KUNGL TEKNISKA HÖGSKOLAN
Royal Institute of Technology
Numerical Analysis and Computing Science

CID-69, KTH, Stockholm, Sweden 2000

Designing Storytelling Technologies to Encourage Collaboration Between Young Children

Steve Benford, Benjamin B. Bederson, Karl-Petter Åkesson, Victor Bayon, Allison Druin, Pär Hansson, Juan Pablo Hourcade, Rob Ingram, Helen Neale, Claire O'Malley, Kristian T. Simsarian, Danae Stanton, Yngve Sundblad & Gustav Taxén

CID
Centre for
User Oriented IT Design

Steve Benford, Benjamin B. Bederson, Karl-Petter Åkesson, Victor Bayon, Allison Druin, Pär Hansson, Juan Pablo Hourcade, Rob Ingram, Helen Neale, Claire O'Malley, Kristian T. Simsarian, Danae Stanton, Yngve Sundblad & Gustav Taxén

Designing Storytelling Technologies to Encourage Collaboration Between Young Children

Report number: CID-69

ISSN number: ISSN 1403-073X

Publication date: April 2000

Reports can be ordered from:

CID, Centre for User Oriented IT Design
Nada, Dept. Computing Science
KTH, Royal Institute of Technology
S-100 44 Stockhom, Sweden
telephone: + 46 8 790 91 00
fax: + 46 8 790 90 99
e-mail: cid@nada.kth.se
URL: http://www.nada.kth.se/cid/

**Paper accepted to CHI-2000, the Hague, the Netherlands,
1-6 April 2000.**

Designing Storytelling Technologies to Encourage Collaboration Between Young Children

KTH authors[2], SICS authors[4], Marland authors[3], Danae Stanton[1], Helen Neale[1], Sue Cobb[1], Victor Bayon[1], Rob Ingram[1], Steve Benford[1], Claire O Malley[1], John Wilson[1]

[1]Nottingham University
Nottingham, UK
+44 115 9515151
{firstname.lastname}
@nottingham.ac.uk

[2]Royal Institute of Technology, KTH
Stockholm, Sweden
{Phone}
{Generic email}

[3]The University of Maryland
USA
{Phone
Generic email}

[4]Swedish Institute of Computer Science, SICS
Stockholm, Sweden
{Phone}
{Generic Email}

ABSTRACT

We describe the iterative design of two collaborative storytelling technologies for young children, KidPAD and the Klump. We focus on the idea of designing interfaces to subtly encourage collaboration such that children are invited to discover the added benefits of working together. This has been motivated by our experiences of using early versions of our technologies in schools in Sweden and the UK. We describe how we have revised the technologies to encourage collaboration and to reflect a number of design suggestions made by the children themselves. We compare the approach of encouraging collaboration with other approaches to synchronizing shared interfaces.

Keywords

Children, Single Display Groupware (SDG), Computer Supported Cooperative Work (CSCW), Education, Computer Supported Collaborative Learning (CSCL)

INTRODUCTION

Collaboration is an important skill for young children to learn. Educational research has found that working in pairs or small groups can have beneficial effects on learning and development, particularly in early years and primary education [10, 14, 15]. Technology offers an opportunity to support and facilitate collaborative learning in many respects [1, 9]. The computer can provide a common frame of reference and be used to support the development of ideas between children. However, neither learning nor collaboration will occur simply because two students share the same computer [9]. Numerous factors must be addressed, not least of which is the learner-machine interface. Today s technology is designed to support either one individual at one computer, or one individual collaborating with another individual at a different computer. However, much if not most, classroom computer use involves pairs or small groups sharing the same computer, especially in primary or elementary schools. What we have come to call *shoulder-to-shoulder collaboration*, as distinct from distributed collaboration, is not well supported with today s interfaces.

In this paper, we explore the design of storytelling technology to help develop collaboration skills in children aged 5-7 years. This is a particularly interesting group to work with because previous research has shown significant changes in the ability to collaborate effectively within this age range [Wood et al., 1995]. Young children find it difficult to collaborate effectively. Informal observation of behaviour in our project has found that the youngest children (aged 4 and 5) have the most difficulty in working collaboratively and cannot work effectively at all in groups greater than 2.

