Article Title: Annotating Distributed Scores for Mutual Engagement in Daisyphone and Beyond

Author Name: Dr. Nick Bryan-Kinns

Author address: Queen Mary University of London, Centre for Digital Music, School of Electronic Engineering and Computer Science, Mile End, London. E1 4NS. UK. nickbk@eecs.qmul.ac.uk

Abstract: Written and drawn annotations of musical scores form a core part of the music composition process for both individuals and groups. This paper reflects on the annotations made in new forms of distributed music making where the score and its annotations are shared across the web. Four kinds of annotation are identified from eight years of studies of mutual engagement through distributed music making systems. It is suggested that new forms of web based music making might benefit from shared and persistent graphical annotation mechanisms.

Article text

Whilst music has an important social role [1] and aural tradition [2], graphical sketches and written notes have formed a key part of many composers' [3] and musical groups' practice [4] for generations. They support idea generation, compositional structuring, communication, discussion, and reflection for both groups and individuals. With the advent of computer mediated music making, especially when groups of musicians are no longer co-present, graphical and textual annotations may take on new roles in managing the creative process. This paper presents observations on the emergent use of sketches and textual annotations in distributed music making where the score, music, and any additional information is shared across the internet. The focus here is not in attempting to replicate naturalistic group music making interaction across the internet using high quality video-links (e.g. [5]), or other high bandwidth communication channels, but rather, exploring what can be learnt from studies of the communication in naturalistic settings, and exploring the design of low bandwidth collaborative music systems which are referred to as *distributed music making* systems. This paper reports on the development, study, and refinement of a specific set of distributed music making tools since 2003, starting with the Daisyphone system.

Distributed Music Making

Föllmer [6], and Blaine and Fels [7] conducted early surveys of collaborative music making experiences using computers, and their origins, highlighting key characteristics from the media used through to systems' learning curves and physicality. Föllmer's survey included tracing the origins of Networked music in Live Performance, Sound Toys, Academic research, Media Artists, and technologically driven developments. Through extensive studies of 'Net Music' systems, Föllmer's research also characterised systems in terms of 'interplay with network characteristics', 'interactivity/ openness', and 'complexity/ flexibility' [8]. Distributed music systems do not attempt to replicate face-to-face interaction and so, in Föllmer's terms, may embrace or even actively exploit network characteristics such as latency, moreover, they tend to be open to a diverse set of participants, and range from quite simple to complex systems for musical expression. Since papers such as Föllmer, and Blaine and Fels, there has been an increase in interest in co-located shared musical experiences, such as the reacTable [9], and distributed music making systems such as Ocarina [10], but we still know relatively little about how people go about using these systems. Barbosa [11] captured this increase through a thorough survey of the field in 2006, and developed the Networked Music Classification Space to distinguish between Local and Remote physical location of participants, and between Synchronous and Asynchronous interaction. Barbosa's work also provides an extensive coverage of different forms of Networked Music systems from distributed composition tools to real-time performance systems, and explored the implications that the nature of networked computers have on the possible forms of music making available. Weinberg's survey [ibid] of the field instead focused on the social imperative of networked music making, and how live performance could be supported through technologically mediated interaction, or Interconnected Music Networks. More recently, Mills [12] provided a contemporary survey of networked music making systems, focusing on software platforms for music collaboration, as opposed to individual systems, reflecting the increasing maturity of the field.

What is particularly interesting about distributed music making is that all communication must be mediated through the system, unlike co-located situations where participants can see, hear, and possibly feel each other (cf. Fels and Vogt [13]). Therefore, making music across the internet with others typically involves at least some shared visual representation of the musical score

being composed or contributed to. The question explored in this paper is: what patterns of use emerge with such shared representations.

From Daisyphone to Daisyfield

In October 2003 the distributed music making tool called *Daisyphone* [14] [15] was launched [16]. At launch it received between 4 and 18 players per day from all over the world. Logs of interaction have been collected since its launch, and there are now 160Mb of log files, or approximately ten million individual interactions.

