPHIN used to nest nodes within other nodes is a single movement of dragging one node onto another and is analogous to the simple motion of sorting papers into a bin. This motion makes it especially easy for users to create relationships as soon as they notice them. The ease of the hypermedia structuring was supported by comments from the groups, who often used metaphors of familiar operations to describe their actions. For example, adopting a metaphor of a book, one group proposed putting the superordinate ideas on the first "page" and then putting the subordinate ideas on a second "page". The greater number of ideas associated with the superordinate concepts follows as a result of the deeper levels since they generally led to wider branches at a lower level in the tree structure.

Groups who used hypermedia were more likely to create network, rather than pure hierarchical structures, and included more ideas in the structures as well. In addition, their ideas and problem approach were judged by experts to be more original. In a similar argument to the availability of the node structure, the prevalence of the network structures could have been due to the availability and ease of the link function, as opposed to the N-condition, where groups must hand-draw arrows in order to indicate a relationship. Links were also used to create unexpected relationships. One group used links between the superordinates to establish a "higher" and "lower" level of superordinates, although they were all contained on the same page.

The claim that the complexity of relationships produced using hypermedia may facilitate problem solving can also be explained by considering some of the models developed in cognitive psychology. Anderson (1983) claims that memory is organized as a network structure and that the spread of activation among associated concepts is an automatic process. Experimental work shows that some memory processes used in problem solving, such as retrieval of information (e.g. McKoon & Ratcliff, 1979), connectedness judgments (e.g. Reder & Anderson, 1980), and schema availability (e.g. Owens et al., 1979) are consistent with a network representation.

5.3 The group strategy

We now consider how hypermedia influenced the strategy that the group used, and individual memories of the information.

H-groups were more likely to use a top-down strategy whereas N-groups tended to develop one superordinate concept at a time. One explanation could be that H-groups knew that they had unlimited space in depth, which could be considered as a "third dimension" by creating and nesting nodes many levels down. Thus, they could plan the top level knowing they would have room for as many subordinate concepts as they needed. Some of the H-groups referred to creating and opening nodes as a space-keeping device which supports this notion. In one group, one member suggested that they put the brainstorming results in a new node, as a "notebook", to create a free page to plan the superordinate concepts. After the group had decided on the superordinate concepts, one member proposed that they create "subordinate" nodes.

A memory test revealed that N-groups recalled more of the total information than the H-groups, against our expectations. The N-groups could have recalled more information because they viewed all of the information a longer amount of time as compared to the H-groups who kept changing the information on their screen by opening and closing different nodes. However, although H-groups viewed the superordinate concepts for less time, they recalled them as well as the N-groups. This could be due to the fact that the lower level concepts were developed in the H-groups by others. One would naturally expect recall to be worse for information developed by someone else. This explanation would suggest that parallel work reduces the memory of other participants' contributions.

5.4 The group process and collaborative style

We discussed earlier that some forms of technology appear to have an effect on collaborative style by increasing the degree of parallel work. Our results are consistent with other studies, suggesting that hypermedia as well facilitates parallel work. In this section, we discuss our evaluation measure and how we discovered correspondences between parallel work and group strategy and product.

Process proposals. In order to gain an understanding of the process by which groups carried out their operations, we measured two aspects of the group process over time: speaking and system use. Temporal patterning in groups can provide much information as to how the group works since it can reflect sets of complex activities such as work flow coordination, matching time allocation to tasks, and synchronization of group members (McGrath, 1993). We found a qualitative difference in the temporal work pattern to exist between the two conditions. Whereas N-groups distributed their speaking and system use generally equally throughout the meeting, H-groups tended to work in a more uneven pattern.

A summary of a typical group process using the hypermedia functionality can be described as follows. During brainstorming, all groups worked collectively. During the planning session, H-groups worked collectively discussing what superordinate concepts (i.e. main concepts) to create. After the superordinate concepts were selected and indicated by being created as nodes, five of the groups divided up the task and worked in parallel, one group soon switching to collective work. After a period of elaborating the superordinates, the groups began to work collectively again, discussing new points, verifying what the other subgroup had done, or clarifying, e.g. under which superordinate a particular idea would be more appropriate.

Task phase and collaborative style

Although it would have been more optimal for H-groups to work in parallel during brainstorming they did not. Interestingly enough, one N-group worked in parallel by dividing the Liveboard into three columns and produced a large number of ideas. Hymes and Olson (1992) observed the effects of parallel work in brainstorming groups by specifically instructing groups to work in parallel. In contrast, our results showed that without specific instruction, groups working with DOLPHIN in a low-structure environment brainstormed collectively. In a future experiment, it would be interesting to compare parallel brainstorming of groups with and without hypermedia functionality.