We introduce an approach to the design of shared interfaces that involves subtly *encouraging* children to explore the possibilities of collaborating, rather than forcing them to do so. At the heart of this approach is the idea that children should be able to work independently if they wish, but could discover that there are added benefits to working together, for example by being able to create new graphics and effects for their stories. This approach lies somewhere between other approaches to the design of shared interfaces that either involve enabling the possibility of collaboration through multi-user access coupled with awareness of others actions [Gutwin 98] and concurrency mechanisms [Greenberg94] or enforcing collaboration by requiring that multiple users act together to drive the interface [e.g., Blaye et al., 1992].

The research described here has been carried out within the KidStory project, a collaboration between researchers, classroom teachers, and children (5-7 years old) from England, Sweden, and the United States. The goal of the project is to develop collaborative storytelling technologies for young children. The KidStory technologies are based on the approach of Single Display Groupware (SDG), where several children interact with a single display using multiple input devices, for example, two independent mice [Buxton 86][Bier 91][Inkpen 97][Stewart 98a][Stewart

98b][Stewart 99]. In its first phase, KidStory has worked with two pre-existing technologies, a shared drawing tool called KidPad [Druin 97] and a shared 3D environment called the Klump (an application of the DIVE collaborative virtual environment platform [Fahl n93]) both initially with one mouse and later with multiple mice. KidStory has also used the methods of cooperative inquiry [6], to involve children as technology design partners in an intergenerational and interdisciplinary design team. To accomplish this, a year-long series of technology design sessions were conducted in two schools with almost 100 children in the England and Sweden.

The following section describes the initial KidStory technologies. We then introduce the idea of designing interfaces to encourage collaboration and describe its use in the redesign of KidPad and the Klump.

THE INITIAL VERSIONS OF KIDPAD and THE KLUMP

We have been working with two collaborative storytelling technologies, KidPad and the Klump. Both enable two or more children to create and tell stories together, but differ in style, KidPad being derived from drawing and the Klump from sculpting or modeling. In the following we describe them as they were at the start of this research, before being extended to encourage collaboration.

KidPad

KidPad is a shared 2D drawing tool that incorporates a zooming interface. Children can bring their stories to life by zooming between drawing elements (see Figure 1). Zooming and spatial structure lie at the heart of KidPad, since they enable children to add narrative structure to their stories by dynamically moving between different parts of a drawing. The KidPad interface is designed around a series of graphical local tools that children pick up and apply using a mouse [Bederson 96]. The tools are:

Crayons — different coloured crayons can be used to create drawing elements;

Arrow —a selection tool to pick up and move objects.

Eraser — this can be used to delete drawing elements;

Magic wand —this is used to create zooms between different drawing elements. The child selects the drawing element to be the start of the zoom followed by the destination element and sees an arrow linking the two;

Hand — the hand is used to activate zooms when the story is being told. Selecting the start point of the zoom initiates an animated zoom to the end point.

Turn alive — this tool animates a story element by causing its outline to ripple, making it appear to be alive.

Bulletin Board — this tool enables children to save stories to a bulletin board.

Toolbox —this special tool is used to organized the other tools, and can be opened or closed.

KidPad is a Single Display Groupware system, which means that it supports several mice plugged into a single computer. Two or more children can independently grab and use different tools at the same time using their own mice. Any free tool can be picked up and the children see each other s cursors. As a result, this initial version of KidPad could be said to *enable* collaboration —the children could choose to work together or individually. Figure 2 shows an example of the KidPad interface.

KidPad is built on top of Jazz and MID, which are both open source Java toolkits. Jazz[1] supports Zoomable User Interfaces by creating a hierarchical scenegraph for 2D graphics, and MI supports multiple input devices for Java.

The creation of a story in KidPad, which involves creating links and zooming between picture/scenes or zooming deeper into the scene, allows the development of non-linear, complex structured stories. These story representations make salient the links between scenes and the overall structure of the story. The focus of the children s attention on these features of the story structure may allow new opportunities for learning, in a different way to the creation of a story using more traditional drawing packages or word-processing package. In lesson based storytelling sessions children may be limited by their writing skills. Young children sometimes have trouble writing a word in the first sentence of a story and this may hinder the childs progress, meaning that their understanding of stories and their creation may be limited. This does not mean that the written word is not important, but the use of KidStory technologies may take on a complementary role and perhaps provide a motivation for learning how to add words and sentences to a story.