Daisyphone provides a shared loop of music (5 seconds; 48 beats) which can be edited by up to 10 people at once. In this distributed and shared score, each participant can create notes using four different voices, and can edit any notes in the score. Shared graphical annotation is also supported around and within the musical loop, or *Daisy*. The left image of Figure 1 illustrates Daisyphone in use on a desktop PC – the score is represented by the circle of dots, and the currently played set of notes is indicated by the grey line radiating from the centre which rotates clockwise over the period of 5 seconds. Shapes represent different kinds of sound (ambient electronic sound palette in C major scale consisting of bass, lead, wash, and percussion), and the shapes in the centre allow participants to select which sound type and volume they create notes with. Colours are uniquely assigned to users to provide a simple sense of identity. In essence, Daisyphone allows a small group of participants to co-create musical scores together which are shared, edited and performed in quasi real-time - participants hear the same audio produced from the shared score, but they might not be listening to exactly the same point in the score at the same time (though remember that the score is only 5 seconds long) and the score is updated in real-time as participants edit it so network delays might mean that changes do not happen at exactly the same time for each participant. The short length of the loop combined with the quasi-real-time synchronisation and the persistence of the score means that this form distributed music making is somewhat novel – a form on online improvised composition. From a user-experience perspective is slightly different to what Barbosa [ibid] refers to as Remote Synchronous Network Music systems, as the client-server updates may not happen immediately or synchronously, meaning that participants hear loops played at different times, and possibly slightly different scores for short periods of time as the updates propagate across the network - it is more akin to Remote *semi*-Synchronous Network Music. The score and graphical annotation sharing is supported using a client-server architecture similar to earlier work on WebDrumII [17].
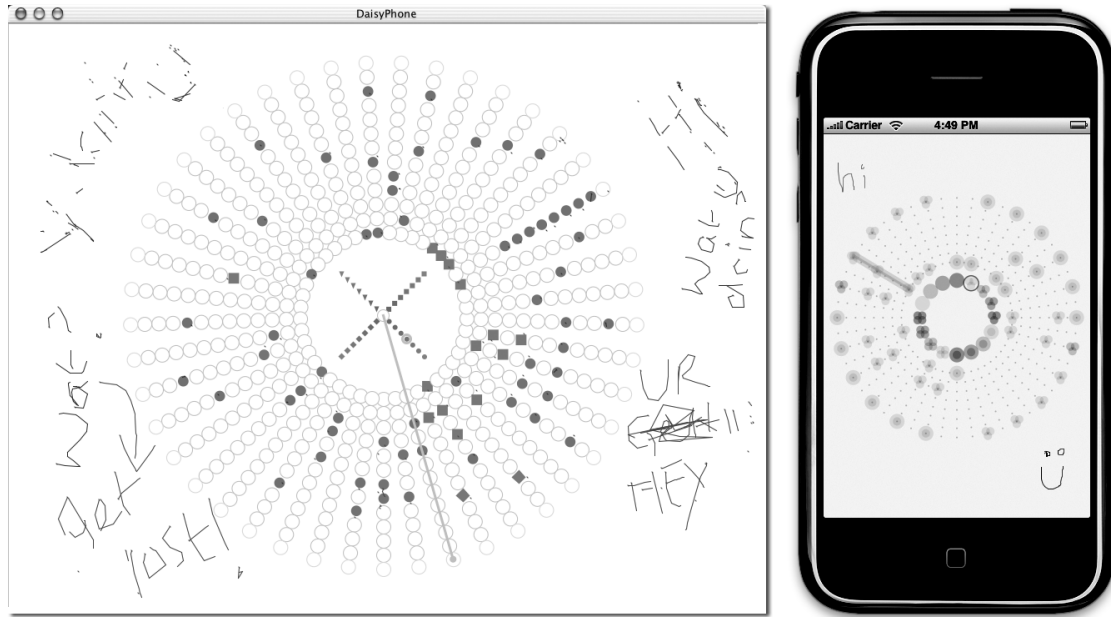
Figure 1: Dasiyphone on a desktop PC and an Apple iPhone © N. Bryan-Kinns

Early versions of Daisyphone ran on web browsers using Java, and its use was explored on desktop PCs and tablet PCs. In February 2009 a version of Daisyphone was launched for the Apple iPhone as illustrated in the right image of figure 1. This interface had less space for annotations, and a revised sound selector interface, but uses the same networking protocols meaning that people could interact with each other using web browsers and iOS devices. Since its launch, Daisyphone has been used in a wide range of environments from controlled experiments to serendipitous public uses, in exhibitions (e.g. [18]), and at live concerts.

