

Mechanisms for collaboration: a design and evaluation framework for multi-user interfaces

Article (Accepted Version)

Yuill, Nicola and Rogers, Yvonne (2012) Mechanisms for collaboration: a design and evaluation framework for multi-user interfaces. *ACM Transactions on Computer-Human Interaction*, 19 (1). pp. 1-25. ISSN 1073-0516

This version is available from Sussex Research Online: <http://sro.sussex.ac.uk/38760/>

This document is made available in accordance with publisher policies and may differ from the published version or from the version of record. If you wish to cite this item you are advised to consult the publisher's version. Please see the URL above for details on accessing the published version.

Copyright and reuse:

Sussex Research Online is a digital repository of the research output of the University.

Copyright and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable, the material made available in SRO has been checked for eligibility before being made available.

Copies of full text items generally can be reproduced, displayed or performed and given to third parties in any format or medium for personal research or study, educational, or not-for-profit purposes without prior permission or charge, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Mechanisms for Collaboration: A Design and Evaluation Framework for Multi-User Interfaces

Nicola Yuill¹ and Yvonne Rogers²

¹ School of Psychology
University of Sussex
Brighton
BN1 9QH
nicolay@sussex.ac.uk

² Interaction Centre (UCLIC)
University College London
London
WC1E 6BT
y.rogers@ucl.ac.uk

Corresponding author: Nicola Yuill
Tel: 44 1273 678630

Abstract

Multi-user interfaces are said to provide ‘natural’ interaction in supporting collaboration, compared to individual and non-co-located technologies. We identify 3 mechanisms accounting for the success of such interfaces: high *awareness* of others’ actions and intentions, high *control* over the interface and high *availability of background information*. We challenge the idea that interaction over such interfaces is necessarily ‘natural’ and argue that everyday interaction involves *constraints* on awareness, control and availability. These constraints help people interact more smoothly. We draw from social developmental psychology to characterize the design of multi-user interfaces in terms of how constraints on these mechanisms can be best used to promote collaboration. We use this framework of mechanisms and constraints to explain the successes and failures of existing designs, then apply it to three case studies of design, and finally derive from them a set of questions to consider when designing and analysing multi-user interfaces for collaboration.

Categories and Subject Descriptors: H5.2 [Information Interfaces and Presentation]: User Interfaces—User-centred design; Interaction styles; Theory and methods; General Terms: Design, Human Factors

Additional Key Words and Phrases: collaboration, constraints, multi-user interfaces

1. Introduction

Multi-user interfaces have become pervasive in the last few years, with a number of commercial and customized systems now available. These include Diamond Touch (Dietz & Leigh, 2001), Microsoft’s Surface, Smart Table and tangibles with multi-user interfaces such as the Reactable (Jorda, 2003). The lightweight and parallel action of touching, the mobility of users, and the increased ability for natural expressions of behaviour such as gesture and posture extend the possibilities for communication and collaboration. User studies have shown how groups of people new to such interfaces find enjoyment in sharing and assembling sets of digital images for a variety of collaborative tasks (Huang & Mynatt, 2003; Ryall et al., 2004; Scott et al., 2003b; Shen et al., 2002). Commercially available multi-touch tabletops, such as the Smart Table are being promoted in terms of how they can help people

work together: 'young students are drawn to its surface, where work and play come together in hands-on, collaborative activities' (www.elementarytechnology.co.uk). The claim is that these new kinds of shareable interfaces provide more opportunities for flexible forms of collaboration compared for example with single-user PCs and mouse input, through allowing co-located users to interact smoothly and simultaneously with digital content.

It has long been recognised that group work can have benefits ('many hands make light work') but also that 'too many cooks spoil the broth': that is, group work does not necessarily gain from the simple sum of each individual's contribution, but there can be 'process loss' due to lack of coordination (Steiner, 1972). The advent of Single Display Groupware (SDG) led to insightful analyses of how such technologies might engender interference e.g. parallel working and conflicts (Stewart et al., 1999), followed by technical developments such as identity-differentiating widgets (Ryall et al., 2005), encouraging designers to think about how identity information might be used to support interleaving of actions by different users. The further development of shareable interface technology has led to a wealth of ideas about the benefits of these new capabilities. For example, Kharrufa et al. (2009, p. 9) list the special features of shared interfaces as 'encouraging externalisation, ...[providing a] rich set of cognitive tools, providing structure to a task, designing for different ability levels and supporting reflection' through structure and logging tools. Rogers and Rodden (2003) propose responses to a cogent set of uncertainties and tensions that arise in shared displays while Scott et al. (2003a) cite the need to support natural interpersonal interactions, smooth transitions and simultaneous actions with tabletops. However, such guidelines, invaluable though they are, tend to give us the 'what' rather than the how and why. The latter require a more psychological perspective, explaining how it is that people manage to coordinate interactions in everyday social interactions. An especially illuminating way to do this is to look at the emergence of collaboration, a thriving area in developmental psychology. Tomasello (2009) premises the distinctive human ability to collaborate on our evolved disposition for shared intentionality, which is based on attention, background knowledge and a drive to cooperate. In this paper we translate these properties into three mechanisms of interaction that can be used to inform the design of shareable interfaces. We can understand how best to design and use such interfaces if we are clear about what it is that gives them their

apparently natural and motivating properties in supporting interaction, and what makes them an apparently superior means of collaboration, compared for example to paper, remote sharing or individual devices. We therefore draw on examples from our own research with young children learning to collaborate and with groups who have difficulties with collaboration, to produce a framework that recognises some of the challenges of collaboration that may be hidden when studying only competent collaboration, and when observing users who recover quickly and silently from breakdowns of collaboration (Easterbrook, 1996).

New multi-user interfaces represent a qualitative shift in supporting collaborative groupwork: the freedom of input enables gesturing, speaking, and touching. These can all be seen, heard and experienced by *others* whereas mouse clicking and key pressing are individual, private acts. While such actions may become largely invisible to those executing them, as they are so familiar, their enaction, in contrast, remains visible to others. Developers of surface technologies and tangibles have been inventive in logging and presenting information about users' interactions, in the hope of improving people's interactions, e.g. making them more equitable (e.g. diMicco et al., 2007, Bachour et al., 2008). Surfaces, tangibles and shareable public displays also enable simultaneous control by multiple users. These technologies therefore provide striking new opportunities for *awareness* -- gesture, body orientation and more so-called 'natural' means of communication, for making salient in displays the *availability* of information supported by the public space provided, and for equitable simultaneous *control*, such as 'entry points' (Rogers et al., 2008) to the technology. In this typology, collaboration might look easy. But if it is so natural, why is it that not every such application is a success? For example, why does increasing awareness of spoken contributions sometimes lead to lower productivity and lack of equitable participation (di Micco et al., 2007)? Why might users show *less* equal contributions with multi-touch applications than with single-touch (Harris et al., 2009)? And do the advantages of such interfaces compensate for the loss of feedback or trace of interactions provided in traditional GUIs (Norman, 2010)?

The answers to these questions are needed to understand, support and explain the effective design of new user interfaces for multiple users (e.g. with tabletops, public displays, tangibles), that can capitalize on the mechanisms we have identified above to facilitate new forms of collaboration, e.g., in decision-making, brainstorming,

planning, learning and creating. In what ways can an interface exploit people's implicit social behaviours for good social effect? Can they be designed to make it more obvious as to how to behave and give appropriate cues? Is it possible to make more salient the cues that can lead to improved understanding, explication of intentions and focus of attention? Can people become more aware of what is good collaboration and conduct through interacting with a given interface? Our goal is to provide a principled way for researchers and designers to be able to make sense of the emerging empirical literature on the benefits of multi-user interfaces, and to enable them to make predictions about their value for collaboration. We present the mechanisms framework to inform the design of interfaces for collaboration through considering social-psychological properties of shared intentionality. These properties can explain why a multi-touch surface that enables simultaneous interaction may be less effective at facilitating collaboration than a single-touch interface, why gesture-based interactions might be better than mouse-based or tangible interactions for collaborative learning and why apparently minor increases in salience of users' actions can have big effects on the connectedness of social interaction.

We first define the three mechanisms, awareness, control and availability of information, based on the psychological literature on processes of social development, and we illustrate them in the light of existing literature on tabletops designed for collaborative work. We then explain the crucial role of constraints on these mechanisms, using the psychological literature on social development to show how constraints support smooth social interaction, and identifying three sources from which constraints can arise, again using examples from the literature on tabletops. We next present three case studies illustrating how the framework of mechanisms and constraints can be used to design for and evaluate collaboration. We finish by showing how the framework could be used by designers of collaborative applications on multi-user interfaces, for a wide range of audiences, from family groups at leisure, through children working collaboratively in formal settings to remedial work supporting people with social, sensory and language impairments.

2. Background: Awareness, control and availability

Much of the research into supporting co-located collaboration has focused on technological approaches to understand group working. We take a different approach, by looking at the psychological literature on the foundations of social interactions and their development, thus focusing on the underlying mechanisms of collaboration in simple everyday interaction. This makes our framework particularly inclusive, and this is reflected in the work that we draw on, including designs to support collaboration not only in competent and mature interactants but also in those facing barriers to collaboration. Technology has a particularly important social role here in supporting collaboration in the broadest range of users. Using the literature on the social development of collaboration, we have abstracted three core mechanisms of behaviour that underlie interactions of users doing collaborative tasks: awareness, control and availability. More specifically, they refer to:

- (i) **Awareness of others** - the degree to which awareness of users' ongoing actions and intentions is present or made visible moment-to-moment
- (ii) **Control of action** - the extent of each user's control over actions and decisions
- (iii) **Availability of information** - the ways in which background information relevant to users' behaviour and to the task is made available, or externalised.