In the planning phase, the collective planning of H-groups using DOLPHIN was similar to other results found with real-world planning (Kraut et al., 1990; Posner and Baecker, 1992). However it is unclear why our results are different than those found in the Capture Lab (Horton et al. 1991), where groups engaged in little planning discussion for writing. The results could be due to different tasks, or to differences in technology.

Top-down strategy and collaborative style

We found a correlation between the use of a top-down strategy and parallel work (see Table 1). We interpret this correlation to mean that the identification of the superordinates right from the start made it clear for the group how to divide up the task, i.e. by assigning superordinate concepts to different group members to develop.

By coding the videotapes, we were able to identify the different task phases of the two conditions. Most of the H-groups had a distinct planning phase, followed by a distinct organizing and idea-development phase, followed again by a planning phase, where members communicated about their individual work. We observed only these three task phases in most of the H-groups and the switch between the task phases appeared to be quite deliberate. In contrast, the shift between planning and development in the N-groups occurred at a much higher frequency. Thus, we see that the frequency in which N-groups shifted between task phases was far higher than in the H-groups, who spent a longer time in the organizing and development phase. If there is an overhead cost for a group in shifting between phases then this would mean an advantage for the H-groups who had a lower frequency of phase shifting. In contrast, the strategy of the N-groups reflects their collaborative style: it was more fluid. Most of the N-groups set up a superordinate concept in a planning phase then developed it collectively. The group then planned a second concept, developed that, and so on. Thus, the groups using hypermedia structures used a very different collaborative style for the task than those groups who did not work with hypermedia structures.

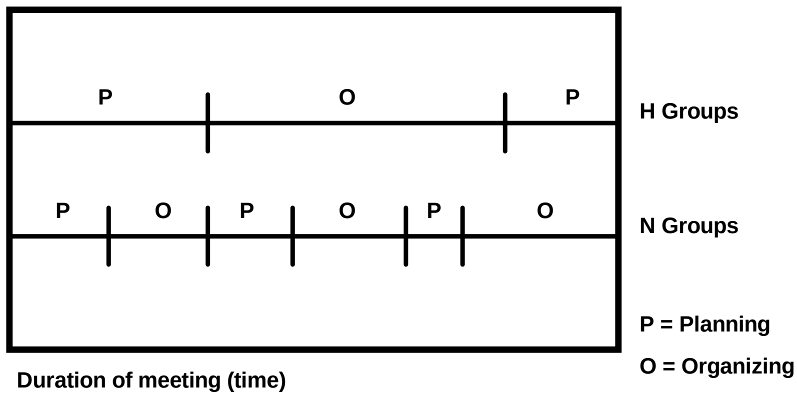

Figure 6: Hypothetical shift of collaboration style over the duration of the meeting for both conditions.

Fig. 6 shows this difference between the conditions in a hypothetical illustration of the frequency in which the collaborative styles shifted between task phases. It illustrates that the H-groups shifted more slowly between task phases thus spending more time in each phase.

The groups' products and collaborative style

Correlating the document measures with the process results reveals a positive relationship: groups who divided their labor produced documents with more levels of depth, with more multiple relationships between concepts, and consequently, with a greater number of ideas associated with the superordinate concepts. Why might such a correlation exist? In the experiment, groups may have accomplished more by dividing up their labor which enabled them to focus on individual concepts. This focus may have become manifest in the creation of more ideas per superordinate concept. We also did not find more inconsistencies among different people's work, even though they worked independently. No significant difference was found in the ratings that the expert raters gave to the local and global structures.

5.5 Evaluating other explanations

We have thus far presented our findings in terms of what we regard to be likely explanations. However, results can often be explained in multiple ways, and it is also sometimes difficult to rule out alternative explanations. Below we discuss two other possible explanations for our results and discuss the likelihood that they can account for our findings.

Explanation 1: Task demands during training

It is of course always possible that subjects behave a certain way because they believe that the experimenter expects them to act that way. We looked back to the training sessions to see if we inadvertently communicated expectations to subjects. First of all, both conditions were given the same training, which involved presenting subjects with concepts which have to do with ordering food in a restaurant. In both conditions the words were presented on the large screen of the Liveboard and subjects were instructed to organize the words any way they chose trying out the functionality that was demonstrated to them. Subjects in the H-Condition were not specifically instructed to open separate windows (corresponding to separate nodes). Although some subjects in these groups experimented with the functionality of opening and closing nodes, they did not split up their labor during the training sessions. It is our belief that the subjects did not make a connection between opening new nodes and splitting up their labor during training. However, the most convincing evidence of all to us was that during brainstorming (the phase directly following training) no groups in the H-Condition worked in parallel. If the groups believed that it was expected of them to work in parallel, then we should have seen division of labor during brainstorming, which actually is optimal for groups. It was also the case that no H-Groups used nodes and links during brainstorming, which also convinces us that during training we did not communicate expectations that they must use the hypermedia structures. The fact that these two things did not occur convinces us that we can rule out differences in the training as an explanation for the division of labor.