In 2009 a richer version of Daisyphone – *Daisyfield* – was developed which allows multiple loops (i.e. Daisys) to be played concurrently, and for participants to arrange Daisys in a two dimensional space which can be graphically annotated. Again, a single shared score of 48 beats is created form the individual Daisys placed in the space, and each participant hears the same audio output. Figure 2 illustrates the Daisyfield interface with three Daisys shown, one being open for editing. The aim of this interface is to provide a richer musical and communicative user interface for exploring distributed music making.
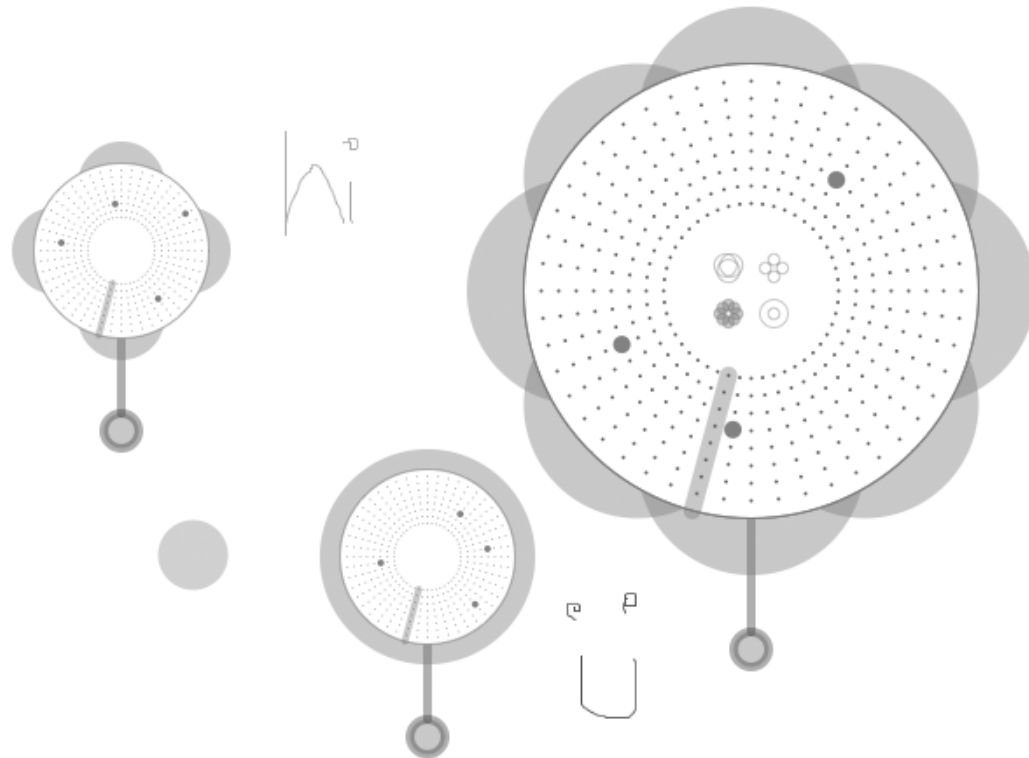
Figure 2: Daisyfield user interface © N. Bryan-Kinns

Mutual Engagement

From studies of Daisyphone and Daisyfield (collectively referred to as *Daisy**), a number of similar patterns of graphical annotations when participants *mutually engage* [19] in distributed music making have been observed which will be discussed later in this paper. Mutually engagement refers to points of interaction where participants engage with each other *and* with the product they are jointly creating – a focussed moment of group flow cf. [20] [21]. Research has explored how different cues to mutual awareness, and different features of the music making system affect mutual engagement, for example, showing through controlled experiments that providing cues to participants' identity increases their mutual engagement, and through post-hoc observations and categorisations of participants' activity that providing persistent musical scores enables learning [22]. The controlled experiments involved participants' creating short loops together over the period of an hour with systematic changes to the user interfaces which were correlated with measures of interaction including counting numbers of contributions and mutual-editing, as well as post-task questionnaires completed by participants. The key differentiator of Daisy* to co-located shared music experiences is that the activity of music making is completely mediated through the shared visual interface. As such, participants are simultaneously undertaking some aspects of the roles of performer, composer, and audience of their collective piece of music, but curiously, they may not know the actual identity of the other participants. This paper steers away from theoretical discussion of the distinctions between performer, composer, and audience in such situations to focus on the forms of interaction that emerge, and instead leaves these topics to the reader to consider. For the purposes of this paper, users of Daisy* are referred to as *participants* – they interact somehow with the technology and each other in the process of making and appreciating music as a group, or musicking [23]. The shared graphical