2.1 Awareness

Awareness is a term used widely in the literature on shared workspaces, for example in the influential work of Gutwin and Greenberg (2002) and in the recognition of situational awareness when coordinating action, notably by Heath and Luff (1992). We use it here to mean the extent to which people, when they interact with new user interfaces that involve seeing, touching and gesturing, have ongoing awareness of the actions, intentions, emotions and other mental states of other interactants. Schmidt (2002) defines awareness as reciprocal practices of monitoring others and designing actions so as to render visible certain aspects of activity. Much of this awareness is implicit but a powerful influence on behaviour. For example, the difference between a driver speaking to a passenger in a car and to a conversant on the other end of a phone is apparent in the smoothness of conversational turns and the responsiveness to external events, such as busy traffic (Haigney and Westermann, 2001).

Several studies have identified subtle ways in which people show awareness in multi-user technologies. For example, Hornecker et al. (2008) found that users display several different signs of awareness when using multi-touch tables: they make running commentaries on their own actions, they anticipate collisions by adjusting their position, and they sometimes elbow others out of the way. These implicit mechanisms of awareness play a clear role in supporting collaboration with multi-user interfaces. In another study, Fleck et al. (2009) observed groups of 3 children completing a classroom seating arrangement on a multi-touch table, and found that completely isolated, individual working was rare and short-lived. Children made their actions and thoughts known by, for example, talking through and demonstrating their ideas at the same time, and their partners were willing to act as audience at such moments. There were also less planned ways in which awareness became shared, when, for example a child would reach across the table, obstructing another child's view, or when users thought aloud about what they were doing. Such actions did not necessarily lead to a response, but kept the group mutually informed about actions and plans.

There is good evidence that multi-user touch surfaces enhance awareness compared to other input devices. Ha et al. (2006) showed that participants had better awareness of others' actions with touch or a stylus rather than a mouse as input device. Hornecker et al. (2008) compared mouse and touch input for collaborating on a planning task and found that awareness of others' activity was increased in the touch condition.

Shared surfaces support more centrally-focused mutual awareness of actions through shared focal attention, as well as peripheral awareness, as for example when digital objects become the focus of shared planning or discussion. Multi-touch surfaces can be considered tangible in the sense that they lure people standing around them to manipulate the digital objects represented on them (Rogers, Lim and Hazlewood, 2006). Where they are used with tangibles (e.g. Falcao and Price, 2009) the physical artifacts support additional cues to awareness. Group members may use them as external thinking props to explain a principle, an idea or a plan to the others that is more effective than using equivalent digital representations (Brereton and McGarry, 2000; Fitzmaurice et al., 1995; Fjeld, 2002; Hornecker and Buur, 2006). In particular,

the act of waving or holding up a physical object in front of others is very effective at commanding attention. The persistence of and ability to manipulate physical artifacts may also result in more options being explored in a group setting (Fernaes and Tholander, 2006), and have been hypothesised to increase peripheral awareness of others' activities, helping collaborators gain a better overview of the group activity (e.g., Scott et al., 2003a).

On the face of it, increasing awareness seems like a strongly positive and distinctive feature of multi-user surfaces, compared with individual devices. However, the power of technology to present large amounts of information and a high degree of simultaneous action by multiple users risks cognitive overload. Multi-user interfaces involve a whole new set of cues about objects to which users direct their actions or intentions. Users of a range of tabletop applications have been found to urge their partners to 'wait' or 'slow down' because they could not grasp all the current action (Harris et al., 2009). Thus, multi-user interfaces offer new possibilities for awareness of ongoing actions, but this awareness is not an unmitigated blessing – more awareness is not necessarily better. It is clear that designers need to think about what sorts of awareness are afforded by multi-user interfaces, and how these might need to be limited by the constraints that we introduce below.

2.2 Control of action

One of the most compelling features of multi-user surfaces is the way in which they give varied opportunities for control by users, e.g. the directness of touch, the flexibility of tangibles on surfaces and the potential for individually-customisable control. By control, we refer to the ways in which users can effect changes in actions within the system and hence decisions within the group. Many applications have focused on the benefit of this for walk-up displays e.g. in museums and social spaces. They offer great scope for control over input devices, digital objects and information, and make technology available to users who may have had difficulty understanding or using the more indirect methods of control in GUIs. Control can be provided or denied by many features of the design. Different input devices provide different levels of control, and software can be designed to increase or constrain control.

Rogers et al. (2009) illustrated the many ways in which users could have control, using the idea of multiple entry points that give all users opportunities for control by different means. For example, users could move items on a surface, place cards and use tangibles to design a garden. Use depends to a large extent on the loci of control: the availability and placement of points of access to the digital content and how easy it is for the group members to move between them.

However, allowing everyone to act at once can pose problems of coordination: interfaces with high levels of control can cause frustration, anger and disengagement. There can also be uncertainty about what the rules of ownership and transfer should be for virtual objects, compared to well-learned rules about sharing toys (Neary et al., 2009; Mansor et al., 2009). Such freedom from custom and constraint may result in uncertainty and parallel working rather than enhanced collaboration. In particular, while multi-user touch interfaces can make it more inviting and easier for group members to interact whenever they want, there can be too many touches, making it less obvious as to how a co-located group should co-ordinate their interactions, and more generally collaborate. This is in contrast to interfaces with more bounded input devices such as traditional PCs, where taking control of the single mouse or keyboard clearly signifies the baton has been handed over for someone else's turn. A mouse can act in some respect like the 'talking stick' that some teachers use as a tangible device to support turn-taking in conversation: the mouse has the strong constraint that only the holder has physical control of the interface. Waving one's arms or diving in to a 'free-for-all' surface is quite different from having to wait for a turn to control a single mouse. It may be that, paradoxically, interfaces with such distributed and multiple sources of control are not as effective at supporting collaboration *per se*.

One way of dealing with difficulties of all acting at once is to control closely how people have to act together to control a task. The StoryTable (Cappelletti et al., 2004) enforces cooperative storytelling in child pairs by requiring users to touch in order to perform certain actions, such as selecting backgrounds and listening to audio. This seemed to increase levels of engagement but it was still possible for one user to dominate. The need to coordinate two users seemed to contribute to lower cohesion in the story produced, compared to individual stories in a low-tech condition. In contrast, the tabletop story-telling application Telltable (Cao et al., 2010) neither enforced nor

encouraged collaboration, and elicited a variety of different behaviours in children to permit turn-taking, equality of interaction and role-sharing.

2.3 Availability of background information

By availability, we mean what information is on hand in the background to influence users' awareness and control. Availability of background information differs from awareness: awareness involves the ongoing, moment-to-moment, generally implicit cues we use in all interaction, whereas availability concerns background information relevant to the task that is accessible for all explicitly and over time. These displays are often designed in explicitly, for example summaries or histories of behaviour that can be made manifest with the aim of influencing behaviour. The boundary between awareness and availability can shift and is fuzzy. A good illustration of this fuzziness is provided by the 'conversation clock' (Karahalios and Bergstrom, 2006). This display shows a circle made up of coloured bars representing on-line audio input for each user, in real time, representing immediate awareness. After one minute, each circular display retreats towards the centre of the table, so that each successive minute is represented in concentric circles, giving a cumulative history of contributions by each user.

Availability of background information is particularly relevant for multi-user interfaces because of the opportunities these technologies give for harvesting and displaying information about the state of play over time and the history of the interactions. For example, surfaces can represent anonymous or tagged information about previous actions, previous states of the problem and summarised contributions of current or previous users. A benefit of groups working around shared surfaces is that relevant information about a task, both in one's own space and in shared space, can be made available to all. Furthermore, this digital history may mitigate some of the conflicts involved in shared work when one user undoes or deletes another's work. For example, Rick et al. (2009) designed Digitile, a tile-pattern application for the Diamond Touch, to provide a running history of previous states. This enabled users to return easily to a previous state of the design. Background information is also useful when there is limited control e.g. users standing at the edge of a large shared display may not be able to act but can follow what is happening, e.g. in the spectators of

TellTable (Cao et al, 2010) and in groupwork in classrooms with large interactive surfaces. The benefits of this mutually shared knowledge can be contrasted with access to information using similar technologies in small displays: iPods are difficult to use as shared platforms, and tend to be handed over rather than shared, while iPads provide a compromise, being quite comfortable for two people to share information, and easy to hand over, like an iPhone, but less comfortable for simultaneous use in larger groups. Small tangibles such as the Audio Ladybugs used in StoryTable for storing recorded segments of stories, allow access to completed segments of work, although children treated the Ladybugs as personal possessions so tended not to share them readily.

Combinations of interfaces can be configured so that users have shared information on a large surface but also receive different pieces of information via personal technology (e.g. iPods), meaning that users can decide how and when to share that information with their peers. Deciding when to tell what to whom can be crucial to success or failure in a collaborative task, and people vary hugely in their ability and willingness to share information judiciously. For example, Fleck et al. (submitted) found wide individual differences in how well groups shared information to solve a collaborative mystery, when they each had different information presented on personal devices to share on a multi-user surface. Through their own communication choices and through features of the technology, different people will have access to different information.