Explanation 2: Multiple windows

Since groups in the H-Condition were able to work in multiple windows, this could explain why the groups worked in parallel when organizing and elaborating their information. Similar results of groups working in parallel in the organizing and development phase were found with systems offering multiple windows (Olson et al., 1993; Horton et al., 1991). In the N-Condition, each group member could use a different view, but each view was of a different region of the same "unlimited" scrollable 2-dimensional large page. It is possible that it is not due to the specific hypermedia functionality, but rather the availability of opening multiple windows that triggered the H-groups to work in parallel. Without further experimental evidence, we cannot rule out this possibility.

However, some of the results lead us to believe that the explanation is more complex than this. First of all, multiple windows were available during brainstorming, but H-groups did not work in parallel. It was only after the H-groups formed the superordinate concepts using hypermedia nodes that five of the H-groups introduced a proposal to divide up their labor. We propose that the clear separation of concepts into nodes served as a device for clarifying how groups could separate the task, i.e. by assigning superordinates to group members. However, the separation of concepts alone was also not enough of an explanation, as evidenced by the one N-group described in section 4.3 who separated their superordinates but then could not find a suitable proposal for dividing up the labor. Although for this group it was clear how the superordinate concepts could be assigned, the means for doing it was not and the group ended up working collectively. It is our view that the N-condition lacked a clear way of separating and highlighting ideas (although in this one case, the group did notice a separation) and that it also lacked a clear method for working in parallel. Although N-groups could work in multiple regions of the same large window, we believe that it is cumbersome to scroll to a separate working area. The fact that five H-groups compared to one N-group proposed working in parallel suggests to us that H-groups noticed the clear separation of ideas in nodes and this may have “triggered” the proposals to divide up the work. For these groups, during their information structuring phase it became clear to them how they could work in parallel. Both facts together suggest to us that a semantic as well as procedural aspect is involved in the explanation. In the next section, we present what we believe is a more comprehensive explanation of the results and discuss the consequences.

6 Hypermedia and the division of labor

Up to now, we have found evidence showing that when using hypermedia structures to produce documents, as opposed to traditional non-linear methods, groups produce a different product, use a different strategy, and tend also to use a different process. In this section, we present a model which we believe can explain our findings in a parsimonious yet comprehensive way.

In order to explain how the use of hypermedia in the DOLPHIN system affects the early phases of group work, we propose the following model representing the relationship between the different processes as shown in Figure 7.

We distinguish between two major sets of activities: the problem solving activity and the writing activity. Problem solving refers here to those cognitive processes involved in planning the group's work and selecting a strategy, especially for the brainstorming and categorization part of the overall task. Writing refers to the actual creation of content, such as ideas, and to the organization of information into an overall structure of what is in the end a hypermedia document. This includes defining categories and assigning ideas to them or creating ideas for previously defined categories. Furthermore, it refers to creating relationships between information elements, thus creating a structured document. Assuming this distinction, we propose that the provision of hypermedia functionality in our experiment via DOLPHIN affects these processes in the following way.

The availability of hypermedia functionality enables the group members to create nodes, to add content to the nodes, and to create links between the nodes. Nodes function in two roles. First, nodes can "trigger" people to think about them as containers which can hold items. This provides a procedural device to put things into boxes so that they can be viewed in a separate window. Second, and more importantly, nodes serve as "semantic" triggers via their name or label. This can be more or less telling or precise but, in any case, it provides additional means for the structuring task which is part of the writing activity as well as of the problem solving activity.

This facilitates a writing style which will result in a highly modular document structure as the final product of the group. Modularity in turn suggests dividing the work between the group members, creating modules in parallel, and integrating them at a later stage. On the other hand, the availability of hypermedia functionality during the planning stage appears to bias the group toward using a top-down strategy for organizing and structuring their ideas. The use of this problem solving style involves choosing the superordinate concepts at the beginning of the task. Adopting this strategy at the outset of the task provides a natural way for groups to divide the work, namely, by assigning the superordinate concepts to group members.

