Submission to International Journal of Human-Computer Studies

Title: Mutual Engagement and Collocation with Shared Representations

Author: Nick Bryan-Kinns

Affiliation: Queen Mary University of London

Address: Interactional Sound and Music Group

Centre for Digital Music

School of Electronic Engineering and Computer Science

Queen Mary, University of London

Mile End, London. E1 4NS

Abstract

In this paper we explore the use of shared representations to support creative activities, focussing on collaborative music making. We examine the effect that user interface features of shared representations have on mutual engagement and show that providing shared awareness mechanisms increases mutual engagement. In particular, we show through an empirical study of 78 participants that providing cues to identity and shared pointers increases mutual engagement between participants, but together these features can overwhelm users. We also demonstrate that support for free-form annotation and spatial interaction with shared representations mediates interaction and helps participants to manage their collaborative activity effectively. In this paper we develop several measures mutual engagement and demonstrate their use to assess the design of shared representations. A key contribution of this paper is the development of a measure of collocation of participant interaction which indicates mutual engagement. The findings of the study have implications beyond the domain of collaborative music making and we outline some design guidelines for mutually engaging shared representations.

Keywords: mutual engagement; collaboration; music making; design; evaluation; creativity.

1. Introduction

Collaboration relies on shared understanding (Clark, 1996) and shared representations in collaborative work (Fischer et al., 2002; Suwa and Tversky, 1997; Hutchins, 1995). Shared representations such as diagrams have been shown to be important to collaborative activities in education (Suthers et al., 2003), the workplace (Cherubini et al., 2007), and beyond - in creative acts such as design (Bly, 1988) and collaborative puzzle solving (Kraut et al., 2002). Shared representations have properties that are central to the collaborative activity (referred to here as *primary* properties), and properties that are not essential to the completion of the activity (referred to here as *secondary* properties). For example, a shared text editor could be considered to provide a shared representation where the text, fonts used, and page layout are primary properties (a document cannot be created without them), whereas additional information such as lists of editors, displays of shared pointers, and so on, are secondary properties (the document could be created without them). As highlighted in literature on the mediating properties of technologically shared representations (cf. Bannon and Bødker, 1997) collaborative tasks can be completed solely with the primary properties of shared representations, but secondary properties provide scaffolding for the interaction and if well designed can improve the efficiency and quality of collaboration.

The domain of Computer Supported Collaborative Work (CSCW) has been exploring the design and evaluation of technologically mediated shared representations for over twenty years, particularly in office-work situations (e.g. see Stefik et al., 1987, for discussion of early experiences with multi-user interfaces). Two key concerns have emerged through the design and use of CSCW systems: the role of *awareness* of who is doing what with the shared representation (e.g. Dourish and Bellotti, 1992), and *management of access* to elements of shared resources (e.g. Abbott and Sarin, 1994). Stefik etl al.'s early work (ibid.) highlights the use of shared pointers (a secondary property of a shared representation where each user sees some representation of other users' mouse pointer location) to support awareness of each collaborators' focus of attention, and flexible spatial layouts to allow users to group and manage contributions to shared representations as they see fit. Gutwin and Greenberg's characterisation of *workspace awareness* (2002) breaks awareness down into three categories focussing on awareness of who is doing things with the shared representation, what they are doing, and where they are carrying out their actions. Managing access to elements of shared representations is a primary property which has been explored on the one hand through sophisticated access control mechanisms, and on the other by simply allowing users to lay elements out in space as they would in naturalistic settings (cf. Tse et al., 2004). In naturalistic groupwork settings, managing access to shared representations and cues to awareness are usually supported by virtue of collaborators' co-presence and shared physical access to the resource being worked on. The challenge for CSCW systems, then, is to explicitly design such awareness and access mechanisms in a way that improves the use of shared representations for all collaborators.

Whilst there has been some research into supporting group creative acts, for example in collaborative virtual environments (Benford et al., 2000), creativity support tools (Shneiderman, 2007), interactive art (Cornock and Edmonds, 1973), and in guerrilla environments (Sheridan and Bryan-Kinns, 2008), there is still substantial work to be done on understanding how shared representations can be effectively used to support collaborative creativity, especially when participants are not collocated. For example, Benford et al.'s work on poetry performances in collaborative virtual environments highlighted the importance of spatial characteristics of shared representations in collaboration, and provided embodiments of users in shared space to convey awareness of action, ownership, and attention. However, their case study of virtual poetry was exploratory, and results may not be applicable to other collaborative creative domains.

2. Shared Representations in Music Making

Making music is a key form of creative human collaborative activity. Indeed, Freeman argues that 'musical skills played a major role early in the evolution of human intellect' (2000). The process of music making from improvisation to composition has been studied and argued about for centuries, and there is not space to explore the debates here. Similarly, we do not explore the debates surrounding the nature and processes of creativity (cf. Boden, 1992) and group creative activities (cf. Mamykina et al., 2002). Instead, we focus on the design and use of shared representations for music making without distinguishing between the nature of composition, improvisation and performance. In particular, we are interested in exploring the observable effects different characteristics of shared representations have on music making. Music making is an interesting domain to study shared representations as compositional elements are frequently translated from one form to another e.g. from written score to musical sounds (Nabavian and Bryan-Kinns, 2006). Moreover, the final product is often not the shared representation that was used to create it e.g. an audio recording does not usually contain the visual score used in its creation and production.

In naturalistic settings such as improvisational 'jams' (cf. Sawyer, 2003) musicians make music together using shared audio and visual resources available to them by virtue of their co-presence - the sounds produced by their instruments and voices (Flor and Maglio, 1997), their physical movement and location (Healey et al., 2005), and information recorded visually or in audio (Nabavian and Bryan-Kinns, 2006). Indeed, in naturalistic situations physical location, access to, and orientation of participants, their instruments, and shared representations, are critical to managing the collaboration (ibid.). The study and exploration of music making with interactive digital technologies has been growing rapidly in recent years with conferences such as NIME (New Interfaces for Musical Expression; Poupyrev et al. 2001) demonstrating the breadth of new forms of music making being developed. However, whilst such research has developed a small number of tools for collaborative music making where participants are not collocated (e.g. CODES, Miletto et al., 2006; and Daisyphone, Bryan-Kinns, 2004) such research often ignores the rigorous evaluation of human interaction with these new devices (Stowell et al. 2009), let alone examining the nature of the collaborative acts supported (cf. Barbosa 2003), or the expressivity of the devices (cf. Dobrian and Koppelman 2006). Moreover, whilst there is ongoing contemporary research on replicating naturalistic interaction through technological mediation such as supporting creativity and group music making at a distance through high quality shared video-links (e.g. Sawchuk et al., 2003), we focus on developing novel means to support collaborative interaction informed by, and in turn informing understandings of the nature of human interaction.

In designing for creative collaborations such as music making, we are particularly interested in moving beyond analysis of the *pragmatic* interaction with shared representations (e.g. Hutchins, 1995, focus on shared information displays for navigation) to identify when people creatively spark together and start to *mutually engage* (Bryan-Kinns and Hamilton, 2009) with each other as well as the shared product. These points of mutual engagement are essentially points of group *flow* (cf. Csikszentmihalyi, 1991); Sawyer (2003) provides ethnographic context for group flow activities - we are interested in how the design of shared representations impacts participants' group flow through observations of their interaction. In addressing this issue, Bryan-Kinns and Hamilton (2009) drew on models of human communication (e.g. Clark and Brennan, 1991) and CSCW research on workspace awareness (e.g. Gutwin and Greenberg, 2002) to propose that user interfaces which are key to supporting the collaborative aspects of mutually engaging interaction are: *Mutual awareness of actions*, *Shared and consistent representations*, *Mutual modifiability*, and *Shared annotation*. Moreover, they showed that providing a shared representation of the music being collaboratively produced in a distributed situation (the primary property of the shared representation) along with cues to the authorship of contributions (a secondary property of the shared representation) and support for shared annotation mechanisms (also a secondary property) increased the mutual engagement between participants. However, considering the highly interactive nature of music making, and the key role spatial orientation and location of participants play in naturalistic group music making there is a need to further explore the design of mutual awareness mechanisms for shared representations of music. In particular, whether awareness of each others' interaction beyond the actual production of music (i.e. awareness being a secondary property of the shared representation) has an effect on mutual engagement, and whether awareness mechanisms have an effect on collaborators' use of shared space (spatial organisation being another secondary property).

3. Understanding Shared use of Space

The importance of participants' physical location and orientation in naturalistic group music making in conjunction with the importance of spatial arrangements of objects as illustrated by studies of how work is organised (e.g. Malone, 1983) leads to a key concern of this paper - how user interfaces features designed to support mutual awareness affect participants' use of the secondary property of spatial layout in the shared representation of music. Intuitively we would expect that as participants become more mutually engaged they tend to collocate their activities closer together – i.e. collocation of shared activity would be an additional indicator of mutual engagement between participants. However, identifying collocated activities in groups on (relatively) small shared spaces is problematic as simply visually inspecting traces of groups' spatial activity is extremely time consuming and the data is difficult to interpret. For example, consider the three traces of shared pointer movement illustrated in figure 1a, 1b, and 1c. Whilst it is clear that there are differences in the spatial distributions of colour (representing four users – light grey, mid grey, dark grey, and black) with users in figure 1b appearing to be more spatially collocated (overlapping and intermixed) than figure 1a, it is difficult to judge how these are different to figure 1c.