Designers have provided explicit displays of information about the accumulated history of interactions, which would normally be subject to the vagaries of each person's unreliable or biased memory. For example, Bachour et al's (2008) Reflect table provided an array of LEDs embedded in a table that showed publicly how much each person was speaking relative to the others, aggregated over time. This shows who has talked most, rather than just users' concurrent awareness of who has the floor, with an aim of reducing verbal domination. Digital objects can be tagged to show who last touched them, or who placed them, providing tracking of activity. For example, Gaver's History Tablecloth (2007) showed the history of objects placed on it, by displaying an increasing glow the longer the object remains in place. Meeting Mediator used various sensing devices to detect social interactions which can be

analysed and aggregated to provide abstract visualisations presented via personal mobile ambient devices (Kim et al., 2008). This can have the effect of reducing domination and making distributed meetings appear more like co-located ones.

However, high availability of background information does not necessarily lead to desired results. DiMicco et al. (2007) performed extensive studies with Second Messenger, a system that represented speaker input on a surface. There were many positive effects, but some negative: for example, over time there was decreased decision-making effectiveness in groups that had previously done this well. The Reflect table has the implicit goal of inviting the participants to “balance their collaboration”. This did not always happen: there were still some who dominated the interaction and others who hardly spoke (Bachour et al., 2008). One possibility is that providing very explicit and mutually-shared background information on contribution heightens self-consciousness, perhaps resulting in anxiety for some and competitive instincts for others, and a concern to present oneself in a particular light.

We can see from this selection of examples that awareness, control and availability of information are features of technology that influence collaborative processes.

However, in order to understand how they might best be used, we need to shift focus to research in social and developmental psychology, where researchers seek to identify the psychological mechanisms through which smooth social interaction develops from infancy onwards. In particular, these mechanisms of collaboration involve mutual knowledge and understanding, and in everyday interaction the smoothness of their operation relies on forms of constraint, which we describe below.

2.4 Mutuality of awareness, control and availability

The three mechanisms of awareness, control and availability underpin interactions at the tabletop interface, but each also involves mutuality. Across a range of disciplines, researchers have stressed the role of commonality, jointness or mutuality in awareness (e.g. Carroll et al, 2006). Developmental psychology in particular highlights the importance of mutuality: for example, Moll and Tomasello (2007) showed that infants showed understanding of an adult’s perspective when jointly engaged in activity with the adult before they could do so when merely observing. Each of the mechanisms we

identify involves mutuality. For awareness, it is not just that I am aware of you and you of me, but we are mutually aware of each other's awareness: our awareness is shared (see Figure 1). Because mutual awareness underpins so much of our behaviour, we tend not to recognise this, but researchers in infancy and primatologists understand well the complexity involved (e.g. Gomez, 1994). Furthermore, we need to recognize that awareness is shared not only through visual channels, but also, for example, through vocal and postural means (Akhtar and Gernsbacher, 2008).




Figure 1: Mutual awareness (from Kozima, 2010)

Mutuality is also a factor in control, which is not just a matter of seeing which individual does what. Psychological research into joint action has recently demonstrated how partners are strongly engaged not just with their own actions but also with their partner's actions: we represent what we are doing, but we also represent the other's actions. Just as our own actions can be subject to interference effects within-individual (e.g. being asked to point in the opposite direction to an arrow), so we ourselves experience interference when observing someone else struggling with the same task (Sebanz and Knoblich, 2003). The role of mutuality in availability has been well-recognised already when considering interfaces, in concepts such as grounding (Clark and Brennan, 1991), through which conversation partners try to establish that what has been said is understood, and in the idea of intersubjectivity (Suthers, 2006), in which participants jointly and dynamically construct and maintain shared interpretations and assumptions about a task. This mutuality means that we have to consider not just how one user is aware of the actions and intentions of other users, but how the group can have shared awareness of others, and the ways in which each user's behaviour is planned with that mutual awareness in mind.

3. The value of constraints for awareness, control and availability

Our selection of literature in developmental psychology, computer-supported collaboration and human-computer interaction identified three mechanisms in behaviour through which multi-user interfaces can support collaboration: mutual **awareness**, mutual **control** and mutual **availability**. It is the shared nature of these features of behaviour that provides opportunities for substantially improved collaboration: ‘coordinated, synchronous activity that is the result of a continued attempt to construct and maintain a shared conception of a problem’ (Teasley and Roschelle, 1993, p.235). As we have seen, though, good collaboration does not always emerge naturally from these aspects of behaviour with multi-user touch surfaces. There can be uncertainty about who should do what, how to progress and how to use available information. There is therefore considerable scope for designers to shape the interface to support collaboration.

Several frameworks have been proposed to guide design for collaboration. For example, Benford et al. (2000) describe a continuum of support. The strongest support is enforcing collaboration, a good example being the SIDES multi-touch table application (Piper et al., 2006). Teenagers with various behavioural difficulties played a board game on a multi-touch table that recognized users and permitted movement of digital objects only in strict rotation of turns. This strongly enforced turn-taking, in that only a specified person can move objects at any one time or in one area. The teenagers were observed to respond more cooperatively when the software forced turn-taking than when a teaching assistant tried to persuade them to take turns. At the other end of the spectrum are approaches that give away control, where ‘users can act independently, are mutually aware and are free to coordinate their actions if they wish’ (Benford et al., p. 560), similar to many of the tabletop applications we have mentioned so far. Benford et al. endorse a middle position of encouraging collaboration, i.e. providing an added benefit, or incentive, for users if they work together. Rogers et al. (2009) describe a ‘shared information spaces’ framework that focuses on the extent to which design constrains or invites participation, through the idea of entry point – something that invites entry to an environment. Both of these frameworks imply a dimension of constraint – the extent to which design denies or provides possibilities or incentives for action.

However, the concept of constraint alone does not provide sufficient guidance in designing for collaboration. Constraint, like awareness and control, is not good or bad in itself. Sometimes removing constraints, such as on turn-taking, can cause discomfort for users. On the other hand, multi-user interfaces such as tabletops relax constraints on awareness so as to enhance mutual understanding in ways that some other technologies do not, without needing additional incentives being built in. For example, the SIDES game benefits from the unconstrained mutual awareness afforded by a multi-touch table, but control needs to be constrained, at least initially, to prevent chaos. In contrast, designing a garden planning task for adult users of the garden works well if participants have many unconstrained entry points, with physical tagged objects around the room and digital objects on the table, such that participants who do not speak much still contribute by placing the freely available objects into the design, even though mutual awareness might therefore be lower (Rogers et al, 2009). The degree of constraint or freedom provided for each of the different mechanisms, of awareness, control and availability, helps designers consider what needs to be constrained or permitted, and for what reasons. It is here that an understanding of how constraints work in naturally-occurring human-human interactions, and how these support smooth collaboration, can inform design for multi-user interfaces.

3.1 Psychological research into constraints in everyday interaction: awareness, control and availability

Consider the way constraints on awareness work in spontaneous everyday interactions. For example, a simple mother-infant interaction occurs when an interesting event triggers pointing by the mother, who naturally orients herself and the baby in her arms to the object of interest, thus providing a constraint on the infant's awareness that makes joint attention more likely. This can support the infant in learning that pointing signifies something of potential interest for joint attention (Trevarthen, 1978). The environment itself provides useful constraints: if there is an unusual event, e.g. a flashing light or movement, people's attention is drawn to it, giving a mutual awareness of the event – a feature psychologists exploit to great effect in studying joint attention in infancy (e.g. Butterworth, 2008). Ways of making meaning visible through constraining attention and interpretation come naturally even to very young children: for example, children as young as 4 use gestures to counter

ambiguity in speech (Kidd and Holler, 2009) and pre-linguistic referential pointing is a powerful communicative device to direct others' attention, used from 9 months of age (Carpenter et al., 1998).

The use of constraints on control is widespread and generally apparent, particularly in educational interactions. In a typical context of scaffolded or apprenticed learning, a tutor will constrain the possible actions or choices so that the learner is supported in picking the correct one: constraint is gradually reduced by the tutor as the learner becomes more able to self-regulate. Driving instruction can move from dual control to single as the learner becomes more competent. Children develop self-control through the adult's gradual relinquishing of control (e.g. Rogoff and Lave, 1984), and become able to provide their own constraints: for example, a common test of self-control is to require a child to sit by a desirable food without eating it: some young children will sit on their hands, or shut their eyes, to provide physical and sensory constraints on action (Mischel and Mischel, 1983). Considering how much to constrain or relax control on tabletops, e.g. in regulating turn-taking or informing about different users' verbal contributions, would depend on how much the designer wanted to rely on or encourage self-regulation in the users.

The role of constraint on availability of information is well-illustrated in educational contexts. In a typical school setting, information is presented in selective and structured ways rather than learners being left to their own devices. Schools also have implicit ground rules, sometimes made explicit as 'golden rules' posted on the classroom wall. In more informal learning, research into children's acquisition of mental state language shows that mothers systematically constrain their vocabulary use to refer to the child's mental states, rather than to those of others, until the child has mastered the particular mental state through their own inner experience (Taumoepeau and Ruffman, 2006). In a broader context, Garrod and Pickering (2004) explain how conversation is helped by the way that interlocutors align their situational awareness, thus constraining interpretations that they might make, suggesting that humans are 'designed for dialogue rather than monologue' (ibid, p. 8).