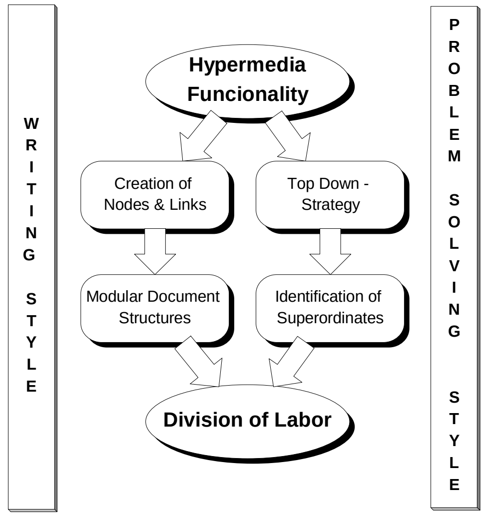

Figure 7: Hypermedia and the Division of Labor

Thus, we have two effects resulting from hypermedia functionality which amplify each other. The node structures provide a clear, visually presented separation of the overall workspace where it is easy to identify and access other members' workspaces. On the other hand, these modular node structures are ideal candidates to represent the superordinate categories. The combination of these effects provides a clear way to divide up the work and work in parallel. There is also a return effect. The ease of sorting and accessing information via the node structures, along with the availability of links, in combination with the top-down strategy has an impact on the characteristics of the document, resulting in a very modular document.

We found some evidence which suggests that groups who divided up the labor tended to produce more in the limited time they had to work on the task (Mark et al., 1995). In a longer task, we expect the variability associated with the factors of this model to become reduced. Although only five of the groups divided up their labor in this experiment (plus one group who proposed to do so), we expect that in a task longer than one hour, a higher percentage of groups would divide up the labor for efficiency. This assertion can be tested in future work.

7 Conclusions

Using hypermedia affects the way that groups work, as we found in the early phases of problem-solving. Groups using hypermedia functionality were more likely to submit proposals on how they could divide up their labor and were also more likely to work in this style. The ease of managing workspaces in terms of division, access, relating work, and sorting information can perhaps explain why hypermedia groups favored dividing up their labor. In contrast, the extra overhead involved with managing separate workspaces might explain why non-hypermedia groups preferred working collectively on the same page. Also associated with hypermedia use is a top-down strategy in which superordinates are chosen at the beginning of the task and which can provide a clear separation of the task. The division of labor in hypermedia groups correlates with the group product in which documents with more relationships between concepts and with deeper elaborations were produced.

7.1 Implications of hypermedia for group work

Process and product differences found with using hypermedia structures can have important implications for group work in general. We discuss these implications below.

7.1.1 The appropriateness of hypermedia

It is our claim that the many-on-many mappings made easier by hypermedia functionality facilitate the formation of multiple problem representations. The intuition that "more is better" may not hold when a problem requires only one optimal solution. However, viewing a problem from a different perspective, or discovering a new relationship between concepts can help overcome what is known as perceptual set, the tendency to encode a problem from one perspective. Overcoming fixed representations can be especially effective when trying to solve real-world problems, such as policy decisions, which are not so clear-cut. Many problem-solving techniques have been developed and are currently in use to help managers break free from remaining fixed on one standard perspective of a problem (e.g. see Bennis, 1989).

Galegher (1990) argues that the success of information systems in supporting group work depends on a thorough understanding of the tasks and the organizational context in which the system will be used. Following this line of reasoning, we argue that it is essential for the success of a cooperative hypermedia system such as DOLPHIN that we gain an understanding of for what tasks hypermedia would be most advantageous. In this paper we focused on the use of the system in the early phases of an ill-defined problem. With such an ill-defined problem, the opportunity for alternative views as facilitated by multiple relationships can be useful. However, we can imagine problems where hypermedia may not be appropriate. There is always the difficulty of transforming a non-linear text back into a linear format, and when the project goal is to produce a concise linear document, such as a memo, then perhaps the conversion cost when planning with hypermedia might be too high. A taxonomy of collaborative hypermedia tasks needs to be identified in order to gain a deeper understanding of which contexts of group work hypermedia would benefit.

7.1.2 Group strategy, memory, and navigation

In a collaborative hypermedia document, we would expect that navigation problems which exist for individual authors would be intensified for the group. First, although there is some evidence to suggest that memory of spatial location in an electronic document may be less important than memory of objects (Jones and Dumais, 1986), memory nevertheless seems to play a role in constructing a mental map of the document. In our study we found that H-groups had a poorer memory of the information than the N-groups. This could be explained by the tendency of H-groups to divide up their labor and thus only view part of the document. If division of labor occurs in groups working with hypermedia, as we suggest, then the navigation through another member's workspace would be difficult due to the unfamiliarity of the material. However, since we did find that H-groups had a good memory of the superordinate concepts (although they were only a few in each case), it could be that H-groups might be able to retain a good global view of the document, although not the details. The finding that H-groups tended to organize their information using a top-down approach could be significant in that groups may find it more natural to access material from a navigation system which is organized in a top-down manner.

Another source of difficulty is in the quantity and complexity of information that we would expect groups to produce. Hypermedia enables users to create multiple connections between the concepts. With collaboration in a long range task, the complexity of structures produced could make the navigation problem quite unmanageable for each member. On the other hand, there are advantages for group recall, since members can pool their memory resources (Yuker, 1955); thus, a collective memory could be an important aid in navigation.