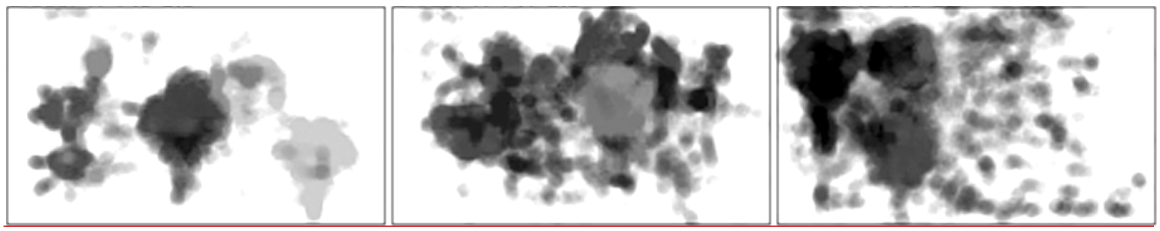

Figure 1a, 1b, 1c: Three traces of shared pointer movement

CSCW's extensive literature on understanding the nature of distributed shared activities has usually focussed on shared document editing (predominantly text editing) and video-conferencing as key forms of interaction. Moreover, the focus has predominantly been on exploring the nature of the collaborative process (Pinelle et al., 2003), and awareness mechanisms (Dourish and Bellotti, 1992) with little exploration of the effect of collaboration support on collaborators' use of spatial properties of shared objects even though collaborative drawing has been extensively explored (e.g. Bly, 1988; Ishii et al., 1993). The major exception to this is the work by Tse et al. (2004) on how people use shared space in Single Display Groupware (SDG cf. Stewart et al., 1999), specifically, how users co-ordinate their actions to avoid interference between their actions such as working atop each other, or opening shared menu items over each others' work. Tse et al. found that in collaboration people naturally partition their workspace to avoid interfering with each others' work. Furthermore, they proposed a number of design guidelines for SDGs which included the proposal that designs should "Let people move artifacts (such as tools) to their preferred locations". However, Tse et al.'s work examined pairs of users, so it is not clear how their findings and suggestions would scale to larger groups of participants. Furthermore, whilst participants shared a single display and used separate input channels, they were physically collocated, so it is not clear how their findings apply when participants are not physically collocated.

Research on collaborative use of interactive tabletops compliments the above research, and has a greater emphasis on the spatial elements of the interaction, probably due to the inherently spatial nature of collocated tabletop interaction. Research on the use of interactive tabletops (for overview see: Scott et al., 2003) has demonstrated their support for collocated groupwork (Tse et al., 2006) through their inherent support for grounding (cf. Clark, 1996) due to the physical collocation of participants around the shared space. Unravelling the spatial nature of the group interaction, Scott et al. (2004) highlighted different patterns of group behaviour with respect to spatial arrangements of shared objects on, and represented in, interactive tables - identifying emerging patterns of personal, group, and storage territories. However, their studies rely on manual cateogrisation of shared object location, and consider each participant to have a fixed ‘origin’ of interaction - their physical position and orientation around a shared table (Kruger et al., 2005). In contrast, participants using distributed shared representations such as Daisyphone (Bryan-Kinns, 2004) typically have no fixed origin of interaction - they have the same orientation (all look at the screen the same way up), and have equal opportunity to interact with the shared space as they are not restricted by their physical position or reach of their arms. Pinelle et al. (2008) developed a measure of the spread of users’ interaction with a shared tabletop, and showed that different interaction techniques had a significant effect on the area of users’ interaction. Moreover, they assessed whether participants’ personal interaction areas overlapped by considering personal territory to be ‘within a radius of 65cm of the seating point of the user’ (ibid.). However, as with other work discussed in this section, the assumption that users have a natural ‘home’ location in the interaction does not hold for distributed shared representations as users are not collocated in physical space, and no restrictions are placed over where users can point to. Kruger et al. (2005) also examined the role of spatial location of shared objects and identified the establishment of personal and public space in a shared tabletop environment.

Two factors contributed to the establishment of personal or public ownership of items - their proximity to the users, and their orientation (objects were parts of jigsaw puzzles whose decals could be oriented towards a user). Again, as distributed users often share the same virtual location and orientation, it is the role of spatial organisation is of interest to us in this paper.

4. Daisyfield

In order to examine the effect of shared awareness of participants' loci of attention, and the effect of mutual awareness on the use of shared space, we developed a new shared music making system referred to as *Daisyfield* which supports multiple distributed participants co-creating music using a single shared representation. The primary properties of Daisyfield's shared representation are multiple short loops (1 minute) of music which are shared between participants. Inspired by studies of naturalistic music improvisation (Healey et al., 2005) and composition (Nabavian and Bryan-Kinns, 2006) which highlight the importance of spatial properties of shared representations, and Tse et al.'s Single Display Groupware design guidelines (2004), Daisyfield allows participants to place and move their loops in a shared visual space. In this version of Daisyfield, the spatial location and orientation of musical loops does not affect the shared sound produced - it is purely a visual secondary property of the shared representation. Daisyfield provides a richer user experience and compositional potential than previous collaborative music making systems such as Daisyphone (Bryan-Kinns, 2004) which only allowed for one loop of music, and very limited use of shared space.

In Daisyfield, participants can edit and move each others' musical loops - each loop is referred to as a Daisy. Figure 2 illustrates the interface in use. In the figure there is a Daisyfield (1250x720 pixel green rectangle shown as dark grey in the figure) containing three shared Daisys which represent three shared loops of music which are played concurrently. The large Daisy to the right has been opened by the user for editing. There is also a circle toward the bottom left of the Daisyfield indicating the location of another users' pointer. Colours in this interface indicate which user created the content – each user is assigned a unique colour. Users click the background to create a Daisy (maximum of 12 on a screen at once). The Daisys can be dragged around the screen, and are dragged off the edge of the window to delete them. Clicking a small Daisy opens it to allow users to add and remove notes by clicking on the dots - the notes are played when the rotating line passes over them. The four shapes in the centre of each Daisy allow for selection of different sounds. The circle below the Daisy represents the volume of the notes in the Daisy - it can be dragged to change the volume. Graphical marks can be made by clicking and dragging in the background. Notes, Daisys, and drawings are shared between all the people interacting with Daisyfield.

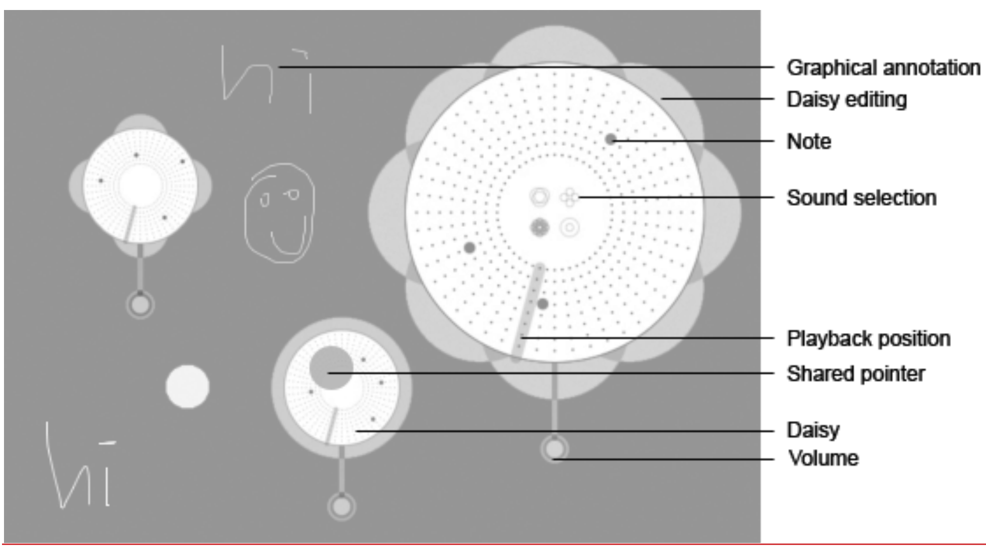

Figure 2: The Daisyfield user interface

Daisyfield's shared representation's primary property is the music being created. Each loop of music with moveable within Daisyfield's shared space (a secondary property of the shared representation) - this allows participants to visually group loops together if they wish. Daisyfield' shared representation also has a secondary property which indicates participants' current *loci of attention* (their current point of interaction with the shared representation) using shared pointers, and a secondary property of annotation which allows participants to communicate through shared drawings. The use of shared pointers has long been contentious in collaboration (cf. Stefik et al., 1987) as on one hand it increases awareness, but on the other can be distracting for users. Taking this into account, Daisyfield's shared pointers are semi-transparent in an attempt to provide awareness whilst reducing distraction. Furthermore, we have provided only simple representations of users' loci of attention - a circle representing the shared pointer as opposed to some avatar style representation - as an initial step in understanding the role of shared representations of users without overly increasing the complexity of the interface. In terms of Schafer and Bowman's frames of reference in spatial collaboration (2004), Daisyfield provides an exocentric view of the shared representation where users are not immersed in the environment, and furthermore, avatars are essentially semi-translucent circles. Clearly there is scope for further investigation of representations of users, and even egocentric views of the shared representations where users are actually immersed in the musical form. However, as discussed previously, we are interested in exploring new forms of interaction rather than re-creating naturalistic settings. In summary, Daisyfield's shared representation consists of the following shared and consistent properties:

Primary properties:

- Shared audio.
- Up to 12 shared musical loops that produce the shared audio.
- Shared editing of musical loops.

Secondary properties:

- Shared positioning of musical loops on a 2 dimensional plane.
- Shared indications of authorship of musical loops.
- Shared graphical annotations on a 2 dimensional plane.
- Shared pointers which indicate loci of attention of other participants.

Daisyfield is written in Java and uses a client-server model to support the networked sharing of data and collaborative interaction. It runs in full-screen mode with no other user interface objects visible or accessible.

5. Experiment

Studies of the simple collaborative music making software Daisyphone show that providing cues to the identity of users, and annotation tools increases mutual engagement (Bryan-Kinns and Hamilton, 2009). In this study we examined whether cues to identity increase mutual engagement with Daisyfield - a richer musical environment than Daisyphone - and the effect that shared pointers had on mutual engagement. We were also interested in the role of shared space in creative collaborations. Building on previous research on mutual engagement and CSCW, we developed two hypotheses:

H1: Mutual engagement would be greater when participants had explicit cues to each other's identity.

H2: Mutual engagement would be greater when participants had indicators of other participants' loci of attention.