Thus, people naturally behave in their everyday lives in ways that constrain action and information so as to interact together effectively. The rules and constraints of when,

where and how to gesture, whose turn it is, what to mention and what to take as read are in everyday life culturally and contextually bound. In contrast, so-called ‘natural’ user interfaces can often lack the constraints, rules, or properties that characterise ‘flowing’ human interaction (c.f. Csikszentmihalyi, 2000), offering too many possibilities, or possibilities that are infelicitous. For example, a multi-touch table on which all can work simultaneously on a complex problem can produce parallel working, where each person makes rapid changes that cannot be apprehended by others who are working simultaneously and not monitoring others’ actions, which happen rapidly and simultaneously. Instead of helping collaborative work, shared surfaces may inadvertently encourage individual working. Providing background information, such as on verbal contributions, gives freedom to use the information but maybe without any shared background assumptions about whether and how users should be equitable, and without any shared understanding of how behaviour might be managed or controlled by the group.

We propose that in order to overcome such problems and to be able to design settings that can enhance collaboration there needs to be an understanding of how to provide different levels and kinds of constraints. This involves working out how to constrain an essentially unconstrained interface and how to factor into design the different constraints imposed by different interfaces. Just as humans in naturalistic social settings implicitly constrain awareness, control and availability in a whole range of collaborative interactions, we argue that designers need to consider how to constrain these features, or how to enable users to constrain them, when designing multi-user interfaces to support collaboration. Everyday social interaction, particularly in the context of teaching, provides a mine of information about how this might be done.

3.2 Sources of constraint

In considering the ways that designers have used features to support collaboration with multi-user interfaces, we can see a variety of degrees of constraint. For example, a strong constraint on turn-taking control would be software that prevented input from specific users. A weaker constraint would be implying to users what they can and cannot do, based on social norms. Having discussed how constraints affect everyday behaviour, we can now turn to how they might be used effectively for each of our

three mechanisms of behaviour. To help in thinking about how to use constraints, we consider three different sources: physical, digital and social. We give examples below of each of these different sources of constraint and then describe how they can be used in relation to awareness, control and availability of information.

(i) Physical constraints

Constraints arise through what is physically possible for us (e.g. a user's reach or hand span, our attentional capacities). One of the benefits of tables for groupwork is their size: they are big enough for all to see easily what others are doing (awareness), but not so big as to prevent reaching all or most digital objects (control). The orientation of a surface (horizontal or vertical) affects collaboration through influencing control (Rogers and Lindley, 2004). Physical features are an important part of background information with or without technology: board game pieces, with their position and colour, for example, act as physical signs of states of play, and tangibles used on surfaces can provide opportunities for representing previous actions or ownership of objects (Brereton and McGarry, 2000). Tabletops draw on physiological mechanisms we already use to make sense of the world. For example, people may avoid clashing arms on tabletops in ways akin to how they avoid bumping into strangers on busy streets (e.g. Nummenmaa, 2010). Technology can also be designed to extend what is physiologically possible (e.g. 'seeing' through touch, Bird et al., 2009). However, there remain constraints on our processing capacity, meaning that we cannot attend simultaneously to all the information that new technologies might provide for us. Thus, designers may limit the amount or orientation of information presented, or may locate buttons so that they are not accidentally activated, to compensate for the cognitive and physical limitations of technology users.

(ii) Digital constraints

Hardware and software can be used to change the constraints on the mechanisms of behaviour we have identified. A clear example of this for control is using software to enable single or multi-touch on surfaces. We can also constrain the physics of awareness through technology e.g. by using mirror projection to present different

views to users in different locations, as in the Lumisight table (Matsuda et al., 2006). Such applications make it possible to manipulate and hence evaluate the role of awareness in collaboration, e.g. by providing different degrees of constraint on what people can see, and assessing how this influences ways the group works together. One of the main strengths of new multi-user interfaces is the digital possibilities they provide for control. Multiple users can control actions simultaneously through touch or stylus, through tangibles and using personal technologies such as iPods. While having multiple entry points to the technology can work well in allowing all to participate, there is a danger in giving too much control, as we illustrated above. Information availability is particularly open to variation through digital means. Representing background information in digital form and using log data allows creative potential for applications to adapt to individual users over time.

(iii) Social constraints

Interfaces are not just physical and digital spaces: people come to them with sets of social norms, mutual understandings, levels of familiarity, implicit awareness of others and motivations. Hence we also need to consider multi-user interfaces as social space. Technology-mediated collaboration is strongly driven by implicit norms about one person not hogging the space, cues as to ownership, such as colour of icons, and conventions about use of personal space (Marshall et al., 2011). Designers can make background information such as the history of interactions or equity of contribution available, but there is not necessarily a mutual and explicit set of assumptions about how to use this information to guide actions. Little attention has been paid to implicit norms about how such information is used – should we be contributing ‘fairly’ or ‘equally’, are users a coherent group or a collocation of individuals, do the ‘digital natives’ dominate? For such interventions to have the desired effects, all participants would need to acknowledge implicitly the function of the explicit representations and to act accordingly and in concert. While some dominant people may notice they are speaking too much and reduce their contributions to the ongoing conversation, others may choose to ignore it. Thus, a participant using the Reflect table (Bachour et al., 2008), who spoke the least in her group, said that she rarely looked at the display and moreover, she did not feel it was important for all members of the group to participate equally. Techniques using implicit norms rely on the group members recognising and

mutually accepting the etiquette implied in design, and one-off groups of unfamiliar people may be less likely to establish shared understanding and ‘social contagion’ than familiar groups (e.g. see Mujde and Teckan, 2009)

Use of one-off groups has restricted consideration of the social psychological aspects of multi-user interfaces. The degree of familiarity of users and their assumptions are important in understanding how people apprehend and interpret the intentions and actions of others. Established peer groups, for example, can be familiar and comfortable as work partners, quick to apprehend others’ plans (awareness), relatively comfortable in exerting control, e.g. elbowing others out of the way, and open to transparent discussion of the social norms of groupwork in the specific setting, e.g. by using peer pressure and reference to class rules to alter plans or complaining about domination (Harris et al., 2009). Studying ‘wild’ settings, such as groups of playmates or relatives in public or private settings, who know each other well enough to use elbows, wrest control and appeal to norms of behaviour, or unacquainted visitors to exhibitions, can help clarify intergroup differences in availability of background information.

3.3 How to use constraints?

How much constraint is the right amount? To answer this, designers need to consider how physical, digital and social features constrain or expand aspects of awareness, control and availability. Low levels of constraint may be most appropriate for very open-ended tasks. High levels of constraint can be unhelpful, if people become frustrated or are too restricted ever to learn or understand independently. There has been a tendency to provide either overt technological constraints or to provide information based on implicit norms, e.g. about equity, without considering fully how that information might be viewed and used by participants. In particular, users may not always take kindly to cooperating with the researchers’ implicit assumptions about being equitable when told they are contributing too much or not enough. We propose designers should also consider providing subtle environmental constraints to encourage collaboration. For example, marked changes in collaboration and equity can be produced simply by changing from a horizontal to a vertical surface. The constraints need to provide just enough support, and not too much, supporting the

smooth use of regulation, not purely at an individual level, but also at the level of the group, and considering the norms of behaviour that apply in the specific setting. For example, if a goal is to help users to learn to take turns, then the framework could be used to design in only weak constraints, presenting challenges so that users develop an understanding of how and why to take turns. Similarly, if the goal is to enable older or more socially-skilled users to play a tabletop game, the constraints can be relaxed by building in opportunities for users to regulate their own turn-taking, for example, by allowing single touch by any user. This provides high constraint in control (only one user at a time) but no constraint on which user, so the group can only work together if there is some agreement about how control is shared. In a situation where users are in competition to contribute, it may be apt to impose a strong constraint to enforce turn-taking.

The level of constraint will be influenced by the skills of the user group. For example, if the goal is to support collaborative play in autistic children, who may be limited in their ability to communicate and show little apparent awareness of others' needs, the framework can provide strong physical constraints on control, e.g. a 'train track' painted on a playground surface to guide movement, and a curved slide that 'lands' the child at the start of a new activity, providing smooth movement into the next activity, but also some opportunity for obstruction of others, supporting the need for children to negotiate some simple rules of priority (Yuill et al., 2007). Strong constraint is sometimes assumed to be counter-productive, but the role of constraints is well-recognised in evolutionary and developmental theory and in supporting creativity (Stokes, 2005). In sum, there is no 'right' level of constraint, but designers need to consider how much constraint is needed on each mechanism for a specific user group, and what the best source is through which to impose that constraint, or to allow users to change constraints themselves.

4. Case studies

We now describe three case studies from our own work showing how the framework can be used to consider how changing constraints on the mechanisms of behaviour can support collaboration for specific groups of users. In particular, we contrast

outcomes for different versions of software and provide empirical evidence of how the changes in constraint of each feature influence interaction with the technology.

4.1 Case Study 1: A software architecture for manipulating awareness, control and availability (SCoSS)

The Separate Control of Shared Space architecture (SCoSS; Pearce et al., 2005, see Figure 2) is designed to support collaboration between multiple users. It provides ‘separate control over an identical version of the task for each [user] within their own private screen space that is visible to both participants’ (Kerawalla et al., 2008, p. 195). It illustrates variations in awareness (user-identified ongoing task states), in control (yoked or shared) and availability of information (about previous joint decisions).