7.1.3 Group process: flexibility and collaborative style

Improved efficiency for the group?

Keeping in mind the costs and benefits of parallel work, one possible benefit of dividing a task is that group members can work in parallel, sometimes accomplishing more. However, it is also important to keep in mind that there is a time expense in communicating about parallel work and in combining the results. Due to the limited amount of time that the group had to structure and develop its ideas (40 minutes), dividing up the task in this experiment was an efficient means for the groups to maximize their contributions. And in fact, those groups who divided up the task and worked on subtasks did produce more.

In contrast, it should be recalled that all the N-groups worked together collectively, most without considering a plan for dividing up the labor. It may in fact have been the case that the best working method for N-groups to use *was* to work collectively. The cost for the N-groups to work in parallel would have been to scroll to separate workspaces, not only to claim a personal workspace, but also group members would need to access others' workspaces to see the results and define relationships. Within the time constraints of the experiment, the coordination costs may have been too high if N-groups were to have worked in different areas of the same flat surface.

When discussing efficiency, it is also important to consider the management of interdependencies of a task such as clarifying redundancies, and combining inputs. However, in this experiment, the shortness of the task limited the amount of interdependent relationships that the group had to deal with.

Task division and communication

Although the division of labor can potentially create the burden of increased communication and management (Galegher, 1990), these burdens need not necessarily be unmanageable. Kraut et al. (1990) found that when labor is divided, communication usually takes place in transition points for new tasks. Such communication often involves encapsulated summaries which is less of a cognitive load than transmitting small details of the project. This communication can also provide useful updates at points in the project where, once group members are finished with one task, they are ready to review their work and move on. When a hypermedia document becomes large, then information summaries at transition points are especially important in order to retain a global view of the project without becoming overwhelmed by details.

As in any group task, one drawback of dividing labor is the potential for redundant work. When members are working in separate areas of a document, there is always the possibility and disadvantage for the group to overlap work. To avoid this danger, a group needs to establish regular communication during the task. Such communication can be coordinated or triggered via group awareness features. Although in our task little redundancy was found, it should be pointed out that the groups created a relatively small hypermedia document. Without further testing, it is difficult to predict the scalability of our results to a large hypermedia document. Despite the clear separation of superordinates among the members, subordinate ideas could certainly overlap, e.g. “information access from home” may be categorized under “technical services” or “loan services”. However, the advantage of hypermedia is that multiple links between concepts can be formed. Establishing a communication channel, perhaps with the assistance of an awareness component in the system, is integral when working with large hypermedia documents in order to ensure that such links are formed.

Reduction of social loafing?

The tendency of group members to work less in a group than they would individually is known as social loafing (Williams et al., 1981). In the case of social loafing, the additive contribution of individuals in a group context does not meet the potential of the sum of the individual members' efforts. Group members may reduce their efforts for a variety of reasons, such as diffusion of responsibility (Latane & Darley, 1970) or the free-rider effect (Kerr, 1983). The group suffers a productivity loss, failing to reach its maximum potential. It is our speculation that in the information structuring phase of a task, if the task is divided up, social loafing may be reduced. In some instances, social loafing is minimized in groups. First, if the work distribution is clear, then a member's contribution will be identified, and identification of individual work can minimize loafing (Williams et al., 1981), especially if the contribution is written and then accessible by all group members. In this case the public nature of the contribution will make it very identifiable, as in a public record. The free-rider effect could also be reduced when members believe that others in the group are working as hard as they are, which should occur with task division and assigning specific tasks to members (Kerr, 1983).

7.2 Limitations of the experimental approach and generalization

Working in an experimental approach always carries with it limitations, mainly that we fail to understand how a task is worked on in an actual work context of a group. Another main problem with an experimental approach is that the results cannot be easily generalized to a broad range of contexts. However, when studying the effects of electronic meeting room technologies, it is important to consider that when groups enter such an environment, they are taken out of their normal working context, and are brought into a strange, often hi-tech environment, usually quite different than their familiar surroundings. It is also often the case that electronic meeting room technologies are used for ad hoc business group purposes, that is, for particular tasks to be performed such as decision-making, or document-writing by specially formed groups such as task forces. Therefore, in this sense, our experimental approach is not so far off the mark to be generalized to such types of groups. We chose the experimental approach specifically to conduct basic research into the effects of hypermedia, in order to isolate the effects of hypermedia structures from nonhypermedia functionality. For this reason, we feel that our approach was valid. What we lose in generalization by working on an abstract level, we gain in a basic understanding of the effects of hypermedia. These basic research results serve as a basis for new research to build upon. It follows that for subsequent research, real-world groups working on their own tasks could provide a different and more specific viewpoint about hypermedia use.