5.1. Independent Variables

Two independent variables were manipulated:

- A between-subjects factor (non-repeated) of *Colours* - whether each participant was associated with a unique colour for their annotations, Daisys, and their shared pointer representation (the *Colours* condition), or each participant had the same colour for their annotations, Daisys, and their shared pointer representations (the *No Colours* condition).

- A within-subjects factor (repeated) of shared *Pointers* - whether each participant saw a representation of other participants' current mouse pointer location (the *Pointers* condition), or did not (the *No Pointers* condition).

5.2. Dependent Variables

Building on previous studies of mutual engagement (Bryan-Kinns and Hamilton, 2009) we developed three categories of kinds of dependent measure to assess the mutual engagement between participants: Activity assessment (ratio scale measures of what participants did), Content assessment (nominal and ordinal scale judged assessments of what participants did), and Participant feedback (ordinal scale self-reporting assessments of what participants did).

Activity Assessment

Objective measures of the interaction with Daisyfield can be derived from numerical analysis of logs of participants' activity with the user interface. Daisyfield logs every mouse pointer movement, click, and drag, with a timestamp and participant ID. We split these measures into three categories:

- **Individual activity**: amount of mouse movement (i.e. Pointer movements on the shared display), amount of movement of Daisys around the shared screen, and amount of annotation.

Note that pointer movements are considered separately to annotations – the underlying assumption is that pointer movements are focussed on the primary property of the shared music.

- **Musical activity**: number of Daisys in final tune, and number of notes in final tune. Previous studies showed that complex tunes indicated a lack of mutual engagement between participants due to the lack of focussed engagement with the activity.

- **Collaborative activity**: amount of mutual modification (changing someone else's content), and measures of the spread of mouse movements, annotations, and Daisys across shared display.

Content Assessment

The aim of the activity was to create a musical jingle which sounded good. We developed a simple coding scheme to judge the quality of the **musically** of the final compositions. To mitigate against subjective assessment of musical quality, we did not evaluate whether the music was *good* per se, but rather whether there was evidence of musicality in the construction of musical motifs and whether these motifs were integrated into a coherent whole. This was judged on a 5-point Likert scale: 1 (Not musical at all); 3 (Some musical motifs and some attempt at integration); 5 (Collaborative effort fitting musical motifs together into one whole piece).

We developed a coding scheme to categorise the **topics** of textual and graphical communication between participants based on similar coding schemes by Applegate et al. (1986), Olson et al. (1993) and Bryan-Kinns and Hamilton (2009). The difference here is that we also categorised the probable topic of graphical marks (e.g. when lines are used to divide a screen into interaction areas) to help understand the role of shared annotation on mutual engagement. This provided a mechanism to compare the focus of interaction between participants, and is divided in five topics:

- **System related** – discussions of Daisyfield system technical problems, or the instructions for the experiment e.g. 'When should we start?'
- **Presence and identity** – statements about presence and identity within Daisyfield, e.g. 'Hi, Bonny'
- **Query presence and identity** – questions about other participants' identity and presence in Daisyfield, e.g. 'Who is yellow?'
- **Quality judgement** – content related to the quality of the music being produced, e.g. 'this sounds awesome!', or a tick mark.
- **Task organisation** – content about the process of completing the task itself, such as role assignment, or spatial layout of Daisys, e.g. 'Mo - make tunes and update progress'
- **Social** – non-task related content, e.g. 'LOL!!' or drawn objects such as cats or faces.

Participant Feedback

We developed our Mutual Engagement Questionnaire (MEQ) based on the questionnaire developed by Bryan-Kinns and Hamilton (2009) to identify mutual engagement with different interface configurations.

The MEQ has two parts:

1) Participant ratings of the quality of the composition and their interaction with each interface.

2) A series of comparisons between the two interfaces aimed at identifying which interface was found most mutually engaging. The comparisons aim to capture: "(a) satisfaction with the

product; (b) feelings of enjoyment or flow (cf. Csikszentmihalyi 1991); (c) sense of collaboration; (d) usability" (Bryan-Kinns and Hamilton, 2009). Participants were asked to choose which interface they felt the following statements were most appropriate to:

(a) 'The best jingle was produced', 'I felt satisfied with the result'

(b) 'I enjoyed myself the most', 'I felt out of control'

(c) 'I felt most involved with the group', 'I understood what was going on', 'Other people ignored my contributions'

(d) 'The interface was frustrating', 'The interface was most complex'

5.3. Procedure

26 trios of Undergraduate students were recruited to take part (78 participants randomly assigned into trios).

Participants sat at a desktop PC and wore headphones. Daisyfield was running on the PC, and a paper information sheet guided them through the tasks. They were asked to spend 5 minutes completing a pre-task questionnaire to collect demographic data and establish their musical and computer literacy. Participants signed a consent form and were advised that they could leave at any time. Participants were then asked to undertake the following activities, and use the timer at the top of the screen to tell them when to move on to the next activity. First they undertook 15 minutes of training. Participants followed instructions on using Daisyfield on their own, trying different patterns of notes, different kinds of sounds, and different volumes to make a good tune. Note that there were no shared pointers in the training mode. Participants were then asked to create two 'Olympic jingles' using Daisyfield with their team-mates (who they did not know and could not see or communicate with). They were free to interpret this brief in a variety of ways, but they were instructed that the composition should be short and aim to have some form of musical ‘hook’. They had 15 minutes per jingle. Finally, they had 5 minutes to complete the post-task questionnaire.

6. Results

24 trios completed the study and undertook each task for the allotted time. The two trios who did not complete the study were prevented from completing the activity by system crashes. 58 participants completed their questionnaire (46 male, 12 female), the average age of the participants was 22, 15 said that they could play a musical instrument, and 12 said that they had composed a piece of music before. None had used Daisyfield before. The most popular category of music style was “Rap/ Hip Hop/ R 'n' B” (39 participants). 29 participants stated that they had used computers for collaborative work giving examples such as “online gaming” and “google docs”. Average computer literacy of participants was self-rated between ‘Intermediate’ and ‘Expert’ with 29 of the participants rating themselves as ‘Expert’. On average participants reported themselves to be online between ‘Once a day’ and ‘Once an hour’.

6.1. Activity Assessment

In this section we report on the results for the dependent variables which measured activity using a ratio scale. Data was collected from the time stamped log files of user interaction with Daisyfield.

Individual Activity Measures

When participants had **Colours** to indicate identity, there were **significantly more Annotations and Daisy Moves** when compared using a Mann-Whitney test ($p < 0.05$, see table 1and figure 3).

Activity	No Colours	Colours	U_A	z	p
Annotation	1933.10	3106.53	180	1.76	0.0392
Daisy Moves	810.52	1070.37	373.5	-1.86	0.0314

Table 1: Mean numbers of Annotations and Daisy Move activities

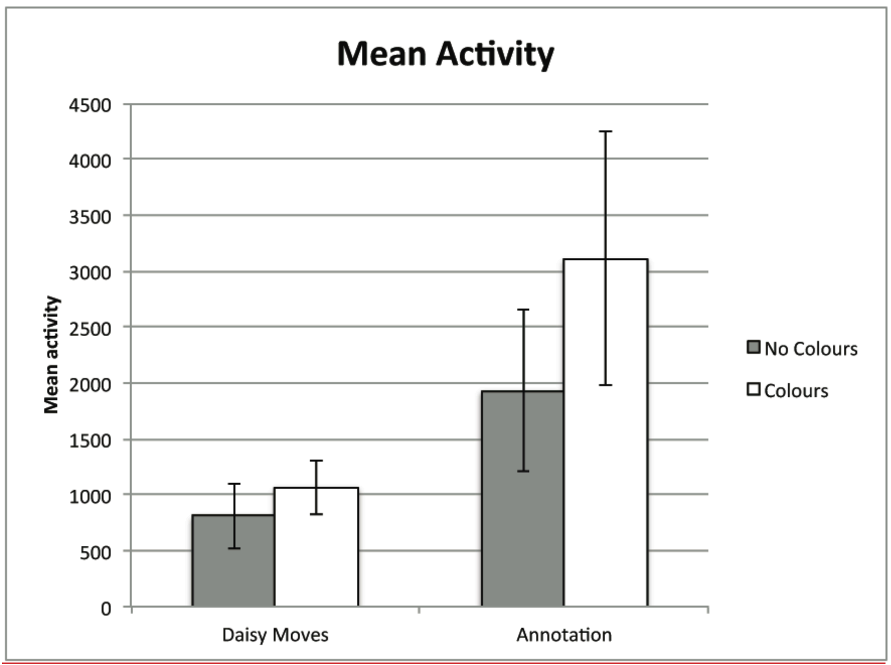

Figure 3: Mean numbers of Annotations and Daisy Move activities

Musical Activity Measure

When participants had shared **Pointers** the final compositions had significantly **fewer Remaining Notes** (i.e. less complex tunes) than when they had no shared pointers; No Pointers mean 253 remaining notes, Pointers mean 173 remaining notes, Mann-Whitney Test $p = 0.0089$ ($U_A = 403$, $z = -2.37$). Figure 4 illustrates an example of low numbers of Remaining Notes on the left hand side (low musical complexity), and higher numbers of Remaining Notes on the right hand side (high musical complexity).

Having Shared **Pointers** also had an effect of number of Remaining Notes when participants had Colours. With **Colours and Pointers, participants had significantly fewer Remaining Notes** (mean = 139) than with No Pointers and Colours (mean = 265); Wilcoxon Signed-Rank Test p = 0.0047 (W = 75, z = 2.6).

The number of Remaining Daisys was not significantly different between conditions (mean = 5). Colours had no significant effect on Pointer Moves, or number of Remaining Notes (i.e. complexity of the music created; mean = 212).