Figure 2. Typical layout for non-SCoSS (left) and SCoSS (right)

Figure 2 shows a non-SCoSS (left) and SCoSS (right) version of a simple task on a DiamondTouch table where users can be identified. The specific task set for pairs working together in an empirical study using adapted PCs by Kerawalla et al. (2008) was to sort 12 words into a 2x2 matrix identified by a multi-coloured grid. Each word varied on two dimensions: a surface quality such as initial letter, and a semantic category, e.g. animal, colour. The dimensions were not given in advance and only one word appeared at a time in the red-bordered word box for categorising (by moving into a coloured cell of the matrix). The task is challenging: users have to develop and test out hypotheses about what the categories might be, and there may be ‘red herrings’ so they may have to backtrack and revise their solution through extensive discussion. In the non-SCoSS version, one copy of the word appears in the word box, and either user can move it into a cell of the single matrix. The non-SCoSS version,

used in a multi-touch setting, has a single ‘we agree’ button, activated by either user, designed for the pair to use when they have reached agreement on a particular word. Thus in non-SCoSS, there is no technological constraint on control: one user could undo the partner’s work, move a word into the matrix and click ‘we agree’ without consulting the partner. Sometimes pairs were uncomfortable with the lack of constraint on control here, and would create their own by allocating roles, e.g. ‘you do the ‘ch’ words’ (Kerawalla et al., 2008).

In contrast, the SCoSS condition has two copies of the matrix, representing what each user thinks. *Control* is constrained as the technology can identify the users e.g. by mouse or by conductive mats: each user can only control their own matrix. They have separate ‘we agree’ buttons, each controllable only by the identified user, and both buttons have to be activated in order to continue with the task. It is therefore much more difficult for one user to dominate, unless by grabbing the mouse in the PC or stepping onto the other user’s conductive mat on the DiamondTouch, both extreme behaviours rarely seen. Users anticipate their partner’s action and will wait to press their ‘agree’ button until the partner is ready. The two matrices represent an opportunity for *awareness* of different points of view: users move words, hands or mice around in their own or the partner’s space in ways that show their partner what they think about a user-specific task representation. Participants have a ready source of awareness of the intentions and deliberations of others, because they observe and use information about the ‘hovering’ of their partner’s mouse, or hand, over their own copy of an item, ‘picking it up’ and moving it to a place on the grid, but not releasing it. This hovering is a clear indication of others’ intentions. Sometimes items can be moved as explicit deictic devices to resource discussion (e.g. ‘Shall we put it *there*?’ with the item being held over a location), sometimes in a way that communicates epistemic states such as uncertainty (e.g. ‘wiggling’ the item over several locations when unsure where it should go), and sometimes items can be ‘waved’ at an intermediate point, apparently as a subtle signal of willingness to negotiate, rather than more direct methods such as making a verbal suggestion about an item’s position or actually placing an item in disagreement with a partner.

Thus, a high level of awareness in SCoSS is supported by the provision of a visual indicator of agreement history that yokes the two users’ beliefs: when both users place

a word in the corresponding cells of their matrix, that word in each matrix turns green. The two interfaces thus also differ in availability of *background information*: agreement (and hence disagreement) is represented in SCoSS, providing a resource for discussion, in contrast to non-SCoSS, where action by one user on the single space deletes or overrides what has already been done.

Use of this interface illustrates and makes the most of the high levels of mutual awareness afforded by multi-user interfaces, in obvious ways (e.g. using green to represent agreement) and also more subtly in behaviour, gesture and apprehension of others' intentions. It also supports technological constraints on control, e.g. reducing the possibility of one person dominating by constraining single-user actions, giving more opportunities than non-SCoSS to discuss possibilities before irrevocable action is taken. Availability of background information is achieved by representing areas of agreement and disagreement with the partner. Newer multi-touch technologies provide additional opportunities for awareness, control and representation of background information for multiple users. Designers can focus on ways in which constraints on these mechanisms of action can be used to support collaboration.

The power of such effects of constraints on control and awareness were shown both in empirical studies of word categorisation, where children using SCoSS showed more complex discussion than in non-SCoSS (Yuill et al., 2009), but also in studies with pairs of severely autistic children (Holt and Yuill, submitted). As expected, the autistic children showed no apparent awareness of others, such as anticipating action or hovering, in the (typical) non-SCoSS architecture. In contrast, in the SCoSS condition, children showed 'active awareness' of others, behaviours not previously apparent to their teachers, such as anticipating planned actions of others by waiting (hovering) for the partner to place an item in the correct box before clicking their own 'agree' button. Differences in social constraints (whether the SCoSS partner was a peer or an adult) also led to changes in collaborative behaviour. The case study illustrates how constraints can be implemented in the software to increase collaborative discussion, reduce domination and highlight awareness of different perspectives, as a resource for joint problem-solving.

4.2 Case study 2: The effects of multi-touch versus single touch on awareness and control (OurSpace)

Our second case study demonstrates how varying constraints on awareness and control has interacting effects on the nature of collaboration. OurSpace is software developed for multi-touch surfaces that represents a space (e.g. a set of offices or a classroom) with digital objects (people, furniture). Groups, usually of three people, arrange the space in a mutually agreed way given various constraints (Figure 3). In the version for classrooms, background information was made available on the icons representing each student, and the information could be used to constrain the solution in different ways, which were up to the user to specify. Thus, colours were used to mark different friendship groups, the presence of a speech bubble represented a talkative student and some students were depicted wearing glasses. In one study, triads of children organised their classrooms in a multi-touch (MT) and a single touch (ST) condition (Harris et al., 2009). In the MT condition, children had lower constraints on *control* than in ST, as all could move icons at the same time. This was expected to decrease *awareness* in groups of 3 users, since children could not apprehend their own and two other users' movements at the same time, in comparison to the ST condition, where all can be focused on a single movement at a time. Users' awareness depends on how many users there are and how many icons are to be moved: in this case the numbers were large enough that awareness was compromised in a MT condition.




Figure 3. OurSpace: Tabletop classroom design application

Qualitative analysis of conversation suggested that, as expected, MT brings lower levels of *awareness* of others' actions than ST: there were frequent, and urgent, requests for others to stop, or wait, when the simultaneous action was too fast to be apprehended, and in fact one group had to stop working temporarily because a child was upset that she 'lost the plot' given the speed of simultaneous action. However, this lack of awareness is replaced by other behaviours such as arms clashing across the table, and users making running commentaries on their actions. Not surprisingly, children spent significantly more time discussing turn-taking in the ST condition than the MT, in which they exhibited more task-focused discussion. Turn-taking talk seemed to replace task-focused talk, since the two were negatively correlated. There were no overall differences in physical or verbal equity between the two touch conditions.

The decision about how to constrain *control* depends on the goals of the activity: if the objective was for children to learn to negotiate about turns, then ST would be appropriate, whereas if the goal was for the most considered solution to the task, then multi-touch would be preferred. The goal of a collaborative task will depend on the shared background knowledge and abilities of the users: younger children can show lower levels of equity because of their inexperience in regulating turn-taking, and this will be exacerbated in a ST condition. However, the school milieu supports fairness and allowing others to speak, meaning that children can speak out about inequity in ways that might remain unspoken for one-off groups with no clear norms. As one child completing OurSpace complained to his two peers, 'Why aren't you letting me arrange the tables? You two are like the bosses and I'm like nothing, you're not letting me do anything'.

OurSpace provides an example of how the different mechanisms interact: placing a constraint on control (using single-touch, preventing simultaneous action) supports better awareness of what each user is doing, compared to the simultaneous actions in multi-touch, but at the same time, it highlights awareness of inequity of turns, so

challenges the users to regulate turn-taking fairly, in the absence of a software constraint that enforces rotation of turns. Designers therefore need to consider how constraints in awareness will influence control, and how this in turn is influenced by availability of background information, such as about equity and turn-taking

4.3 Case Study 3: Tangible interfaces (Augmented Knights' Castle)

Our third example shows how the framework can be applied beyond shared surfaces, to understand the design of tangibles and digitally-augmented objects in supporting collaborative interaction. Tangibles offer immediate opportunities for high levels of awareness, control and the application of background knowledge based on the norms that apply to physical objects.

The Augmented Knights' Castle (AKC: Figure 4; Hinske et al., 2008) is a Playmobil® medieval castle playset containing figures fitted with RFID tags, such that the figures 'speak' specified sounds when they are in particular locations of the set – for example, the Red Knight says 'Let us attack!' when in the castle courtyard, or the dragon roars when encountering the princess.



Figure 4: The Augmented Knights' Castle in play

Children's free play with the augmented playset was contrasted with that in a switched-off version (the KC version). There are striking differences between the two

conditions, for both typically-developing children and for children with autism (Farr, and Hinske, 2010; Yuill et al., submitted), in that there is significantly less solitary and more social play in the augmented version. In terms of *awareness*, the playset in both versions allows all-round views, allowing each child to see the focus of attention among their playmates through eye gaze and body orientation. More crucially for the quality of the play, the speaking figures in the augmented version apparently provide a boost to mutual awareness: simply because the figures can be expected to produce interesting and relevant sound effects, a child's bid to gain the attention of peers by showing a figure and looking at or vocalising to others is significantly more likely to be successful with the technology than without. This apparently minor boost seems to have a snowball effect on the quality of play and in subsequent story-telling, producing more social play, more balanced narratives and more creative stories than in the KC condition. This demonstrates the potential for powerful cumulative effects of relatively fleeting aspects of mutual awareness. Similar effects might be found with other technologies if a user's interaction with a digital object is highlighted through sound or visual effects (e.g. if someone shifts a key piece of a design).