The Klump

In contrast to the drawing based approach of KidPad, our second storytelling tool, the Klump is based on a modeling approach. the Klump is a collaborative 3D storytelling tool based around an amorphous 3D object (in fact, a textured deformable 3D polygon mesh) that can be stretched, textured and coloured and that makes sounds as it changes and is manipulated. Figure 3 shows an image of the Klump after it has been stretched and textured to create a character

[1] *Jazz is available at http://www.cs.umd.edu/hcil/jazz*

[2] MID is available at http://www.cs.umd.edu/hcil/mid

Figure 1: A sequence of views inKidPad as we zoom into a simple story (from left to right, and then top to bottom)

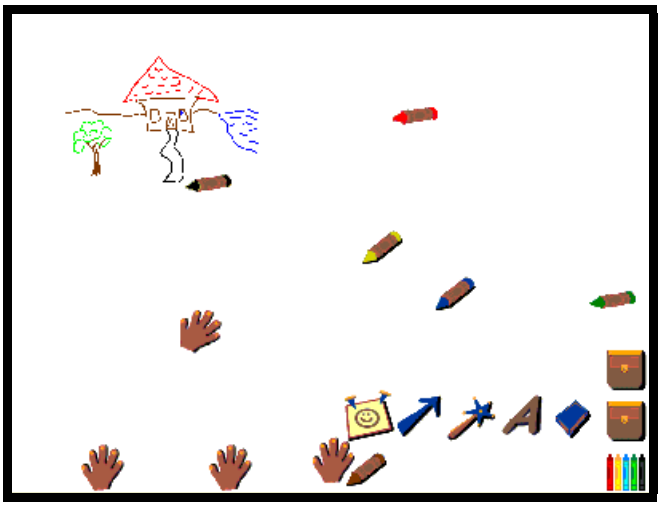

Figure 2: The initial version of KidPad showing all the toolboxes open at once with four simultaneous users.

Figure 3: the Klump, a deformable 3D modeling object

As with KidPad, two or more children can manipulate the Klump at the same time. the Klump is intended to be a more improvisational storytelling tool than a structured one. It is the real-time exploration of the properties of the Klump that leads to the creation of simple stories. Again, by supporting synchronous multi-user access and by displaying the children s cursors to one another, the Klump enables collaboration

*EMBED*The initial version of the Klump can be manipulated in the following ways:

Stretching —a point on the surface of the Klump can be grabbed using the mouse and can be pulled about to deform its shape. There is an option to switch between pulling out a single vertex and a group of vertices, thereby changing the kind of deformation that occurs. The single vertex option pulls out a thin volume of the Klump, whereas the group of vertices pulls out a fat volume. There is also a button to return the Klump back to its original roughly spherical shape.

Texturing - a variety of pre-defined textures may be applied to the surface of the Klump by selecting buttons on the interface. These textures allow different facial expressions to be added to the font-side of the Klump and enable its background colors to be changed, giving it a sense of character.

Rotating — the textureon the surface of the Klump can be grasped and rotated around to a new position.

Finally, the Klump makes a variety of sounds to reflect these different manipulations.

REDESIGNING KIDPAD AND THE KLUMP TO ENCOURAGE COLLABORATION

The core technical innovation of this paper is the idea of designing interfaces to encourage or invite children to collaborate. This has been motivated by our experiences of using the initial versions of KidPAD and the Klump in two schools, one in Sweden and one in England, during the 1998-1999 school year as part of a wide program of activities that included:

- **contextual inquiry** — sessions to observe how children work with existing storytelling technologies (e.g., crayons and paper) and how they collaborate and tell stories.
- **participatory design** — initial sessions to establish the children in the role of design partners and co-inventors of technology, followed by sessions with KidPAD and The Kump aimed at eliciting specific design suggestions.
- **evaluation of the technologies** —observations of how the children used the initial versions of KidPAD and The Kump.

Over the course of the year, the combination of these activities has resulted in more than fifty sessions in schools involving more than one hundred five and seven year olds. At the peak of this activity, there were weekly participatory design and contextual inquiry sessions.