7.3 Future work

In future work, we plan to investigate how hypermedia can affect later phases of problem solving, as well as medium range and longterm collaboration, and public vs. private work. The role of hypermedia as a "medium" for supporting collaborative work, such as in coordination, communication, and record-keeping is planned as a future direction. We would also like to test the generalizability of our results with ad hoc groups by working with "real-world" groups. We also plan to investigate new hypermedia features of the more recent versions of DOLPHIN which include navigation support, awareness, and history lists. In terms of evaluation, we would next like to observe other variables playing a role for designing these meeting environments as we did with respect to the availability of personal and public information devices (Streitz et al, 1997). We plan to continue this line of research by designing different configurations of what we call “roomware”. While these evaluations are concerned with local meetings in one room, we have extended the application scenarios to distributed settings. We now also support so-called "virtual meetings" by coupling remote meeting rooms and external experts via ATM-based networks (Johannsen et al., 1996). It will be another challenge to evaluate these more complex settings.

Acknowledgments

A number of people played important roles in making this study possible. The authors would like to thank the following people for their invaluable help and support: Jörg Geißler, Petra Rexroth, Chris Neuwirth, Matthias Will, Elke Teich, Adelheit Stein, Wiebke Möhr, Ajit Bapat, and finally Lutz Kirchner, Christian Schuckmann, and Jan Schümmer. Many thanks to Jennifer Beck-Wilson for proof-reading the document.

References

Anderson, J. (1983). *The Architecture of Cognition*. Cambridge: Harvard University Press.

Bennis, W. (1989). *On Becoming a Leader*. Reading, MA: Addison-Wesley.

Baecker, R. (Ed.) (1993). *Readings in Groupware and Computer-supported Cooperative Work*. San Mateo: Morgan Kaufmann.

Bapat, A., Geißler, J., Hicks, D., Streitz, N., Tietze, D. (1996). From Electronic whiteboards to distributed meetings: Extending the scope of DOLPHIN. *Video Proceedings of the ACM Conference on CSCW (CSCW ‘96)*, Boston, November, 1996.

Chang, E. (1989). Protocols for Group Coordination in Participant Systems. In M.M. Taylor, F. Neel, and D. G. Bouwhuis (Eds.). *The Structure of Multimodal Dialogue*. Elsevier Science (North Holland).

Craik, F., Lockhart, R. (1972). Levels of processing: a framework for memory research. *Journal of Verbal Learning and Verbal Behavior, 11*, pp. 671-684.

Dennis, A.R., Valacich, J.S., and Nunamaker, J.F. (1990) An experimental investigation of the effects of group size in an electronic meeting environment. *IEEE Transactions on Systems, Man, and Cybernetics*, 20, pp. 1049-1057.

Elrod, S. et al. (1992). Liveboard: a large interactive display supporting group meetings, presentations and remote collaboration. *Proceedings of CHI'92*, Monterey, May 3-7, 1992, ACM Press, pp. 599-607.

Ellis, C.A., Gibbs, S. J., and Rein, G. L. (1991). Groupware: Some issues and experiences. *Communications of the ACM*, 34 (1), pp. 39-58.

Galegher, J. (1990). Intellectual teamwork and information technology: the role of information systems in collaborative intellectual work. In J. S. Carroll (Ed.), *Applied Social Psychology and Organizational Settings.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Gick, M., Holyoak, K. (1983). Schema induction and analogical transfer. *Cognitive Psychology,* 15, pp. 1-38.

Grønbaek, K., Hem, J., Madsen, O., Sloth, L. (1994) Cooperative hypermedia systems: A Dexter-based architecture. *Communications of the ACM, 37 (2),* pp. 64-74.

Haake, J., Neuwirth, C., Streitz, N. (1994). Coexistence and transformation of informal and formal structures: Requirements for more flexible hypermedia systems. *Proceedings of ACM European Conference on Hypermedia Technology (ECHT'94).* Edinburgh, pp. 1-12.

Horton, M. , Rogers, P., Austin, L., Brimm, D., McCormick, M. (1991) The impact of face-to-face collaborative technology on group writing. In J. Nunamaker, Jr. and R. S. Sprague, Jr., *Proceedings of the Twenty-Fourth Annual Hawaii International Conference on System Sciences,* Kauai, Hawaii, Jan. 7-11, 1991. IEEE Computer Society Press, pp. 298-307.

Hymes, C. M. and Olson, G. M. (1992). Unblocking brainstorming through the use of a simple group editor. In J. Turner and R. Kraut, (Eds.), *Proceedings of CSCW '92,* Toronto, Oct. 31-Nov. 4, 1992, ACM Press, pp. 99-106.