In terms of *control*, each figure has a potential range of sounds that are activated only when a user picks the figure up and places it in a particular location. Because there can be many figures and locations, control is fairly unconstrained and distributed among users (multiple entry points). Control for KC and AKC is somewhat unconstrained, and distributed among users, with many more entry points than users, and thus engenders a situation where multiple narratives and cross-cutting actions are typical.

For a free play setting, this unconstrained control seems to work very well: both AKC and KC groups showed high levels of equity in their narrative conversation. The AKC effects, however, are not entirely unconstrained: because of the RFID technology, characters can utter sounds appropriate to a specific location. Users understand the contingency of figure-location combinations producing sound, so there is more or less systematic exploration of these combinations. In theory, in the KC 'anything goes': users have no constraints, other than physiological, on the sounds they themselves can produce in each location. However, the lack of specific pre-programmed pairings was associated with reduced creativity: stories told after the KC sessions were less creative

than those in the AKC. The moderate constraints of providing sounds apparently helped to suggest possible actions in an area, to teach vocabulary items and speech patterns and to prompt ideas for storytelling. The moderate constraint in the AKC condition was accompanied by significantly greater shifting between types of narrative (speaking in character, telling the story, making suggestions for the story) than in the KC condition. Notably, in the KC condition, there was a significantly higher level of debate about the story, meaning that often no single narrative ever got off the ground: the total lack of constraint in what characters might say seems to have produced chaos.

Constraints on *availability* are not a strong feature for the AKC, as in many other tangibles: there is no visible history for the user of who has done what, e.g. who has configured or triggered a particular sound effect, so the history is opaque (except for the researcher, who has a useful web-based log of activity, that could be used in displays for users). The character figures themselves are saturated with meaning, e.g. the actions, personality and preferences likely in a knight, a dragon or a princess figure, and these provide some constraints that help users in KC and AKC to construct scenarios. However, the AKC links specific sounds to characters, which can prompt elaboration on a given theme. For example, children heard particular archaic speech patterns, e.g. ‘My Lord!’ and extended these creatively to other characters (‘My Queen!’, ‘My Knight!’), and developed stories using actions that the characters’ speech suggested.

An important feature of the set is that it is played with similarly to a traditional toy, so the norms that guide and constrain behaviour transfer readily from typical non-augmented play, e.g. that there is no set structure, that the figures do not belong to a particular person and therefore should be shared unless a claim is made. The underlying assumptions of play are brought into relief through comparing play by typically-developing children and children with autism: for the latter, assumptions about the affordances of the objects are often not shared: for example children with autism may use the play figures atypically, e.g. as hammers or wedges rather than as characters, affecting the possibilities of shared understanding in play, and this lack of constraint is a challenge for design with this user group.

This case study shows how the mechanisms can be extended to digitally-augmented tangible objects. As these are usually based on cultural artefacts that already exist, there is robust background information available about how the objects are to be used, functionally and socially. In the AKC, this is built on by having speech and sound that elaborate on the historical context, prompting imitation, learning and elaboration. Tangibles also tend to provide high levels of awareness, since movements of objects in space can be readily seen, and digital augmentation enables further signals of action, e.g. audio effects, visible trails of movement, haptic cues. Tangibles also offer relatively unconstrained control, or access points. In the AKC, having the constraint of sounds only playing for specific figure-location pairings supports structured exploration. These three mechanisms work for tangibles in a similar way to Hornecker (2005) idea of ‘embodied facilitation’, with its three concepts of embodied constraints on control, multiple access points to awareness, and tailored representation based on users’ prior experience.

To summarise the effects of varying constraint in each mechanism, Table 1 shows the different behaviours observed in the two versions of each case study, illustrating how these can be exhibited or inhibited depending on constraints on the different mechanisms.

Condition	Application	Awareness	Control	Availability
Case Study 1 (dyads)	SCoSS	-- hover hand -- wiggle object in own space -- 2 displays give shared, user-specific focus for gesture and talk	--anticipate other’s action -- wait to press ‘agree’ --agreed choices are ‘frozen’	Visual sign of past agreement (greenness) used as resource for discussion
	NonSCoSS	simultaneous action -- moves can go unnoticed by partner	--1 user can dominate -- undo other’s work --autonomously click ‘we agree’	No traces of past agreement

Case Study 2 (triads)	OurSpace Singletouch	sequential action: single shared focus no clashes	sequential action: more negotiation of turns	No traces of who did what, so undo others' work, classroom assumption of fairness/ equity
	OurSpace Multitouch	simultaneous action – distributed attention -- clash of arms -- running comment	simultaneous action: free for all. More task-focused talk	
Case Study 3 (groups)	Augmented Knights' Castle	--gross movements of objects in space --sound effects track objects -- 'boosts' attention	multiple entry points, sounds constrained by figure-location: more exploration	Free play so low constraint, but sounds prompt elaboration
	Knights' Castle, not augmented	shared visual awareness but no augmentation by sound	Multiple entry points, but 'anything goes': no sounds to explore	Free play so low constraint, no prompts through sound

Table 1. Group behaviours relating to each mechanism in the two versions of each case study. The more constrained in each design pair is shown shaded.

5. Design Implications

We have shown how the mechanisms framework can be used to explain how constraints can elicit behaviours supporting collaboration. We now outline how the mechanisms framework can be used to inform the design and evaluation of technologies for different researcher or developer goals for collaborative working.

- Constraints: technological, physical or social?

We have highlighted different sources of constraint to help designers think about the many different ways in which constraints can be modified. For example, awareness can involve any sense modality, alone or in combination, and haptic and audio channels could be used more extensively, particularly in the case of tangibles. The focus is often on technological sources of constraint but it is important to be aware of

the role of constraints that operate without being explicitly designed in, such as implicit norms about group behaviour. Underlying assumptions of how to behave have powerful effects which we sometimes see better when we move beyond the lab or design studio, and beyond the easily available user. Considering how variations in constraints of different types influence the mechanisms of behaviour can support the systematic empirical investigation of design on collaborative interaction and provides a strong bridge between HCI and psychology.

- Awareness, control and availability: questions to ask

When considering what users will be aware of from moment to moment, we need to consider what information is at hand and how it is constrained. This requires consideration of orientation of display, presence of personal ‘territories’ for users, size of surface in relation to users, amount of information a user is able to apprehend, possibility of users interfering with each other’s awareness and the role of personal technologies that might require sharing, or might inhibit mutual awareness. Designers need to consider how control by different users can be coordinated or constrained, how this might be orchestrated by the technology, by the users, or by the technology having a role in moving control from the technology to the group, and the extent to which users might undo each other’s work or might need to negotiate agreement. For availability, we need to consider what information is made manifest, and whether there are implicit norms underlying the use of such information. The norms may be implicit but not shared. Multi-user technology has the capacity to present otherwise hidden information to the users’ mutual awareness, and hence users’ potential to control how this information is used to support better collaboration.

- Designing for different user groups

We have placed a particular focus in our review on user groups beyond the typical. Designing for users who are less able to self-regulate naturally leads to the idea that the technology has to constrain in order to ensure, e.g. that users behave equitably, take turns and focus on the information that is most important. However, there may be cases, particularly in pedagogical design, when designers may want to relax constraints, perhaps gradually, so that users develop the ability to self-regulate, either through their own behaviour or through altering the way that the technology operates throughout an interaction session. This gives technology an important role in

enabling groups to gain greater control over their collaborative work. Designers need to be aware of the cultural practices that users associate with the multi-user interface, and these may be particularly rich for tangibles, e.g. in assumptions about what actions the objects afford and how they are to be shared. For older children, the issue of social control over the interface appears particularly important, so for example control through software might be more acceptable than control by an authority figure. We might expect adult users to abide more closely by social norms, but these will vary by social context, for example a game presented in a work or play setting. Users in mixed groups, such as families or serendipitous groups of people in public spaces, raise different considerations, and could perhaps benefit from tailored user-specific constraints. The constraints and mechanisms framework also invites reflection on how design might use different ways of supporting awareness, control and information availability in groups where there might be restrictions, e.g. in children with autism who will have difficulty with awareness of others, or those with sensory impairments. Here, considering how awareness might be enhanced, e.g. by sound, could both enhance the ability of such users to collaborate with others, and also supports our understanding of the nature of such impairments and how they compromise non-augmented interaction.

- Beyond the tabletop

We have focused mainly on interactive surfaces, but the framework can be applied to other multi-user technologies such as networked tablets, tangibles and combinations of personal and shared technologies. All these provide wider scope for manipulating awareness, control and availability, as shown in our case study of a tangible, since these can be shared or asymmetrical, as in the case of a group using a shared surface along with individual devices, or a group each with individual tablets that are networked with each other and a shared display. Where there are individually-controllable devices, we need to consider what triggers users to share information rather than keeping it to themselves.

- Multi-touch or single touch?

Varying how many people can act on a multi-user interface affects awareness and control. Single-touch surfaces constrain awareness by heightening shared attentional focus, useful for tasks where users' actions are highly interdependent. They also

constrain control, which might help focus attention but prompts a need to negotiate turn-taking. If this negotiation is difficult, the designer could choose to constrain control by using identification of users to enforce turn-taking. In multi-touch, all can act simultaneously, so several events can occur simultaneously in users' peripheral field of attention, reducing awareness of what others are doing. Control is less constrained than in single-touch and users might find this liberating, if actions are not very interdependent, or frustrating, if one user's action is contingent on another's. Multi-touch can work well when seeking to support serendipitous effects and creative solutions.