Children were observed with respect to collaborative behavior and their ability to use the technology to tell stories. Children and teachers were encouraged to provide feedback on these technologies that would instigate changes in design. Although by the spring, small-group and whole-class collaborative storytelling activities were being performed using these technologies, the difficulty some children had in collaborating was evident.

Interfaces that encourage collaboration were proposed as a way of addressing this problem. Such interfaces should provide opportunities for children to discover the positive benefits of working together. Ideally, this should be achieved in as subtle and natural a way as possible, avoiding forced solutions. Encouraging collaboration is more proactive than only *enabling* it as was the case with the initial versions of KidPad and the Klump described previously. On the other hand it is not as extreme as strictly requiring collaboration, for example, demanding that two children have to press a button together to achieve an action, an approach that might be given the label enforcing collaboration .

We now describe how KidPad and the Klump were redesigned according to the lessons learned from the various schools sessions. Our overall strategy was to introduce design changes that satisfied two criteria:

- first they should encourage collaborative activity, reflecting the project s educational agenda and reacting to the observations noted previously.
- second, they should be based on the children s own design suggestions, emerging from the participatory design process.

Our general approach has been to use the more frequent occurring of the children s ideas as the basis for deciding on new functionality, but to realize this functionality through the approach of encouraging collaboration . For example, a common suggestion for KidPad was to provide more colors. Our response has been to allow more colors to be generated by children collaborating to mix together the existing colors.

Redesign of KidPad

The basic approach that we followed in redesigning KidPad to encourage collaboration was to support tool mixing . By this, we mean that when two (or sometimes more) children each use mixable tools at about the same time and place, the tools give enhanced functionality.

As a concrete example of this approach, consider the operation of the crayons in KidPad. The initial version provided three colours. A frequent design suggestion from the children was to provide more colours. We immediately added three more crayons, but that wasn t enough. Our final solution is to enable children to collaborate and combine their crayons to produce new colours. If two children draw with two crayons close together, then the result is a filled area between the two crayons whose colour is the mix of their crayons. In this case, the children are not prevented from drawing as individuals, but they can gain additional benefit (new colours and filled areas) by working together.

In its strictest interpretation, the approach of encouraging collaboration without enforcing it would require that a single child could achieve on their own any action that two children could achieve together, but that the two would do so in an easier, more efficient or more fun way. However, a more relaxed interpretation, is that a single child can carry out all of the major classes of action supported by the tool, but that by working together, two children can achieve subtle extensions to and variations on these actions. Thus, a single child or two children working independently can create a fully functioning drawing in KidPad, but two children collaborating can create an enhanced one, for example, with more colours. This is the approach that we have adopted in revising KidPad and the Klump.

Applying our approach involves examining combinations of actions to look for interesting benefits and effects. We can consider all actions combined with themselves, for example, what happens when two selection tools are used together in KidPad? We can also consider how actions combine with other actions, for example, what might happen if one child rotates the Klump in the Klump while another stretches it? In each case, we look for effects that are natural and useful rather than contrived.

As described above, crayons in KidPad now work this way by drawing a filled in area between the two crayons using a colour that mixes the two crayon s colours. By introducing collaborative colour mixing, we added 15 mixed colours with the six crayons, and filled areas while encouraging collaboration and without adding any new tools. (see figure 4). Tools get a special hilight when they are moved close enough to a mixable tool so the children know when they can get special mixing behaviour. Also, we added a special duplicating tool which made copies of other tools so several children could use the same tool type simultaneously. Figure 4 shows the redesigned interface with two children using mixed crayons.

Figure 4: Redesigned KidPad interface with mixed crayons being used. Note that inactive tools are faded. There are three active crayons, and two are currently being used to create a mixed area.

We built in mixing capability for multiple user of all tools, except the magic wand and toolboxes. In every case, we tried to add a special behavior that acts as if it is a natural extension from a single user behaviour. We felt this design ideal to be important in order to make it as easy as possible for children to anticipate what the mixed behaviour might be. The mixing behaviour we added is:

Crayons — As described above.

Arrow —Two or more children can squash and stretch selected drawing objects.

Eraser —One user can erase bits of a drawing object, but two children can erase an entire drawing object at once.

Hand —Two or more children can zoom the view in and out by moving their hands further apart, or closer together, respectively.

Turn Alive —Two or more children can control the animation properties of a wiggling object by moving the turn alive tools closer together or further apart.