Ishii, H., Kobayashi, M. Grudin, J. (1993). Integration of interpersonal space and shared workspace. In T. Malone & N. Streitz (Eds.), *ACM Transactions on Information Systems, Special issue on CSCW.* 11(4), pp. 349-375.

Johannsen, A., Haake, J., Streitz, N. (1996). *Telecollaboration in Virtual Organizations: The Role of Ubiquitous Meeting Systems.* GMD Arbeitsbericht No. 974.

Jones, W. P., Dumais, D. T. (1986). The spatial metaphor for user interfaces: Experimental tests of reference by location versus name. *ACM Trans. Office Inf. Syst.*, 4 (1), pp. 42-63.

Kerr, N. L. (1983). Motivation losses in small groups: A social dilemma analysis. *Journal of Personality and Social Psychology*, 45, pp. 819-828.

Kraut, R.E., Egido, J., & Galegher, J. (1990). Patterns of contact and communication in scientific research collaboration. In J. Galegher, R.E. Kraut, & C. Egido (Eds.), *Intellectual teamwork: The social and technological bases of cooperative work.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Landow, G. (1989). Hypertext in literary education, criticism, and scholarship. *Computers and the Humanities*, 23, pp. 173-198.

Latane, B. & Darley, J. M. (1970). *The Unresponsive Bystanders: Why Doesn't He Help.* New York: Appleton-Century-Crofts.

Mantei, M. (1988). Capturing the Capture Lab Concepts: A case study in the design of computer-supported meeting rooms. *Proceedings of the Conference on CSCW (CSCW '88),* Portland, September 26-29, 1988, ACM Press, pp. 257-268.

Mark, G. (1992). *The generation of idea proposals: an illustration in a negotiation setting.* EDS Center for Advanced Research Technical Report, Cambridge, MA. Available from author.

Mark, G., Haake, J., and Streitz, N. (1995). The Use of Hypermedia in Group Problem Solving: An Evaluation of the DOLPHIN Electronic Meeting Room Environment. In Marmolin, H., Sundblad, Y., Schmidt, K. (Eds.) *Proceedings of the Fourth European Conference on CSCW,* Stockholm, Sept. 10-14, Kluwer Academic Publishers, pp. 197-213.

Mark, G., Haake, J., and Streitz, N. (1996). Hypermedia Structures and the Division of Labor in Meeting Room Collaboration. In M. Ackerman (Ed.), *Proceedings of the ACM Conference on CSCW* (CSCW '96), Boston, ACM Press, Nov. 16-20, 1996, pp. 170-179.

Marmolin, H., Sundblad, Y., Pehrson, B. (1991). An analysis of design and collaboration in a distributed environment In L. Bannon, M. Robinson, and K. Schmidt (Eds.), In *ECSCW '91, Proceedings of the Second European Conference on CSCW,* Amsterdam, Kluwer Academic Publishers, pp. 147-162.

Mayer, R, and Greeno, J. (1975). Effects of meaningfulness and organization on problem solving and computability judgements. *Memory & Cognition, 3,* pp. 356-362.

McGrath, J. E. (1993). Time, interaction, and performance (TIP). In R. Baecker (Ed.), *Groupware and Computer-Supported Cooperative Work*, San Mateo: Morgan Kaufman, pp. 116-129.

McGuffin, L. & Olson, G. M. (1992). *ShrEdit: A Shared Electronic Workspace.* Technical Report No. 45, University of Michigan, Cognitive Sciences and Machine Intelligence Lab, 1992.

McKoon, G. & Ratcliff, R. (1979). Priming in episodic and semantic memory. *Journal of Verbal Learning and Verbal Behavior* 18, pp. 463-480.

McLeod, P. (1992). An assessment of the experimental literature on electronic support of group work. *Human Computer Interaction, 7*, pp. 257-280.

Newell, A. (1980). Reasoning, problem solving, and decision processes: The problem space as the fundamental category. In R. Nickerson (Ed.), *Attention and Performance VIII,* Erlbaum.

Nielsen, J. (1995). *Multimedia and Hypertext: The Internet and Beyond.* NY: Academic Press.

Nunamaker, J.F. et al. (1991). Electronic meeting systems to support group work. *Communications of the ACM*, *34(7*), pp. 40-61.

Olson, J., Olson, G., Storrosten,M., Carter, M. (1993). Groupwork close up: A comparison of the group design process with and without a simple group editor. In T. Malone & N. Streitz (Eds.), *ACM Transactions on Information Systems, Special issue on CSCW*, 11(4), pp. 321-348.

Osborn, A. F. (1957). *Applied Imagination*. New York: Charles Schribner’s Sons.

Owens, J., Bower, G. H., and Black, J. B. (1979). The "soap opera" effect in story recall. *Memory and Cognition,* 7, pp. 185-191.