- The value of availability of background information and making norms explicit
Ambient displays have great potential for conveying background information in real time, e.g. on contribution, previous actions, use of space. However, making such information transparent is not enough: designers have to consider how this information is to be used. Such displays rely on the group accepting and abiding by the social norm that underlies the design. In this sense it is a low constraint, since people can simply choose to ignore the display. When there is little mutuality about the background assumptions and no easy way to step outside the task to discuss them, this sort of approach seems doomed to failure. Previous research largely ignores the norms that prevail in multi-user settings, in part because studies are often carried out in relatively artificial lab settings. The potentially important but hidden effects of availability suggest the need to consider users' underlying assumptions. Providing such background information could be used as an opportunity for focusing users' awareness on group processes, allowing them to reflect as a group on how they might be guided by this information.

- Low versus high constraints?

There are some situations where low constraint can be better for collaboration, for example, if people are conducting an activity on which they can each work independently, and need little awareness while doing so of other people's activity. Examples include solving puzzles and games. Towards the end of a puzzle, where only one or two pieces need to be fitted or clues to be solved, one person's actions may then constrain another's more closely, so it is easier if they are constrained to slow down, e.g. by switching the software into single touch mode.

In a task with sustained periods of relatively independent or parallel working, it would be frustrating to have only a single entry point such as single-touch, mouse or tangible. If participants do not need to follow everything that their partners are doing, and if tasks are relatively independent, as in division of labour, or in free play, multiple control and multiple entry points can work well.

By explicating the mechanisms of awareness, control and availability, and examining interaction closely to see how design can support these mechanisms in collaborative interaction, we want to encourage designers to consider each mechanism in terms of how the particular technology supports it and to decide whether each needs constraint or relaxation to support the desired form of interaction.

6. Conclusions

The mechanisms and constraints framework is intended to support the design and evaluation of multi-user interfaces to support collaboration across a range of devices, users and settings. We proposed three mechanisms of behaviour that can support or obstruct collaborative working: *awareness* of one's own and others' gestures, actions, movements and mental states, *control* over interactions with the interface and availability of the *background information* and expectations that users bring with them. All of these involve mutuality: not just an individual's awareness, or control, but each person's recognition of the other's beliefs and intentions. Each of the mechanisms can be subject to higher or lower levels of constraint, and these constraints can arise from three different sources: the constraints imposed by physical capabilities, by what is allowed in the technology and by the social rules, often implicit, that underpin users' approaches to the technology. We presented three case studies to illustrate the ways in which constraints on these mechanisms can work for or against better collaboration. In each case, varying the technological constraints in these mechanisms provides hard evidence of how collaboration is affected. There is no single ideal level of constraint: it depends on the aims of the designers, the interactions of the three mechanisms and the composition of the group. We offer the framework to designers and evaluators as a way of generating ideas, codifying observations and reflecting on how to support collaborative work and play.

REFERENCES

- AKHTAR, N., AND GERNSBACHER, M. 2008. On Privileging the Role of Gaze in Infant Social Cognition. *Child Development Perspectives*, 2, 59–65.
- BACHOUR, K., KAPLAN, F., AND DILLENBOURG, P. 2008. Reflect: An Interactive Table for Regulating Face-to-Face Collaborative Learning. In *Times of Convergence. Technologies Across Learning Contexts*. Lecture Notes in Computer Science, 39-48. Springer.
- BENFORD, S., BEDERSON, B. B., AKESSON, K., BAYON, V., DRUIN, A., HANSSON, P., HOURCADE, J. P., INGRAM, R., NEALE, H., O'MALLEY, C., SIMSARIAN, K. T., STANTON, D., SUNDBLAD, Y., AND TAXEN, G. 2000. Designing storytelling technologies to encourage collaboration between young children. *Proc. CHI'2000*. ACM, 556–563.
- BIRD, J., MARSHALL, P., AND ROGERS, Y. 2009. Low-Fi Skin Vision: A Case Study in Rapid Prototyping a Sensory Substitution System. *Proc. BCS-HCI'09*. 55-64.
- BRERETON, M., AND MCGARRY, B. 2000. An observational study of how objects support engineering design thinking and communication: implications for the design of tangible media. *CHI Letters* 22, 17–224.
- BUTTERWORTH, G. 2008. Joint Visual Attention in Infancy. In *Theories of Infant Development*, G. Bremner and A. Slater, Eds. Blackwell Publishing Ltd, Malden, MA, USA., 317-354.
- CAO, X., LINDLEY, S., HELMES, J. AND SELLEN, A. 2010. Telling the whole story: Anticipation, inspiration and reputation in a field deployment of TellTable. Proceedings of CSCW 2010, *Computer Supported Cooperative Work (CSCW'10)*, ACM, 251-260.
- CAPPELLETTI, A., GELMINI, G., PIANESI, F., ROSSI, F. AND ZANCANARO, M. 2004. Enforcing Cooperative Storytelling: First Studies, *Proc. ICALT, Fourth IEEE International Conference on Advanced Learning Technologies (ICALT'04)*, 281-285.
- CARPENTER, C., NAGELL, K., TOMASELLO, M., BUTTERWORTH, G. AND MOORE, C. 1998. Social Cognition, Joint Attention, and Communicative Competence from 9 to 15 Months of Age. *Monographs of the Society for Research in Child Development* 63, 4.
- CARROLL, J., ROSSON, M, CONVERTINO, G. AND GANOE, C. 2006. Activity awareness and teamwork in CSCW. *Interacting with Computers*, 18, 1, 21-46.
- CLARK, H. H., AND BRENNAN, S. E. 1991. Grounding in Communication. In *Perspectives on socially shared cognition*, L. B. Resnick, J. Levine, and S. D. Teasley, Eds. APA Books, 127–149.
- CSIKSZENTMIHALYI, M. 2000. *Beyond Boredom and Anxiety: Experiencing Flow in Work and Play*. Jossey-Bass.

- DI MICCO, J. M, HOLLENBACH, K., PANDOLFO, A. AND BENDER, W. 2007. The impact of increased awareness while face-to-face. *Human-Computer Interaction* 22, 47-96.
- DIETZ, P., AND LEIGH, D. 2001. DiamondTouch: A multi-user touch technology. *Proc. UIST'01*, ACM, 219–226.
- EASTERBROOK, S. 1996. Coordination Breakdowns: How flexible is collaborative work? In P. Thomas. Ed. *CSCW: Requirements and Evaluation*, 91-106. London: Springer-Verlag.
- FALCAO, T.P. AND PRICE, S. 2009. What have you done! The role of 'interference' in tangible environments for supporting collaborative learning. *Proc. of the 9th international conference on Computer supported collaborative learning, (CSCL'09)* 325-334.
- FARR, W., YUILL, N., AND HINSKE, S. 2010. The Augmented Knights Castle and Social Interaction in Children with Autism. *Proc. of the 9th International ACM SIGCHI Conference on Interaction Design and Children (IDC'10)*.
- FERNAEUS, Y., AND THOLANDER, J. 2006. Finding design qualities in a tangible programming space. *Proc. CHI'06*, ACM, 447-456.
- FITZMAURICE, G., ISHII, H., AND BUXTON, W. 1995. Bricks: Laying the foundations for graspable user interfaces. *Proc. CHI '95*, ACM, 442-449.
- FJELD, M. 2002. Physical and virtual tools: Activity theory applied to the design of groupware. *CSCW*, 11, 153–180.
- FLECK, R., ROGERS, Y., YUILL, N., MARSHALL, P., CARR, A., RICK, J., AND BONNETT, V. 2009. Unpacking collaboration around the tabletop: implications for collaborative learning. *Proc. of Interactive Tabletops and Surfaces (ITS'09)*. 189-196.
- FLECK, R., YUILL, N. AND RICK, J. (submitted). Co-located Sharing in a Multi-Device Environment.
- GARROD, S., AND PICKERING, M. J. 2004. Why is Conversation So Easy? *Trends in Cognitive Sciences* 8, 8–11.
- GOMEZ, J. 2010. The Ontogeny of Triadic Cooperative Interactions with Humans in an Infant Gorilla. *Interaction Studies* 11, 353-379.
- GUTWIN, C., AND GREENBERG, S. 2002. A Descriptive Framework of Workspace Awareness for Real-Time Groupware. *CSCW* 11, 411-446.
- HA, V., INKPEN, K., WHALEN, T., AND MANDRYK, R. 2006. Direct Intentions: The Effects of Input Devices on Collaboration around a Tabletop Display. *IEEE Tabletop '06*, 177-184.
- HAIGNEY, D., AND WESTERMAN, S. 2001. Mobile (Cellular) use and driving: A critical review of research methodology. *Ergonomics* 44, 132-143.
- HARRIS, A., RICK, J., BONNETT, V., YUILL, N., FLECK, R., MARSHALL, P., AND ROGERS, Y. 2009. Around the table: Are multiple-touch surfaces better than single-touch for children's collaborative interactions? *Proc. CSCL'09*, 335-344.