annotations of Daisy* then become both the instructions for the composition and the medium of social and work organisation as is discussed in the rest of this article.

Annotations

Whilst it has been shown through controlled, lab-based experiments manipulating support for annotations, that people can create music using Daisy* without graphical annotations [24] (i.e. purely through the shared score), they do appear to be much more mutually engaged when they have some form of additional communication channel, as indicated by participants' responses to post-study interviews and questionnaires. The interesting question is: what do people communicate through shared annotations when they are engaged in distributed music making. Through previous studies (Bryan-Kinns et al., ibid), and drawing on the work of CSCW researchers [25] (particularly workspace awareness – participants ongoing understanding of who has done what, and where) five categories of use of shared annotations with Daisy* were identified: Presence, Making it happen, Quality, Social, and Localisation. These categories are illustrated in this section using examples drawn from the Daisy* corpus of logfiles. Whilst CSCW research tends to be very 'office-based' in its analysis of interaction, resulting in categorisations of collaborators' discussion topics in terms of design issues, project management, meeting management, etc., this section illustrates that process of identifying categories of textual discussion can usefully be applied to understanding distributed music making systems.

Presence

Daisy* does not provide any indication of who is present in the collaboration, so unsurprisingly, quite a lot of initial annotations in a composition session revolve around finding out who is in the session, and making others aware of your presence. Figure 3 illustrates some example statements of presence (e.g. "ITZ ME FAZ"), and queries about who is present (e.g. "Anyone here?"). As participants have unique colours, the statements of presence serve both as indicators of who is collaborating, and also to identify future contributions as belonging to a specific contributor. Finding out who is also present in Daisy* is particularly important for serendipitous web based use where chance encounters lead to music making.

Figure 3: Annotations about presence © N. Bryan-Kinns

Making it happen

After identifying themselves to each other, participants usually get down to the task of making music together. Again, as Daisy* does not enforce any roles on users, it is up the participants to organise themselves somehow. Figure 4 illustrates some of the annotations participants use to co-ordinate their activity from score oriented instructions such as "Fill it up", to identifying the purpose of parts of the score e.g. "main", and task division between participants e.g. "1 ball each?" (referring to the loops in Daisyfield). Again, the free-form nature of the shared annotations in Daisy* allows for different groups to work in ways that suit them – a dynamically reconfigurable task management emerges from the participants and is not enforced by the system. This is particularly well illustrated by the use of Daisyfield's larger shared annotation area to divide the space into areas for different activities as illustrated in the three screenshots of figure 5. The left most screenshot shows participants dividing their shared space into composition (left) and discussion areas (right), the centre screenshot shows participants dividing their space into three columns (one for each participant), and the right hand screenshot shows division of the shared space into three from the centre point. This behaviour illustrates the importance of providing large shared annotations spaces for participants to construct their own instructions for the management of composition of shared scores. It is interesting to note that participants very rarely use the graphical annotations as a way of sketching out musical ideas, but rather use the space to co-ordinate the process of composition.