Pedersen, E., McCall, K., Moran, T., Halasz, F. (1993). Tivoli: An electronic whiteboard for informal workgroup meetings, *Proc. of the InterCHI'93 Conference,* Amsterdam, pp. 391-398.

Posner, I. R. and Baecker, R. M. (1992). How People Write Together. *Proceedings of the Twenty-Fifth Hawaii International Conference on System Sciences, Vol. IV*, Kauai, Hawaii, January 7-10, 1992.

Reder, L. M. & Anderson, J. R. (1980). The role of elaboration in the comprehension and retention of prose. *Review of Educational Research* 50, pp. 5-53.

Rein , G. L., Ellis, C. A. (1989). The Nick experiment reinterpreted: Implications for developers and evaluators of groupware. *Office: Technology and People, 5(1)*, pp. 47-75.

Schein, E. H. (1990). *Organizational Culture and Leadership.* San Francisco: Jossey-Bass.

Schuler, W., Hannemann, J., Streitz, N. (Eds.) (1995). *Designing User Interfaces for Hypermedia.* Heidelberg: Springer.

Schwartz, S. (1971). Modes of representation and problem solving: Well evolved is half solved. *Journal of Experimental Psychology,* 91, pp. 347-350.

SMART (1997). http://smarttech.com/prodhome.htm

Sproull, L. and Kiesler, S. (1991) *Connections: New Ways of Working in the Networked Organization*. Cambridge: MIT Press.

Stefik, M. et al. (1987). Beyond the chalkboard: computer support for collaboration and problem solving in meetings. *Communications of the ACM*, *30(1)*, pp. 32-47.

Streitz, N. (1985). *Subjektive Wissensrepräsentationen als Determinanten kognitiver Prozesse.* (Subjective knowledge representations as determinants of cognitive processes). Doctoral Dissertation. Technical University of Aachen.

Streitz, N. (1994). Putting objects to work: Hypermedia as the subject matter and the medium for computer-supported cooperative work. In M. Tokoro & R. Pareschi (Eds.), *Object-Oriented Programming.* (ECOOP’94). Lecture Notes in Computer Science. Heidelberg: Springer, pp. 183-193.

Streitz, N., Geißler, J. Haake, J., Hol, J. (1994). DOLPHIN: Integrated meeting support across Liveboards, local and remote desktop environments. In R. Furuta and C. Neuwirth (Eds.), *Proceedings of the Conference on Computer-Supported Cooperative Work (CSCW’94),* Chapel Hill, Oct. 22-26, 1994, ACM Press, pp. 345-358.

Streitz, N., Haake, J., Hannemann, J., Lemke, A., Schuler, W., Schütt, H., Thüring, M. (1992). SEPIA: A cooperative hypermedia authoring environment. *Proceedings of the 4th ACM European Conference on Hypertext (ECHT'92),* Milan, pp. 11-22. (updated and revised version in: R. Rada (Ed.), Groupware and Authoring. London: Academic Press, 1996. pp. 241-264)

Streitz, N., Rexroth, P., Holmer, T. (1997). Does „roomware" matter ? Investigating the role of personal and public information devices and their combination in meeting room collaboration. In Prinz, W., Rodden, T., Hughes, J. and Schmidt, K. (Eds.) *Proceedings of ECSCW'97,* Lancaster, Sept. 7-11, 1997, Kluwer Academic Publishers, Dordrecht.

Tang, J. C. (1991). Findings from observational studies of collaborative work. *Int. J. Man-Machine Studies*, 34, pp. 143-160.

Tatar, D. G., Foster, G., and Bobrow, D. G. (1991). Design for conversation: lessons from Cognoter. *Int. J. Man-Machine Studies* 34, pp. 185-209.

Watson, R. T., DeSanctis, G., and Poole, M. S. (1988). Using a GDSS to Facilitate Group Consensus: Some Intended and Unintended Consequences. *MIS Quarterly* (12:3), September, 1988, pp. 463-478.

Whittaker, S., Brennan, S., and Clark, H. (1991). Coordinating activity: An analysis of interaction in CSCW. *Proceedings of CHI'91*, New Orleans, April 27-May 2, 1991, ACM Press, pp. 361-367.

Williams, K. B., Harkins, S., & Latane, B., (1981). Identifiability as a deterrent to social loafing: Two cheering experiments. *Journal of Personality and Social Psychology*, 40, 303-311.

Wingfield, A., Byrnes, D. (1981). *The psychology of human memory*. NY: Academic Press.

Yuker, H. E. (1955). Group atmosphere and memory. *Journal of Abnormal and Social Psychology*, 51, pp. 17-23.

Zigurs, I. and Kozar, K. (1994). An Exploratory Study of Roles in Computer-Supported Groups. *MIS Quarterly,* September, 1994, pp. 277-297.