Redesign of the Klump

In redesigning the Klump to encourage collaboration, we have focused on combining the actions of stretching and texturing with themselves.

Stretching —the initial version of the Klump enabled toggling between two modes of stretching, pulling out a single vertex and pulling out a group of vertices. The revised version enables a single child to pull out only a single vertex on their own. However, if two children synchronously pull out two vertices that are close together on the Klump s surface, the result is to pull out a whole group of vertices. Thus, the added benefit of collaborating is to be able to make a different shaped deformation.

Texturing —our redesigned version of the Klump enables the children to apply a limited number of textures to its surface by pressing buttons. The textures represent happy and sad faces as well as background textures for the three primary colours. These may be applied independently so as to combine each of the two faces with the three background colours. However, by pressing some buttons together, the children may arrive at new combined textures. Three new faces become possible: laughing (pressing happy and happy), surprise (pressing happy and sad) and crying (pressing sad and sad). In addition, the background colours can be selected together to make new combined colours (similar to combining the crayons in the revised KidPad). A single user can also select the combined textures by selecting one button and then another a short time after (while the first is seen to rotate), but it requires speed and skill.

We have also extended the sounds made by the Klump to provide feedback as to when collaborative effects are being triggered, for example, by saying cool and yippee .

Figure 5 shows the revised Klump interface. In the centre we see the Klump, currently with its laughing face on a red background. To its left are the two buttons that are used to apply happy and sad face textures. To its rights are the three buttons for applying the colors. Above the Klump are two buttons that toggle between using a mouse for stretching and using it for rotating. The red button at the bottom returns the Klump to its original shape.

Figure 6 shows the difference between single-user and collaborative stretching. On the left we see the results of a single user stretching the Klump, pulling out a single vertex. On the right we see a collaborative stretch that pulls out a group of vertices, making a larger deformation.

Figure 7 shows the different facial expressions that can be obtained using the two buttons at the left of the interface. Faces 1 (happy) and 2 (sad) are obtained by a single user pressing the button. Faces 3 (laughing), 4 (surprised) and 5 (crying) are obtained when two users select combinations of the buttons at once (happy and happy gives laughing, happy and sad gives surprised, sad and sad gives crying).