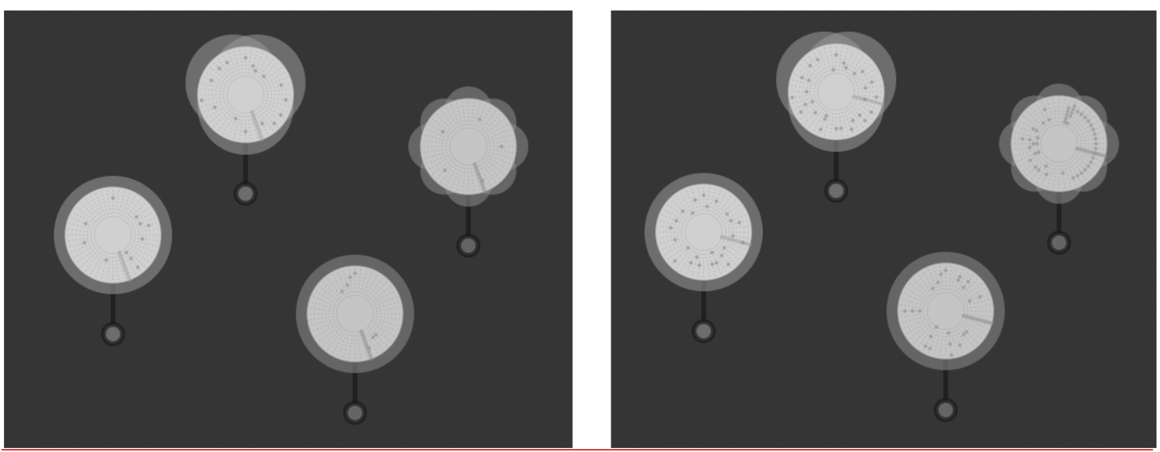

Figure 4: Examples of a) Low Musical Complexity, b) High Musical Complexity

Collaborative Activity Measures

We found no significant differences in the amount of co-editing between participants.

We developed the Collocation Measure (CM) to indicate how closely participants located their activity to each other. It is based on the perimeter of the triangle joining the centre of three users' activity. CM is calculated by finding the mean *x* and *y* pixel values for a particular activity for each participant, then calculating the length of the line bounding each users' mean point (i.e. in a triangle), and finally dividing by maximum possible length of line. For the 1250x720 Daisyfield screen, the maximum distance possible between three users is 3412.5. CM ranges from 0.0 to 1.0, with CM = 0.0 when users are furthest apart, and CM = 1.0 when users' mean interaction points are closest together (i.e. collocated means resulting from spatially overlapping activity). For example, figure 5 illustrates two screens in which the activity points for users A, B, C are indicated by letters, and the mean activity points by labeled circles with bold letters. The perimeter of the triangle joining the three mean points is calculated and used to generate the Collation Measure (CM), with high collocation illustrated on the left, and low collocation illustrated on the right.

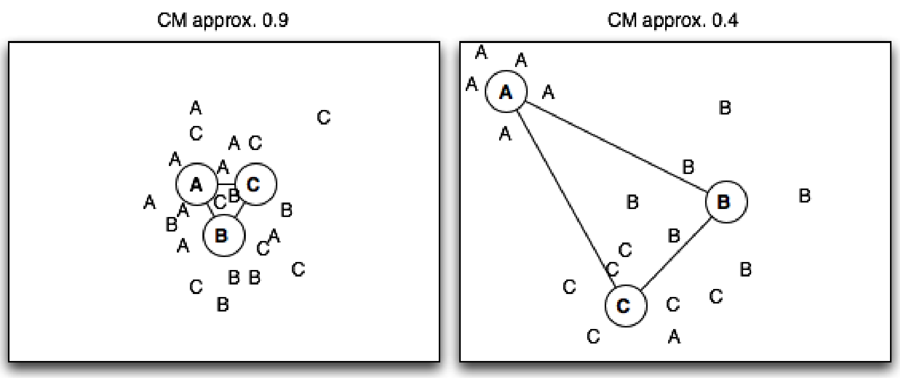

Figure 5: Example Collocation Measures for points A, B, C

With **Colours**, participants' **Annotations, New Daisys,** and **Pointer movements were closer together** when comparing CM using a Mann-Whitney test ($p < 0.05$, see table 2 and figure 6). We examined Annotation, New Daisy, and Pointer Movement collocation as they are the most location dependent activities that we could log. We only considered groups where each person contributed towards the activity being measured. Figure 7 gives an example of the location of New Daisys on a Daisyfield (each circle represents a new Daisy – some may have been deleted in the collaboration) with high CM on the left, and low CM on the right. Figure 8 shows a visual trace of the total pointer movements for each group (each colour represents a different user regardless of the experimental condition). Visually inspecting this figure confirms the CM results – participants with No Colours appear to focus their interaction predominantly on their own Daisys, not moving to other participants' Daisys in other parts of the screen. Furthermore, users' focus of interaction tends to remain in one area with No Colours, implying that they do not move their Daisys to other parts of the screen.

Activity	CM No Colours	CM Colours	U_A	z	p
Annotation	0.698	0.763	197	-1.67	0.0475
New daisy	0.665	0.719	332	-1.67	0.0475
Pointer movement	0.735	0.799	176	2.01	0.0222

Table 2: Colocation Measures for Annotation and New Daisy activities

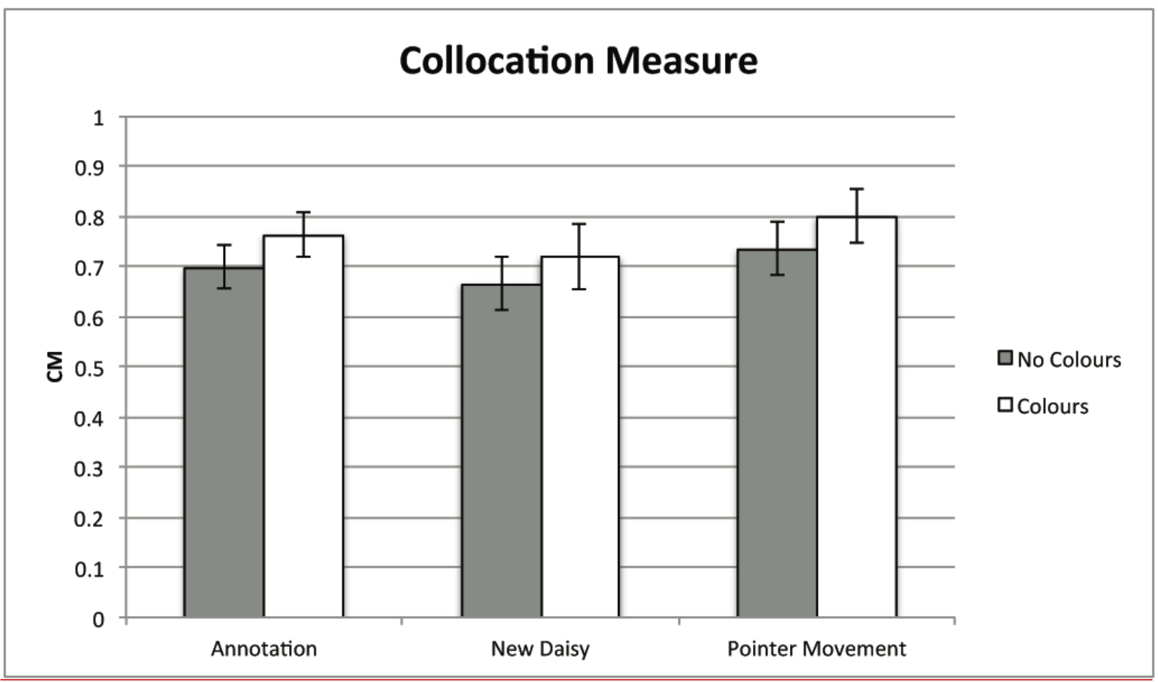

Figure 6: Collocation Measures for Annotation, New Daisys, and Pointer Movement

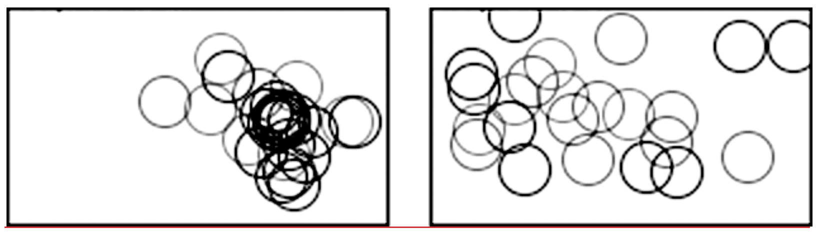

Figure 7: Example of Collocation of New Daisys – a) High Collocation, b) Low Collocation