- HEATH, C., AND LUFF, P. 1992. Collaboration and Control: Crisis Management and Multimedia Technology in London Underground Line Control Rooms. *CSCW* 1, 69-94.
- HINSKE, S., LAMPE, M., YUILL, N., PRICE, S., AND LABNGHEINRICH, M. 2009. Kingdom of the Knights: evaluation of a seamlessly augmented toy environment for playful learning. *Proc. of the 8th International Conference on Interaction Design and Children*. 202-205.
- HOLT, S., AND YUILL, N. (submitted). Facilitating classroom interaction and awareness of a partner in children with autism and typical development: effects of a dual control computer paradigm.
- HORNECKER, E. 2005. A Design Theme for Tangible Interaction: Embodied Facilitation. *Proc. ECSCW'05*. 23-43.
- HORNECKER, E., AND BUUR, J. 2006. Getting a Grip on Tangible Interaction: A Framework on Physical Space and Social Interaction. *Proc. CHI'06*, ACM, 437-446.
- HORNECKER, E., MARSHALL, P., DALTON, N., AND ROGERS, Y. 2008. Collaboration and interference: Awareness with mice or touch input. *Proc. CSCW'08*. ACM, 67-176.
- HUANG, E., AND MYNATT, E. 2003. Semi-Public Displays for Small, Co-located Groups. *CHI Letters*, 5, 49-56.
- JORDA, S. 2003. Sonigraphical Instruments: From FMOL to the reacTable. *Proc. of the 3rd Conference on New Interfaces for Musical Expression (NIME'03)*.
- KARAHALIOS, K., AND BERGSTROM, T. 2006. Visualizing Audio in Group Table Conversation. *IEEE Proc. TableTop*, 131-134.
- KERAWALLA, L., PEARCE, D., YUILL, N., LUCKIN, R., AND HARRIS, A. 2008. "I'm keeping those there, are you?" The role of a new user interface paradigm – Separate Control of Shared Space (SCOSS) – in the collaborative decision-making process. *Computers and Education*, 50, 193-206.
- KHARUFFA, A.S., OLIVIER, P. AND LEAT, D. 2009. Digital Mysteries: Designing for Learning at the Tabletop. *Newcastle University Computer Science Technical Report CS-TR-1171*.
- KIDD, E., AND HOLLER, J. 2009. Children's use of gesture to resolve lexical ambiguity. *Developmental Science* 12, 6, 903-913.
- KIM, T., CHANG, HOLLAND AND PENTLAND, A. 2008. Meeting Mediator: Enhancing group collaboration and leadership with sociometric feedback. *CSCW*, 457-466.
- KOZIMA, H. 2010. <http://www.myu.ac.jp/~xkozima/carebots/index-eng.html>
- MANSOR, E., DE ANGELI, A., AND DE BRUIJN, O. 2009. The Fantasy Table. *Proc. of the 8th International Conference on Interaction Design and Children (IDC'08)*, 70-79.
- MARSHALL, P., MORRIS, R., ROGERS, Y., KREITMAYER, S., AND DAVIES, M. 2011. Rethinking 'multi-user': an in-the-wild study of how groups approach a walk-up-and-use tabletop interface, *Proc. CHI'11*, ACM, 3033-3042.

- MATSUDA, M., MATSUSHITA, M., YAMADA, T., AND NAMEMURA, T. 2006. Behavioral analysis of asymmetric information sharing on Lumisight table, *TableTop. First IEEE International Workshop*, 7.
- MISCHEL, H., AND MISCHEL, W. 1983. The development of children's knowledge of self-control strategies. *Child Development* 54, 603-619.
- MOLL, H., CARPENTER, M. AND TOMASELLO, M. 2007. Fourteen-month-olds know what others experience only in joint engagement. *Developmental Science*, 10, 826-835.
- MUJDE, P. AND TECKAN, A. 2009. The role of familiarity among group members in collaborative inhibition and social contagion. *Social Psychology*, 40, 3, 111-118.
- NEARY, K., FRIEDMAN, O. AND BURNSTEIN, C. 2009. Control of Permission. *Developmental Psychology*, 45, 3, 873-876.
- NORMAN, D. 2010. Natural user interfaces are not natural. *Interactions* 17, 3, 6-10.
- NUMMENMAA, L., HYÖNÄ, J., AND HIETANEN, J. K., 2009. I'll walk this way: Eyes reveal the direction of locomotion and make passersby look and go the other way. *Psychological Science*, 20, 1454-1458.
- PEARCE, D., KERAWALLA, L., LUCKIN, R., YUILL, N., AND HARRIS, A. 2005. The task-sharing framework: A generic approach to scaffolding collaboration and meta-collaboration in educational software. In *Towards sustainable and scalable educational innovations informed by the learning sciences*, C. Looi, D. Jonassen, and M. Ikeda, Eds. IOS Press, Singapore.
- PIPER, A. M., O'BRIEN, E., MORRIS, M. R., AND WINOGRAD, T. 2006. SIDES: a cooperative tabletop computer game for social skills development. *Proc. CSCW'06*. ACM, 1-10.
- RICK, J., ROGERS, Y., HAIG, C. AND YUILL, N. 2009. Learning by doing with shareable interfaces. *Children, Youth and Environments*, 19 (1) 321-342.
- ROGERS, Y., LIM, K., AND HAZLEWOOD, W. R. 2006. Extending Tabletops to Support Flexible Collaborative Interactions. *Proc. of Tabletop*, 71-78.
- ROGERS, Y., LIM, Y., HAZELWOOD, W., AND MARSHALL, P. 2009. Equal Opportunities: Do shareable interfaces promote more group participation than single user displays? *Human-Computer Interaction* 24 (2) 79-116.
- ROGERS, Y., AND LINDLEY, S. 2004. Collaborating around vertical and horizontal large interactive displays: Which way is best? *Interacting with Computers*, 16, 1133-1152.
- ROGERS, Y. AND RODDEN, T. 2003. Configuring spaces and surfaces to support collaborative interaction. In *Public and situated displays: social and interactional aspects of shared display technologies*, K. O'Hara, Ed. Springer, Berlin, 45-79.
- ROGOFF, B., AND LAVE, J. 1984. *Everyday Cognition: Its development in social context*. Harvard University Press.
- TEASLEY, S. AND ROSCHELLE, S. 1993. Constructing a joint problem space: the computer as a tool for sharing knowledge, in S. Lajoie and S. Derry, *Computers as cognitive tools*, Hillsdale: LEA.

- RYALL, K., FORLINES, C., SHEN, C., AND MORRIS, M. 2004. Exploring the Effects of Group Size and Table Size on Interactions with Tabletop Shared-Display Groupware. *CSCW*, 6, 284–293.
- RYALL, K., ESENTER, A., EVERITT, K., FORLINES, C., MORRIS, M.R., SHEN, C., SHIPMAN, S. AND VERNIER, F. 2005. iDwidgets: Parameterizing Widgets by User Identity. *Proc. of Human-Computer Interaction - INTERACT 2005: IFIP TC13 International Conference*, 1124-1128.
- SCHMIDT, K. 2002. The problem with “awareness”. *CSCW* 11 (3–4) 285–98.
- SCOTT, S., GRANT, K., AND MANDRYK, R. 2003a. System guidelines for co-located, collaborative work on a tabletop display. *Proc. ECSCW*, 159–178.
- SCOTT, S., MANDRYK, R., AND INKPEN, K. 2003b. Understanding children's collaborative interactions in shared environments. *JCAL* 19,(2), 220-228.
- SEBANZ, N., AND KNOBLICH, G. 2003. Representing others' actions: just like one's own? *Cognition* 88, 3, B11-B21.
- SHEN, C., LESH, N., VERNIER, F., FORLINES, C., AND FROST, J. 2002. Sharing and building digital group histories. *CSCW*, 324–333.
- STEWART, J., BEDERSON, B., AND DRUIN, A. 1999. Single Display Groupware: A Model for Co-present Collaboration. *Proc. CHI'99*, ACM, 286-293.
- STOKES, P.D. 2005. *Creativity From Constraints: The Psychology of Breakthrough*. Springer.
- SUTHERS, D. 2006. Technology affordances for intersubjective meaning-making: a research agenda for CSCL. *International Journal of Computer Supported Collaborative Learning* 1, 315–337.
- TAUMOEPEAU, M. AND RUFFMAN, T. 2006. Mother and infant talk about mental states relates to desire language and emotion understanding. *Child Development*, 77, 2, 465–481.
- TOMASELLO, M. 2009. *Why We Cooperate*. MIT Press, Boston, Mass.
- TREVARTHEN, C. 1978. The concept and foundations of infant intersubjectivity. In S. Bråten (Ed.) *Intersubjective Communication and Emotion in Early Ontogeny*, (pp. 15-46). Cambridge: Cambridge University Press.
- YUILL, N., PEARCE, D., KERAWALLA, L., HARRIS, A., AND LUCKIN, R. 2009. How technology for comprehension training can support conversation towards the joint construction of meaning. *Journal of Research in Reading* 32, 1, 109-125.
- YUILL, N., HINSKE, S. AND WILLIAMS, S. (submitted). Using an augmented toy to demonstrate the role of joint attention in children's cooperative play and narrative.
- YUILL, N., STRIETH, S., ROAKE, C., ASPDEN, R., AND TODD, B. 2007. Designing a Playground for Children with Autistic Spectrum Disorders – Effects on Playful Peer Interactions. *Journal of Autism and Developmental Disorders* 37, 6, 1192-1196.

Acknowledgements. The work for this paper was supported by a grant from the Engineering and Physical Sciences Research Council for the ShareIT project (EP/F017324/1). We thank all members of the ShareIT project for their contributions, our many research participants and the Riddles project team. We benefitted from insightful comments on the paper from many colleagues and in particular from the 3 anonymous reviewers, and from Sam Holt, Paul Marshall and Julie Coultas.

Statement re prior publication: This document has not been published or submitted in this or any similar form to any prior publication. It includes summary description and analysis of other work the authors have published, and these papers are clearly referenced as appropriate.