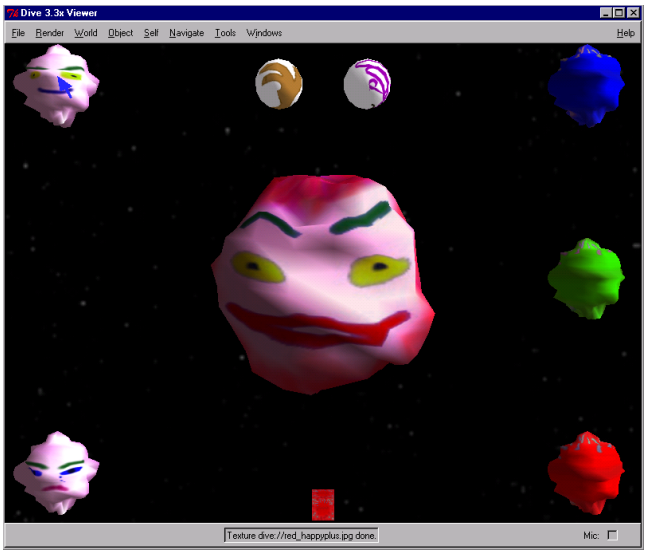

Figure 5: the revised Klump interface

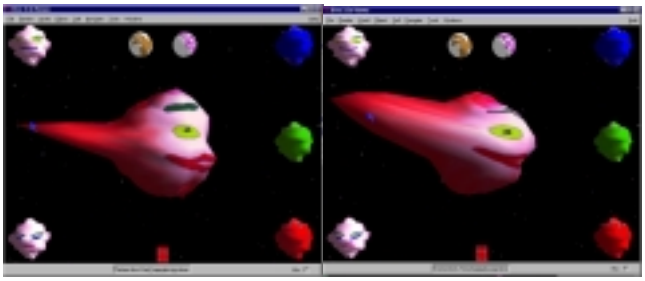

Figure 6: single user and collaborative stretching

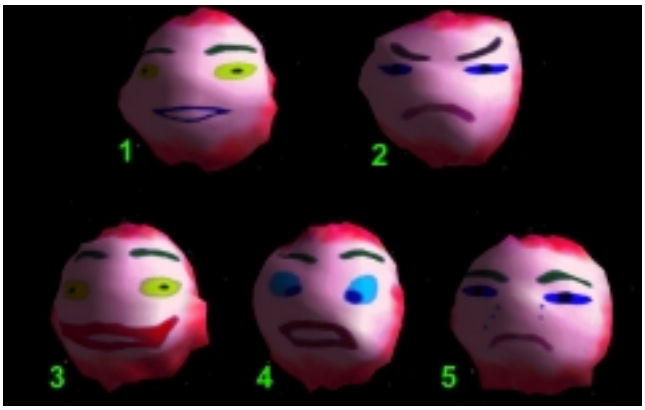

Figure 7 : facial expressions for the Kump

Initial reflections on the revised interfaces

The revised versions of both KidPad and KidDIVE were informally tested with a small group of children that are design partners at the University of Maryland s Human-Computer Interaction Lab. This formative evaluation showed that it took experience with KidPad and KidDIVE for children to make use of the collaborative tools. For example, in a one-hour session where two boys (ages 10 and 8) used the Klump, it took almost 25 minutes for the children to make use of the collaborative features in a non-combative way. (These children on a previous occasion had used a less collaborative version of the Klump for a 20 minute session). During this one-hour session with the revised version of the Klump, the two boys began by exploring the new interface. Within two minutes, the older boy was showing his younger partner different features he had found, See, you can rotate it. They quickly began imagining shapes and forms in the Klump: He has a brain tumor. That s cool! It looks like a ice cream cone.

After this initial exploration, the children were shown the software features that supported collaboration. They had not previously found any of these on their own. When the boys began to explore, they initially became fascinated with the zooming feature. Both children could type on the keyboard and fly through space to look at the Klump and controls from different angles. This led to 10 minutes of competitive behaviours: Hey, where are we? I want to go up! I want to go down! Hey, we re stuck! However, within 25 minutes, this competitive behaviour changed to a more collaborative one: Hey, let s go to the face. (both press on the face and face changes) Cool! Let s do colours! 1 2 3press! (both press on different colours and the color changes) Let s do it again!

The session ended with the boys sharing their thoughts about the system. Their overwhelming favorite was to fly through space together and change the faces and mix colours. They suggested that for future versions they would want to see multiple klumps in space so that they could talk to each other. They also wanted to add a North Star so that they wouldn t get lost when flying through space together.

Another formative study was done with six children (4 boys/2 girls; ages 7-10) using KidPad. For an hour and a half session, the three children who had previously worked with KidPad (a single-mouse version) showed strong differences in their use of collaborative tools, then the three other children who had never seen KidPad before. The children formed two teams, and each team worked on a computer with three mice. The children that already had used KidPad formed one group, and the children that hadn t used KidPad formed another group. After introducing KidPad and the new collaborative tools to the group, the children freely explored the tools for 20 minutes. Then, the children were asked to create a story with at least three scenes to zoom to and from.

The experienced children had little trouble creating a story. They collaborated throughout the process, making extensive use of the collaborative tools before starting the story, trying out the different possibilities. However, interestingly enough, they did not use the collaborative tool behaviors in the actual story creation.

The children that used KidPad for the first time had a harder time collaborating to create a story. They tended to experiment with the tools, including the collaborative tool behaviors. Most of what they did however was scribbling. From this team such comments could be heard: Take this! Take this! Take what?, It s not your computer!, Everybody stop!, Who am I?, Who s who?, Stop doing that!. We found that this group had a difficult time identifying each other s cursors and agreeing on what to collaborate on. They never seemed to move out of the independent stage, while the other team moved from independent to peer-learning to finally collaborative.

After the storytelling activity ended, the children were asked for their feedback on what they liked and didn t like about KidPad. The children overwhelmingly agreed that they liked the crayons (including the collaborative behavior of crayons) and the turn-alive tool. One child mentioned she liked taking other people s tools, although others seemed to be annoyed when this occurred. The children also mentioned they liked cooperating as long as the rest of the children wanted to cooperate. Among the dislikes, the first time users of KidPad said they disliked their team and disliked cooperating with their teammates.