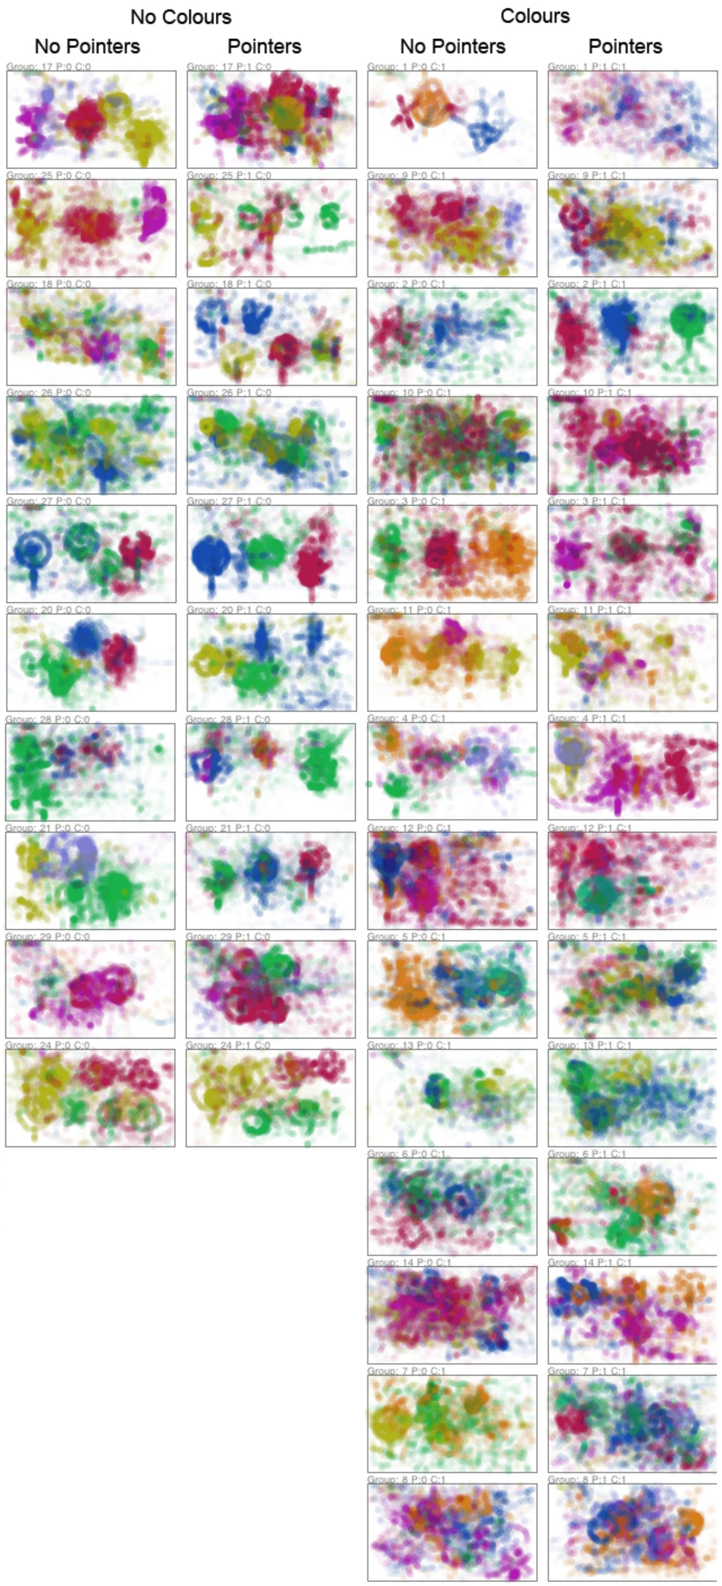

Figure 8: Overview of pointer movements for each group

Shared Pointers had no significant effect on collocation of Annotations, New Daisys, or Pointer Movements.

6.2. Content Assessment

In this section we report on results for dependent variables focussing on the content of participants' interaction. The data was gathered from human analysis of playbacks of the logs of interaction.

Musicality measure

Two independent judges rated the final compositions using our musicality rating scale. Inter-judge reliability was tested using Cohen's Kappa and found to be good (0.77). With **Colours compositions** were judged to be **more musical** (mean rating = 3.60) than without Colours (mean rating = 2.92); Mann-Whitney test ($U_A = 369.5$, $z = -1.86$, $p = 0.0314$). Shared Pointers had significant secondary effect on ratings of musicality overall as illustrated in figure 9. With **Colours and shared Pointers participants had significantly lower rated compositions** (mean rating = 3.0357) than with Colours and No Pointers (mean rating = 4.1786) when tested using a Wilcoxon Signed-Rank Test ($p = 0.0018$, $W = 66$, $z = 2.91$). With **No Colours and shared Pointers participants had significantly higher rated compositions** (mean rating = 3.9) than with No Colours and No Pointers (mean rating = 1.95) when tested using a Wilcoxon Signed-Rank Test ($p = 0.0027$, $W = -55$, $z = -2.78$).

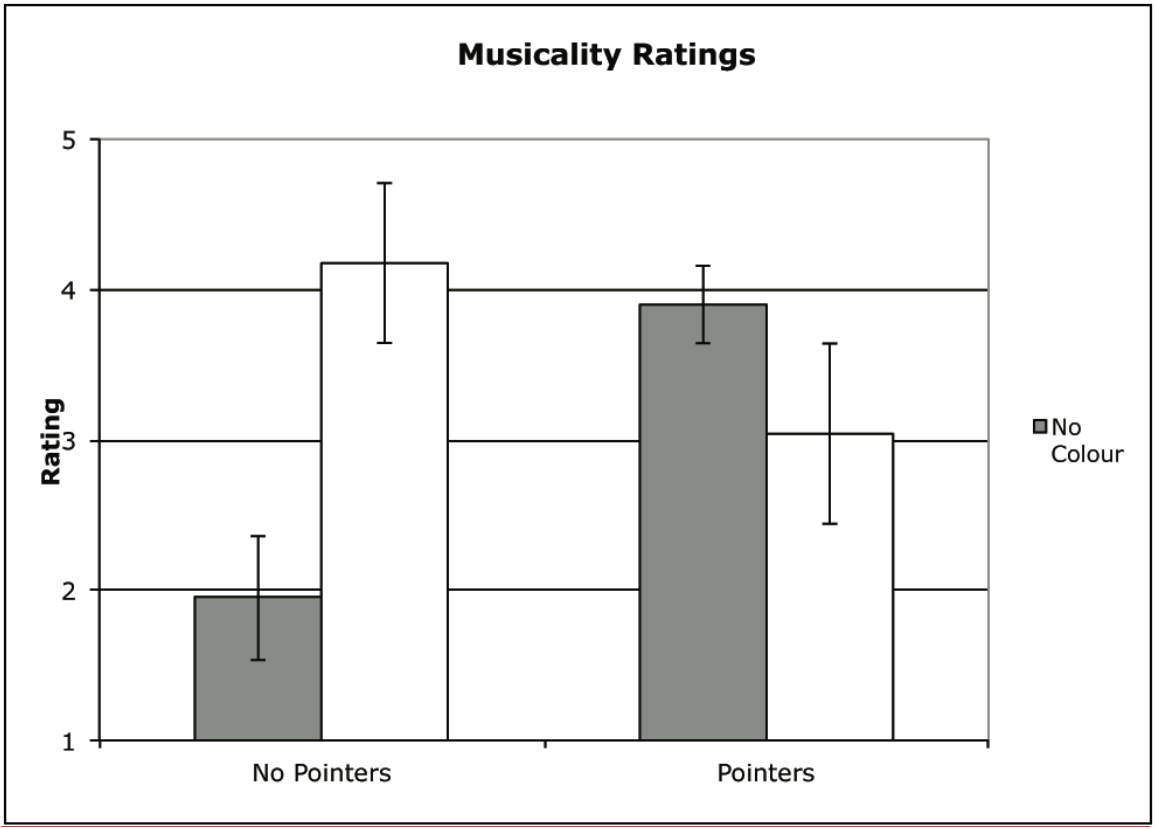

Figure 9: Ratings of musicality

Topics of Annotation

There were significantly **more contributions about Presence/ Identity with Colour** than without Colour (Mann-Whitney test; $U_A = 93.5$, $z = -2.18$, $p = 0.0146$), as illustrated in figure 10. There were significantly **more textual annotations with Colour** (mean = 17.6) than with No Colour (mean = 8.2) when tested using a Mann-Whitney test ($U_A = 90.5$, $z = -1.98$, $p = 0.0239$). None of the other differences were statistically significant. There were no statistically significant differences between topics of annotations with and without Pointers. There were no significant differences between the numbers of letters in each annotation (mean = 8.43), nor number of drawings per group (mean = 2.41).

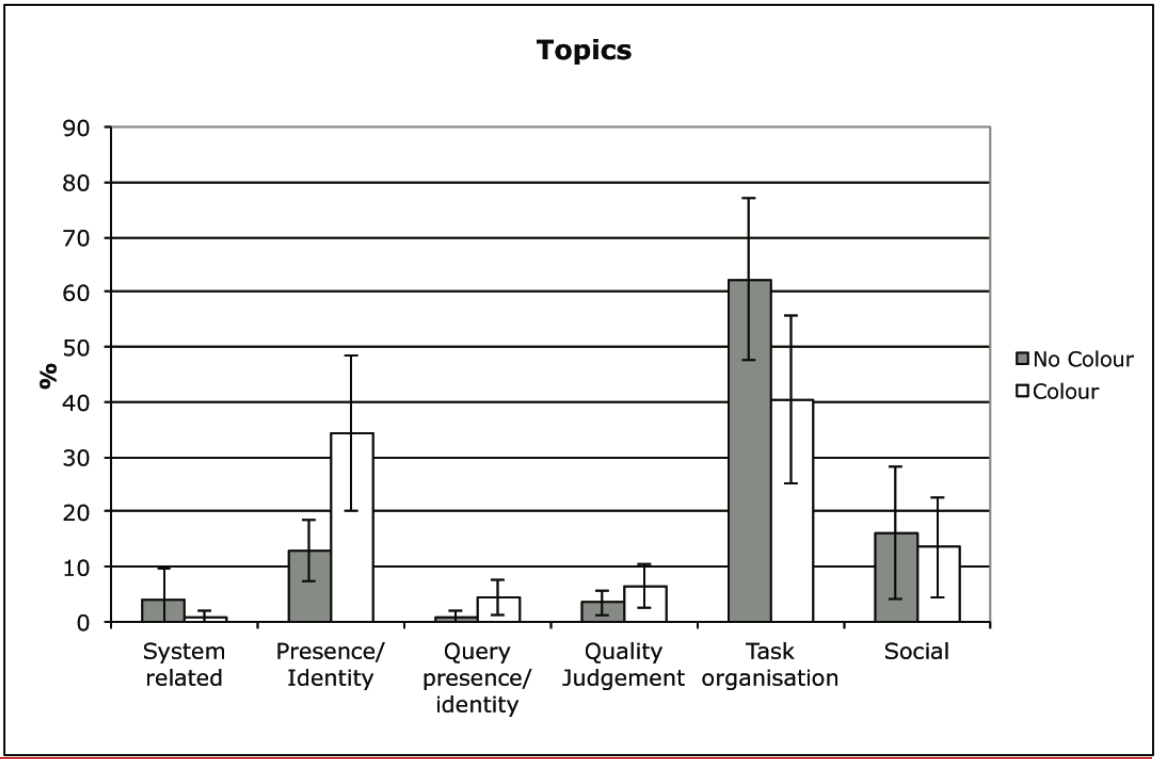

Figure 10: Percentage of topics of annotation with and without Colours

We identified four types of use of drawing in support of Task Organisation topics in nine groups as outlined below:

1) Dividing the space into three vertical sections - one space for each participant to put their Daisy(s) into as illustrated in figure 11 (observed in 4 groups).

2) Dividing the space from the centre, and then using one section per participant (observed in 1 group) e.g. figure 12.

3) Dividing the space into a work area and a discussion area as illustrated in figure 13 (observed in 2 groups).

4) Creating a space for role assignment and comments (observed in 2 groups) e.g. figure 14.

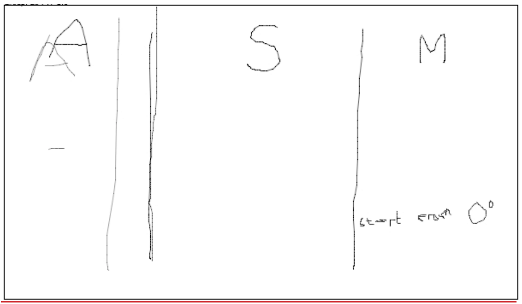

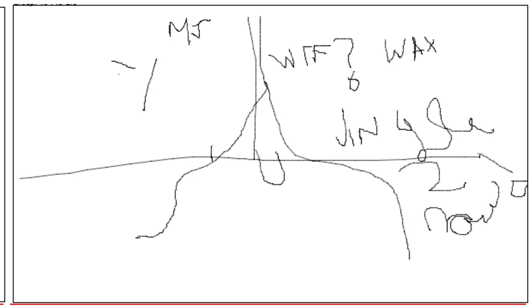

Figure 11: Vertical division of space. *Figure 12: Division of space from centre*

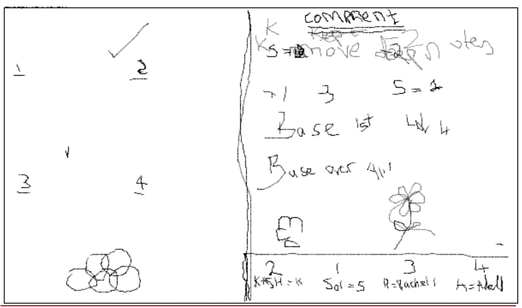

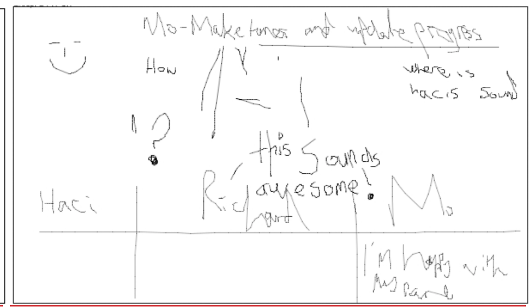

Figure 13: Division into work and discussion. Figure 14: Role assignment space.

We also identified 40 drawing (non-text) uses of the graphical annotations. There were no significant differences in the kinds of drawings used by different groups, so this reflects a group or individual's working style. Fourteen of these drawings were classified as Social topics (figurative drawings not connected with the activity such as flowers as illustrated in figure 13). Twenty six drawings were identified as arrows used to localise an associated textual annotation (Query Presence: 2, Task organisation: 24), as illustrated in figure 15.

Figure 15: Drawn arrow which associates text with musical loop

6.3. Participant Feedback

Table 3 and figure 16 detail the results of the participant questionnaires with significantly different results underlined with preference highlighted in bold using a Chi^2 test ($df = 1$; $p < 0.05$).

	No Colours, No Pointers	No Colours, Pointers	Colours, No Pointers	Colours, Pointers
The best Jingle	12	18	**<u>19</u>**	<u>7</u>
I felt most involved with the group	<u>8</u>	**<u>20</u>**	**<u>18</u>**	<u>7</u>
I enjoyment myself the most	12	18	17	9
I felt out of control	**<u>20</u>**	<u>8</u>	10	13
I understood what was going on	<u>6</u>	**<u>17</u>**	11	10
Other people ignored my contributions	10	9	<u>4</u>	**<u>14</u>**
The interface was most complex	14	9	11	9
The interface was frustrating	**<u>16</u>**	<u>5</u>	8	12
I felt satisfied with the result	8	18	16	7

Table 3: Questionnaire results

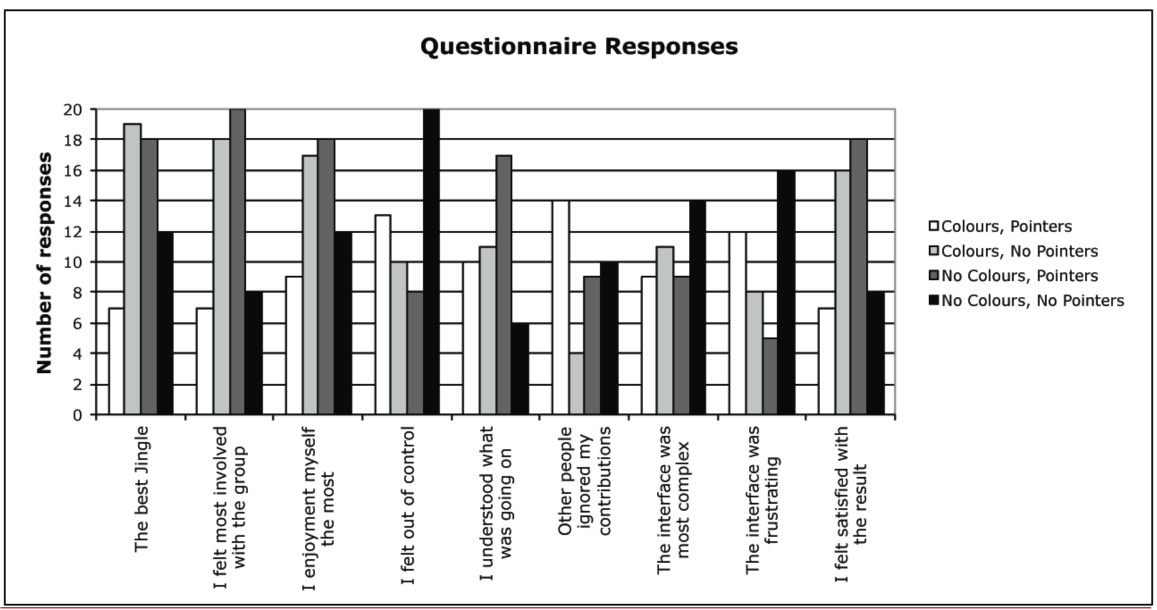

Figure 16: Questionnaire results

6.4. Summary of Results

To summarize, with **Colours** we found:

- **significantly more Annotations and Daisy Moves,**
- participants' **Annotations, New Daisys,** and **Pointer movements were significantly closer together** when comparing CM,
- significantly **more Annotations relating to Presence/ Identity,**
- significantly **more Textual Annotations,**
- significantly **better compositions.**

With **Pointers** we found that participants had significantly **fewer Remaining Notes** (i.e. less complex tunes).

Participants reported being more mutually engaged in the **Colours condition with No Pointers**, and in the **No Colours condition with Pointers**. This result was consistent with ratings of the musicality of the composition.

7. Discussion

Both Hypothesis H1 (Mutual engagement would be greater when participants had explicit cues to identity) and Hypothesis H2 (Mutual engagement would be greater when participants had indicators of other participants' loci of attention) are partially supported by our findings. H1 is supported by evidence that participants worked closer together with Colours, had better rated compositions, and made more annotations than with No Colours. H2 is supported by the finding that with Pointers participants had less complex final compositions. However, the participant feedback and ratings of musicality of compositions indicated that there is an interaction between Pointers and Colours - it seems that participants are most mutually engaged when they have either Colours or Pointers, but not both. A possible reason for this is that participants experience some form of **cognitive overload** - with both Pointers and Colours the shared representation simply had too many secondary properties distracting from the activity, i.e. there was too much non-task information presented on the screen (similar to the design issues with collaborative systems identified by Stefik et al., 1987), whereas with Pointers or Colours only, sufficient awareness information about others' activities was given for participants to mutually engage with each other. Having No Colours and No Pointers was not as mutually engaging, so there is a clear need for mutual awareness mechanisms.

The only significant reported differences in Interface frustration, Interface Complexity, and Understanding of what was going on were with No Colours - with No Colours, having No Pointers was frustrating, complex, and participants found it difficult to understand what was going on. It seems from this that participants were frustrated by lack of mutual awareness with No Colours and No Pointers. However, there was not a similar pattern when participants had Colours. This may be because there was sufficient awareness information just with the colours, but participants reported feeling that other people ignored their contributions when they had Pointers and Colours - maybe there was an expectation that others would take more notice of their work when they could see each others' loci of attention. Possibly these shared activity indicators lead to increased expectations on others - a form of social contract that is not satisfied by participants' actions. This may explain why Colours and Pointers are low rated as an experience.

Previous studies (Bryan-Kinns and Hamilton, 2009) showed that having different colours to indicate different participants increased their mutual engagement. By implication, our results additionally allow us to infer that mutual engagement is additionally indicated by significantly more collocation of Annotations, New Daisys, and Pointer movements. Our results also confirm Bryan-Kinns and Hamilton's finding (2009) that less complex final musical compositions indicate mutual engagement.

The results indicate that the presence or absence of Colours and Pointers affected different sets of dependent variables. In particular, Colours affected the amount of annotation, amount of movement of Daisys, and how collocated interaction was, whereas Pointers affected how complex the final musical compositions were. Furthermore, with Colours there were significantly more Annotations relating to Presence/ Identity, and more Annotations about Task Organisation (but this was not significant). One explanation for this difference could be that having shared pointers may provide an on-going reminder of the focus and activity of other users (typically editing the notes into a coherent form) and so there was more focus on the joint product, whereas conveying identity through colour encouraged participants to mark their presence in the shared representation and more explicitly organise their work. By implication it may be that more explicit user expressions of presence (e.g. participants writing their names on the shared representation) are an indication of mutual engagement between participants.