RELATIONSHIP TO OTHER WORK INTO SYNCHRONISING SHARED INTERFACES

So far, we have introduced the idea of interfaces that encourage collaboration within the context of educational applications. We now consider its broader relationship to

CSCW technologies, especially how it compares to other approaches to synchronizing shared interfaces

How to synchronise shared interfaces has been a major concern for CSCW research. This has predominantly focused on distributed groupware where multiple users share a common workspace, for example a shared document, 2-D sketch tool or 3-D virtual world, using separate displays connected over a computer network. In such cases, the problem of synchronization can be broadly broken down into two parts.

How to synchronize what different users see? One of the first approaches was WYSIWIS (What You See Is What I See) where different users at different displays were forced to see the same part of the virtual workspace [Stefik87]. Experience with WYSIWIS led to less strictly coupled approach called relaxed WYSIWIS where different user s views could diverge [Stefik87]. Systsmes adopting this approach typically introduce additional functionality to support users in being aware of where others are looking and what they are doing. This may take to form of various awareness widgets, such as radar views in D workspaces [Greenberg98] or visible user embodiements (avatars) in 3D systems [benford95].

How to synchronize object manipulations? Many CSCW systems allow user s to collaboratively manipulate objects, changing their state. Examples include jointly editing a shared document or grasping and moving objects in a virtual world. This raises the problem of how to prevent conflicting updates. The most common solution is some form of locking, including simple turn-taking protocols, optimistic locking, non-optimistic locking and serialization protocols that allow participants to interleave their actions at various granularities [Greenberg94]. Another option is social locking where given sufficient mutual awareness, user s may be able to negotiate mutual access with minimal system intervention.

We suggest that these various strategies can be located along a collaboration continuum according to the extent to which they constrain individual actions and enforce collaboration. One extreme of the continuum involves *enforcing collaboration*, where the users are locked in step with one another. WYSIWIS and strict turn-taking can be found here. So can the idea of requiring children to work together to drive an educational application [Blaye et al., 1992]. At the other is enabling collaboration, where the users can act independently, are mutually aware are free to coordinate their actions if they wish. Relaxed-WYSIWIS and social locking can be found here.

Our approach of *encouraging collaboration* lies somewhere between the two. It is not so strict as to require users to work together, but it provides some explicit motivation for them to do so in terms of added benefit. As noted earlier, encouraging collaboration can be interpreted in different ways. The case where a single user could achieve any action, but multiple users can achieve it a way that is easier or more fun lies towards he enabling end of the continuum. The case where a single user can carry out each general class of action, but where multiple users can achieve enhanced actions lies towards the enforcing end. The overall classification is shown in figure 8.

It should be noted that a single CSCW system can combine different approaches for different actions. For example, collaborative virtual environments often enable collaboration for viewpoint control (each user steers their own viewpoint, but is made aware of others viewpoints through their embodiments), but enforce it for object manipulation (there is a turn-taking or coarse locking protocol for who can grab a virtual object).

This discussion raises the question of how the approach of encouraging collaboration might be useful out of the contexts educational applications. One possible application area is in entertainment and games applications where participants might choose to collaborate, pooling abilities and resources to mutual benefit. Another more subtle approach might be in situations where participants can benefit by sharing costs. People increasingly have to pay for the use of network resources, for example in video and audio streaming. Users who agree to collaborate, for example to receive or manipulate the same information might be rewarded by sharing the costs between them.

Summary and future work

In summary, we have proposed a new approach to designing shared interfaces that is intended to support children in learning to collaborate. The approach, called encouraging collaboration, allows children to work as individuals, but gives added benefits if they choose to work together. We have demonstrated this approach applied to the design of two storytelling technologies within the more general framework of participatory design within UK and Swedish schools. Finally, we have argued that our approach can be situated along a continum of collaboration mechanisms along with other common approaches from CSCW and have suggested that it may have further applications for adults within areas such as gaming, entertainment and the more general sharing of resources.