7.1. Secondary Properties of Shared representations

An interesting feature of Daisyfield is that the spatial position of Daisys in Daisyfield has no effect on the audio produced – location is a visual secondary property of the shared representation. So, any use of the spatial aspects of Daisyfield does not affect features of the music produced, but instead contributes to some other aspect of the interaction such as managing the process of music production, assignment of roles, social communication, and so on. Similarly, graphical annotations, colours of Daisys, and shared pointers are secondary properties which do not change the music per se, but are used to co-ordinate the activity of producing the musical composition and maintain awareness of the group's activities. In designing these secondary properties of shared representations it is important to consider the amount of screen real estate consumed by each feature. For instance, when not being edited, Daisys are 100 pixels in diameter, so twelve Daisys can easily fit into one Daisyfield. Furthermore, the size of Daisys allows for plenty of spatial separation between groups of Daisys if needed. Larger Daisys, or a smaller Daisyfield would have made grouping Daisys harder for participants, and our Collocation Measure would have produced less useful results. Similarly, the size of shared pointers were 40 pixels in diameter, but as discussed above, this may have been too prominent as they seemed to distract users when combined with Colours in the interface. Possibly reducing the size of the shared pointers, or making them more transparent would make them more acceptable to users (as demonstrated by increased participants' mutual engagement).

Tse et al. (2004) demonstrated that pairs of participants naturally spatially partition their work when undertaking co-drawing activities with no explicit indication of identity. We found similar behaviour when trios had No Colour – they tended to spatially lay their Daisys out away from each other. It would be interesting to re-run Tse et al.'s studies with different colours for different participants and examine whether this caused more collocated drawing as we found with Daisyfield. An important difference between Tse et al.'s studies and our study of Daisyfield is that in Tse et al.'s study graphical annotations were a primary property of the shared representation, whereas in Daisyfield annotations are secondary properties. It may be that colours and spatial layout have different effects on primary and secondary features of shared representations. Also, whilst we found similar result to Tse et al. for the No Colours condition, it is worth noting that Tse et al.'s studies were with collocated pairs of participants, whereas we examined remotely located trios participants which certainly has an effect on both the group dynamics (with pairs of participants each participant knows that if they have not done something, then the other person has done it whereas with trios it would not be obvious who had made which contributions) and available communication channels (collocated participants can aurally and gestural communicate using rich naturalistic channels). Furthermore, our participant activities were creative and audio in nature whereas Tse et al.'s were primarily focussed on the non-creative tasks of graphical tracing and reproduction. The similarity of the findings in significantly different participant situations suggests that our findings have potential applicability outside the domain of collaborative music making.

The Collocation Measure developed in this paper gives a numerical indication of how closely trios of participants work together. In particular, the CM is used in this paper to show that when participants are provided with colours to indicate authorship, they work closer together, or to put it another way, they rely less on spatial separation to distinguish their contributions from each other. It would be interesting to study whether the same finding held in non-creative and non-music domains as exemplified by Tse et al.'s work (2004).

The CM is itself an advance on the measures developed by Tse et al. (2004) which only consider pairs of participants. It also improves on measures by Pinelle et al. (2008) who consider physically collocated participants with static points of origin. However, the CM may not be scalable to larger groups than trios as it simply considers the length of the perimeter joining participants' mean points of interaction. A more sophisticated approach may need to be developed to handle groups with four of more participants, for instance, by considering bounding shapes of interaction, and examining the intersection of bounding shapes. Future research will explore this issue.

The underlying assumption that pointer movement provides an indication participants' loci of attention also needs to be further explored and unpacked in future work. In this paper we have assumed that collocated pointer movements indicate mutual engagement with the musical aspects of the interaction, but participants could just be moving their pointers together for fun, or some other reason. Conducting comparative experiments with and without any shared visual representations may give us some insight into whether point location correlates with loci of task attention.

7.2. Users' Access and Understanding of the Shared representation

Line drawings were used by nine of the groups to manage the Task Oragnisation of their music composition – essentially using the secondary visual properties of the shared representation to help manage their access and manipulation of the primary musical properties of the shared representation. Two kinds of use of drawings can be identified from the four kinds of diagrams reported in the results: 1) division of screen real estate to mark spaces for specific participants (a form of emergent access control), 2) division of screen real estate for specific activities such as making comments versus creating music (a mechanism to manage the collaborative editing process). These uses of drawings to mark spaces for specific activities is similar to the findings of Scott et al. (2004) and Kruger et al. (2005) who found that participants using interactive tables tended to divide the interaction space into personal, group, and storage territories. However, as Scott et al. and Kruger et al. were examining the use of tabletops, private space was dictated by participants' physical orientation around the tabletop whereas Daisyfield participants have no fixed point of origin. It is particularly noteworthy that the use of drawings was not affected by Pointers or Colours, but reflected groups' working styles. It would be interesting to explore how such groups would react to a shared representation which did not allow drawing (e.g. which had text only communication rather than shared graphical annotations). It may be that in such situations participants would assign areas of the screen using words such as 'top-left' or 'centre' etc., and may even develop their own text tags to indicate different kinds of comment much like users of twitter create their own hashtags. Furthermore, the shared annotation in Daisyfield is persistent – that is, annotations cannot be deleted (e.g. with an eraser) and do not disappear over time (e.g. as text would in a scrolling text chat window). The persistent display of annotation history provides a stable shared context, and allows participants to make localisation marks, but may be difficult to use or maintain over extended periods of time as the annotation space becomes full of marks and potentially confusing.

Examining the use of the secondary properties of Daisyfield's shared representation indicates that providing spatial positioning and graphical annotation helps users to understand and manage the information for the activity (shared loops of music in this example). Maintenance of shared knowledge about the activity and roles became more explicit when participants were assigned unique colours to identify their contributions – there were more annotations expressing identity and relating to task management when participants had Colours. In addition, the flexible nature of the shared graphical representation allowed participants to localise their annotations – to write some text and then draw an arrow to indicate what they were referring to. This feature of shared graphical annotation is a powerful mechanism allowing participants to easily make specific references to objects and parts of objects in their collaborative activity. All our participants were novices with the Daisyfield system, but it would be fascinating to examine how the shared representation supported learning in groups of users, and sharing of expertise.

7.3. Implications for Design

This study showed that providing shared awareness mechanisms in conjunction with free-form spatial communication mechanisms significantly improved participants' mutual engagement. This has direct implications for the design of similar collaborative music systems in fields such as NIME (Poupyrev et al. 2001). The measures of mutual engagement and collocation also contribute to the evaluation of engagement, and group expressivity with such systems (cf. Dobrian and Koppelman 2006). In a broader sense, the role of shared graphical annotations has implications for the design of CSCW systems in general, especially for creative collaboration (cf.

Benford et al., 2000). Our findings also point to potential problems when designing and implementing shared awareness mechanisms – the potential for participant cognitive overload. Careful consideration therefore needs to be given to the prominence of awareness displays, and their potential to interfere with other elements of the display. Our recommendations from this study for the design of distributed collaborative creativity systems are outlined below and reflect guidelines proposed in CSCW awareness research such as Gutwin and Greenberg (2002):

- Provide secondary properties which indicate the **identity** of participants. This could be as simple as assigning each participant a unique colour as we have done in this paper, or assigning textual names, or photos of participants, or avatars. However, as outlined below, increasing the richness of the identity indicators may actually cause information overload.

- Provide indicators of loci of **attention** of participants, for example, using shared pointers as we have exemplified in this paper to provide moment-by-moment indication of participants' current pointer location. Other indicators might include emphasising loci of attention over time - highlighting where participants are most active, or highlighting where participants are working together most, for example using heat maps of activity to indicate where most attention is spent. Comparing moment-by-moment indicators to time-lapse indicators would be a fruitful avenue for future research on designing effective secondary properties of shared representations.

- Support participant driven **grouping** of elements of the product. In Daisyfield spatial layout was a secondary property of the shared representation which allowed participants to visually arrange elements of the shared music product as they saw fit without affecting the music itself. Other grouping strategies could be compared and contrasted to identify effective use of space as a secondary property, for example, restricting participants to certain areas of the shared representation, or automatically visually aligning similar contributions, or allowing participants to hide contributions from each other.

- Support flexible and localisable shared **annotations** as a secondary property of the shared representation. We found that allowing people to draw around the musical compositions, and to point to parts of the compositions, increased their mutual engagement. However, drawing can be a time consuming activity, especially with a mouse, so it would be useful to examine the effect on mutual engagement of providing support for textual annotations, for example as text chat, or placed in the shared space, with or without graphical annotations.

- **Minimise the amount of non-task information** displayed to users. Having a shared representation with too many secondary properties such awareness information can be detrimental to mutual engagement – secondary properties should be designed to provide scaffolding for the collaboration, and not be too intrusive. This clearly creates a design tension between primary and secondary properties of shared representations which needs to be researched further.

From our study we propose that the following indicate increased mutual engagement between participants, and so can be used to help identify which design features actually contribute to enjoyable creative collaborations:

- More **collocated** activity. The study showed that when participants grouped their annotations, and music contributions together, or moved their pointers in the same areas, they were more mutually engaged.

- More **coherent** and better quality final products. The quality of the final product reflects the level of mutual engagement between participants, even for tasks where participants may have little or no experience of the domain (in this case, music making).

- More use of secondary properties of the shared representation to take **ownership** of the product. Our analysis of annotations indicated that when participants were mutually engaged they wrote their names and checked the presence of others, indicating a sense of ownership with the product, and focus on collaboration rather than individual activity.

- User **preference** (identified using the Mutual Engagement Questionnaire). The key to the success of the MEQ is its use of comparison between interfaces rather than subjective ratings of interfaces which we have found do not provide a strong enough indicator of mutual engagement.

8. Conclusions

In this paper we showed that providing collaborators undertaking creative activities with shared representations that include secondary properties such as mutual awareness indicators increased their mutual engagement. In doing so we also highlighted the difficulty of designing secondary properties which do not overload the user and lead to decreased mutual engagement. The role of flexible and usable shared space in creative collaborations was demonstrated through support for shared graphical annotations and spatial layout which we argued helped users to mediate their interaction with the shared representations. Future work will explore larger groups of collaborators and different creative activities. We are particularly keen to further develop our measures of mutual engagement, especially those concerned with spatial interaction in shared representations.

9. Acknowledgements

This work was partly funded by EPSRC grant EP/H042865/1 (Collaborative Cross Modal

Interfaces).

10. References

Abbott, K. R., and Sarin, S. K. (1994). Experiences with workflow management: issues for the next generation. In *Proceedings of CSCW '94*. ACM, New York, NY, USA, pp. 113-120.

Applegate, L., Konsynski, B.R., and Nunamaker, J.F. (1986). A group decision support system for idea generation and issue analysis in organization planning. In *Proceedings of CHI '86*. Austin, TX: ACM, pp. 16–34.

Bannon, L., and Bødker, S. (1997). Constructing common information spaces. In Proceedings of European Conference on Computer-Supported Cooperative Work (ECSCW'97), Norwell, MA: Kluwer Academic Publishers, pp. 81-96.

Barbosa, A. (2003). Displaced soundscapes: a survey of network systems for music and sonic art creation. *Leonardo Music Journal, 13*, pp. 53–59.

Benford, S., Reynard, G., Greenhalgh, C., Snowdon, D., and Bullock, A. (2000). A poetry performance in a collaborative virtual environment. *IEEE Computer Graphic Applications 20 (3)*, pp. 66–75.

Bly, S. A. (1988). A Use of Drawing Surfaces in Different Collaborative Settings. In *Proceedings of CSCW'88,* pp. 250-256.

Boden, M. A. (1992) *The Creative Mind*. London: Abacus.

Bryan-Kinns, N. (2004). Daisyphone: The Design and Impact of a Novel Environment for Remote Group Music Improvisation. In *Proceedings of DIS 2004*, Boston, USA, pp. 135-144.

Bryan-Kinns, N., & Hamilton, F. (2009). Identifying Mutual Engagement. *Behaviour & Information Technology*. DOI: 10.1080/01449290903377103.

Cherubini, M., Venolia, G., DeLine, R., and Ko, A. J. (2007). Let's go to the whiteboard: how and why software developers use drawings. In *Proceedings of CHI '07*. ACM, New York, NY, USA, pp. 557-566.

Clark, H.H., 1996. *Using language*. UK: Cambridge University Press.

Clark, H.H. and Brennan, S.E., 1991. Grounding in communication. In: L.B. Resnick, J. Levine, and S.D. Behrend, eds. *Perspectives on socially shared cognition*. Washington, DC: American Psychological Association, pp. 127–149.

Cornock, S., and Edmonds, E. (1973). The Creative Process Where the Artist Is Amplified or Superseded by the Computer. *Leonardo, 6 (1),* pp. 11-16.

Csikszentmihalyi, M. (1991). *Flow: the psychology of optimal experience*. New York: Harper Collins.

Dobrian, C. and Koppelman, D. (2006). The 'E' in NIME: musical expression with new computer interfaces. In *Proceedings of new interfaces for musical expression (NIME).* Paris, France: IRCAM, Centre Pompidou, pp. 277–282.

Dourish, P. and Bellotti, V. (1992). Awareness and coordination in shared workspaces. In *Proceedings of ACM conference on computer-supported cooperative work (CSCW092), Toronto, Canada*, pp. 107–114.

Fischer, F., Bruhn, J., Gräsel, C., and Mandl, H. (2002). Fostering collaborative knowledge construction with visualization tools. *Learning and Instruction*, 12, pp. 213–232.

Flor, N. V. and Maglio, P. P. (1997). Emergent Global Cueing of Local Activity: Covering in Music. In *CSLC Proceedings*, pp. 45-52.

Freeman, W. J. (2000). A neurobiological role of music in social bonding. In (eds) Wallin, N., Merkur, B., and Brown, S. *The Origins of Music*, Chapter 22, Cambridge MA: MIT Press, pp. 411-424.

Gutwin, C. and Greenberg, S. (2002). A descriptive framework of workspace awareness for real-time groupware. *Computer Supported Cooperative Work, 11*, pp. 411–446.

Healey, P. G. T., Leach, J., and Bryan-Kinns, N. (2005). Inter-Play: Understanding Group Music Improvisation as a Form of Everyday Interaction. In *Proceedings of Less is More — Simple Computing in an Age of Complexity*, Microsoft Research Cambridge.

Hutchins, E. (1995). How a Cockpit Remembers Its Speeds. *Cognitive Science, 19 (3)*, pp. 265–288.

Ishii, H., Kobayashi, M., and Grudin, J. (1993). Integration of interpersonal space and shared workspace: ClearBoard design and experiments. *ACM TOIS, 11 (4),* pp. 349-375.

Kraut, R. E., Gergle, D., and Fussell, S. R. (2002). The Use of Visual Information in Shared Visual Spaces: Informing the Development of Virtual Co-Presence. In *Proceedings of CSCW'02*, New Orleans, Louisiana, USA, pp. 31-40.

Kruger, R., Carpendale, M. S. T., Scott, S. and Greenberg, S. (2004). Roles of orientation in tabletop collaboration: Comprehension, coordination and communication. *Journal of CSCW, 13(5-6)*, pp. 501-537.

Malone, T. (1983). How do people organize their desks?: Implications for the design of office information systems. *ACM Transactions on Office Information Systems, 1(1),* pp. 99–112.

Mamykina, L., Candy, L., and Edmonds, E. (2002). Collaborative creativity. *Communications of the ACM 45 (10),* pp. 96-99.

Miletto, E. M., Flores, L. V., Rutily, J., and Pimenta, M. S. (2006). CODES: Supporting Awareness in a Web-based Environment for Collective Music Prototyping. In *Proceedings of IHC 2006, Brazil*, pp. 50-55.

Nabavian, S., and Bryan-Kinns, N. (2006). Analysing Group Creativity: A Distributed Cognitive Study of Joint Music Composition. In *Proceedings of Cognitive Science*, pp. 1856-1861.

Olson, J. S., Olson, G. M., Storrøsten, M., and Carter, M. (1993). Groupwork close up: a comparison of the group design process with and without a simple group editor. *ACM Transactions on Information Systems, 11 (4),* pp. 321–348.

Pinelle, D., Gutwin, C., and Greenberg, S. (2003). Task Analysis for Groupware Usability Evaluation: Modeling Shared Workspace Tasks with the Mechanics of Collaboration. *ACM Transactions on Computer-Human Interaction, 10(4)*, pp. 281-311.

Pinelle, D., Nacenta, M., Gutwin, C., and Stach, T. (2008). The Effects of Co-Present Embodiments on Awareness and Collaboration in Tabletop Groupware. In *Proceedings of Graphics Interface Conference 2008, Windsor, Canada*, pp. 1-8.

Poupyrev, I., Lyons, M. J., Fels, S., and Blaine, T. (2001). New interfaces for musical expression. In *Proceedings of CHI '01 extended abstracts on Human factors in computing systems*, Seattle, USA. New York, USA: ACM, pp. 491–492.

Sawchuk, A. A., Chew, E., Zimmermann, R., Papadopoulos, C., and Kyriakakis, C. (2003). From

Remote Media Immersion to Distributed Immersive Performance. In *Proceedings of ETP '03, Berkeley, California, USA*.

Sawyer, K. R. (2003). *Group Creativity: Music, Theater, Collaboration*. New Jersey: LEA.

Schafer, W. A., and Bowman, D. A. (2004). Evaluating the effects of frame of reference on spatial collaboration. *Virtual Reality, 7*, pp. 164–174.

Scott, S. D., Grant, K. D., and Mandryk, R. L. (2003). System Guidelines for Co-located, Collaborative Work on a Tabletop Display. In *Proceedings of European Conference Computer-Supported Cooperative Work (ECSCW 2003), Helsinki, Finland*.

Scott, S. D., Sheelagh, M., Carpendale, T., and Inkpen, K. M. (2004). Territoriality in Collaborative Tabletop Workspaces. In *Proceedings of CSCW'04, Chicago, Illinois, USA*, pp. 294-303.

Sheridan, J. G. and Bryan-Kinns, N. (2008). Designing for Performative Tangible Interaction. *International Journal of Arts and Technology. Special Issue on Tangible and Embedded Interaction, 1 (3-4),* pp. 288-308.

Shneiderman, B. (2007). Creativity support tools: accelerating discovery and innovation. *Communications of the ACM, 50 (12),* pp. 20-32.

Stefik, M., Bobrow, D. G., Foster, G., Lanning, S., and Tatar, D. (1987). WYSIWIS revised: early experiences with multiuser interfaces. *ACM Transactions on Office Information Systems 5 (2),* pp. 147–167.

Stewart, J., Bederson, B. B., and Druin, A. (1999). Single display groupware: a model for co-present collaboration. In *Proceedings of the CHI '99*. ACM, New York, NY, USA, pp. 286-293.

Stowell, D., Robertson, A., Bryan-Kinns, N., and Plumbley, M. D. (2009). Evaluation of live human-computer music-making: quantitative and qualitative approaches. *International Journal of Human–Computer Studies, 67*, pp. 960–975.

Suthers, D., Girardeau, L., and Hundhausen, C. (2003). Deictic Roles Of External Representations in Face-To-Face And Online Collaboration. In *Designing for Change*, Kluwer Academic Publishers, Amsterdam, The Nederlands, pp. 173–182.

Suwa, M., and Tversky, B. (1997). What do architects and students perceive in their design sketches? A protocol analysis, *Design Studies*, 18 (4), October 1997, pp. 385-403.

Tse, E., Histon, J., Scott, S. D., and Greenberg, S. (2004). Avoiding Interference: How People Use Spatial Separation and Partitioning in SDG Workspaces. In *Proceedings of CSCW'04*.

Tse, E., Shen, C., Greenberg, S., and Forline, C. (2006). Enabling Interaction with Single User Applications through Speech and Gestures on a Multi-User Tabletop. In *Proceedings of AVI '06, Venezia, Italy*, pp. 336-343.