REFERENCES

1. Barfurth, M.A. (1995) Understanding the Collaborative learning process in a technology rich environment: The case of children s disagreements. *Proc CSCL 1995*,
2. Bederson, B. B., & McAlister, B. (1999). Jazz: An Extensible 2D+Zooming Graphics Toolkit in Java., *University of Maryland Computer Science Tech Report #CS-TR-4015.*
3. Bederson, B. B., Hollan, J. D., Druin, A., Stewart, J., Rogers, D., & Proft, D. (1996). Local Tools: An Alternative to Tool Palettes. *UIST 96*, pp.169-170.
4. Bier, E. A., & Freeman, S. (1991). MMM: A User Interface Architecture for Shared Editors on a Single Screen. *UIST 91,* pp. 79-86.
5. Blaye, A., Light, P. & Rubtsov, V. (1992) Collaborative learning at the computer: How social processes interface with human-computer interaction. *European Journal of Psychology of Education*, *7*(4), 257-267.

6. Buxton, W., & Myers, B. A. (1986). A Study in Two-Handed Input. *CHI 86,* pp. 321-326.
7. Druin, A. (1999). Cooperative Inquiry: Developing New Technologies for Children with Children. *CHI 99*, pp. 223-230.
8. Druin, A., Stewart, J., Proft, D., Bederson, B. B., & Hollan, J. D. (1997). KidPad: A Design Collaboration Between Children, Technologists, and Educators. *CHI 97*, pp. 463-470.
9. Fahl n, L. E., Brown C. G., Stahl, O., Carlsson, C., A Space Based Model for User Interaction in Shared Synthetic Environments, *InterCHI' 93*
10. Greenberhg, S. & Marwood, D., Real Time groupware as a Distributed System: Concurrency Control and its Effect on the Interface, *CSCW' 94*, 207-217
11. Gutwin, C. & Greenberg, S., Design for individuals, Design for Groups: Tradeoffs betwenbetween Power and Workspace Awareness, *Proc, CSCW' 98*, 207-217.
12. Inkpen, K., Booth, K. S., Klawe, M., & McGrenere, J. (1997). The Effect of Turn-Taking Protocols on Children' s Learning in Mouse-Driven Collaborative Environments. *In Proc Graphics Interface (GI 97)* Canadian Information Processing Society, pp. 138-145.
13. O Malley, C (1992) Designing Computer Systems to support peer learning, in *European Journal of Psychology of Education,* Vol. VII, No. 4, 339-352.
14. Rogoff, T. (1990) *Apprenticeship in Thinking: Cognitive development in social context.* Oxford University Press, Oxford.
15. Stefik, M., Bobrow, D., Foster, G. Lanning, S., & Tatar, D., WYSIWIS Revised: Early experiences with multi-user interfaces, ACM TOIS, 5(2), 147-167, 1987.
16. Stefik, M., Foster, G., Bobrow, D., Kahn, K., Lanning, S. & Suchman, L., Beyond the Chalkboard: Computer Support for Collaboration and problem Solving in Meeting, CACM, 30(1), 32-47, 1987.
17. Stewart, J. (1998). Single Display Groupware. Doctoral dissertation, University of New Mexico, Albuquerque.
18. Stewart, J., Bederson, B. B., & Druin, A. (1999). Single Display Groupware: A Model for Co-Present Collaboration. *CHI 99*, pp. 286-293.
19. Stewart, J., Raybourn, E., Bederson, B. B., & Druin, A. (1998). When Two Hands Are Better Than One: Enhancing Collaboration Using Single Display Groupware. CHI 98 *Extended Abstracts*, pp. 287-288.
20. Topping, K. Cooperative learning and peer tutoring: An overview. *The Psychologist,* 5(4), 151-157, 1992
21. Wood, D. & O Malley, C., Collaborative learning between peers: An overview. *Educational Psychology in Practice*, 11(4), 4-9, 1996.
22. Wood, D., Wood, H., Ainsworth, S. & O Malley, C., On becoming a tutor: Toward an ontogenetic model. *Cognition and Instruction*, *13*(4), 565-581, 1995

HYPERLINK**ACKNOWLEDGEMENTS**

KidStory is funded under the ESPRIT i^3 Experimental Schools Environment initiative of ESPRIT. We are deeply grateful to our partners at the Albany Infant School in Nottingham, England and at R gsvedsskolan in Stockholm, Sweden. We would like to thank our summer evaluation team of children at the University of Maryland s HCIL. We also appreciate the initial work of Jason Stewart on KidPad and Single Display Groupware, and the initial DIVE work of researchers at the Swedish Institute of Computer Science