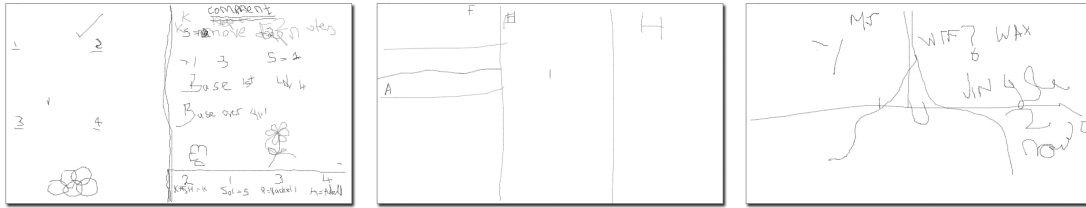**Figure 4: Making it happen © N. Bryan-Kinns**

Figure 5: Screen division in Daisyfield © N. Bryan-Kinns

Quality

Once the process of making music was underway, numerous comments about the quality of the music being produced e.g. "this sounds awesome" and "do u like it?" as illustrated in figure 6 can be observed in controlled studies, observations of public events, and logs of online usage. These emerging discussions are critical to the success of the collaboration and provide a mechanism for participants to reflect on their music making, and to engage in quality judgements – moving the focus of the discourse to the key topic of making good music together.
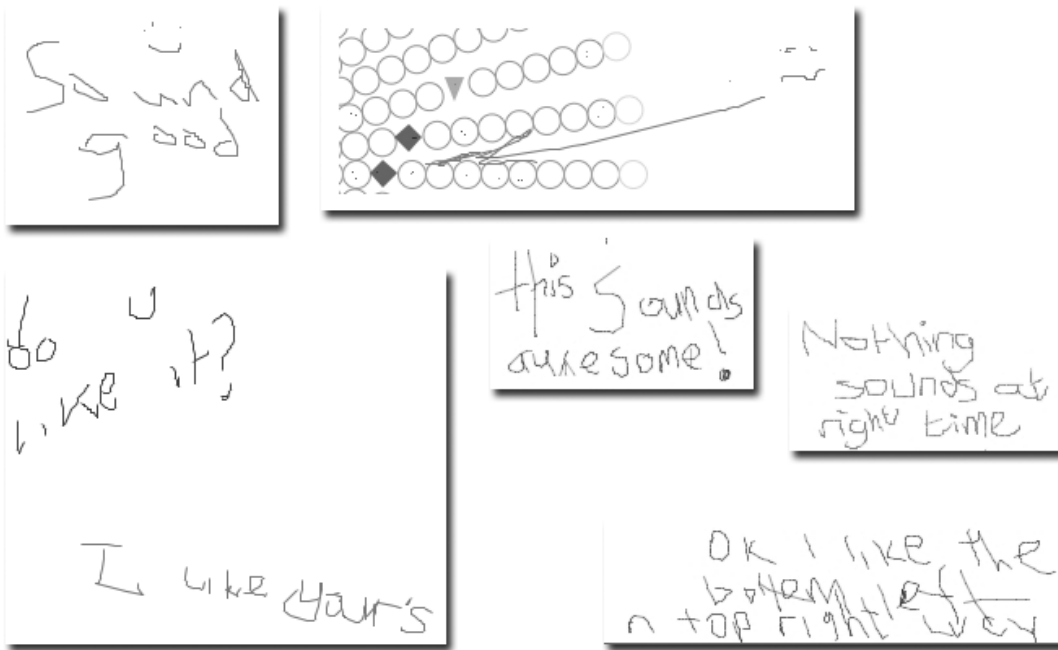

Figure 6: Annotations about quality © N. Bryan-Kinns

Social

Aside from actually making the music and managing their online presence, participants routinely engaged in social, or non-task related discussions. Figure 7 illustrates these words or images that are not connected to any of the categories of annotation above, for example, drawings of faces, or "LOL" not directly connected to any parts of the shared score. Many of these annotations were used in humorous ways e.g. drawing a funny looking fish, or smiley faces. These light hearted contributions form the social glue of the interaction rather than contributing to the musical score per se.

Figure 7: Social annotations © N. Bryan-Kinns

Localisation

Across the four categories of textual annotations outlined above, a key use of the shared annotations in Daisy* is localisation – identifying important areas, parts, and words, and connecting them together. Typically this is achieved by drawing arrows as illustrated in figure 4 – one participant has written "check it" with an arrow pointing to a highlighted part of the score. Similarly, in figure 6 a participant has connected a sad-smiley face to a particular piece of the score to indicate their feelings about it. These forms of localisation are critical to the mutually engaging refinement of the shared scores – they allow participants to directly embed their thoughts in the shared representation, and highlight pertinent parts of the shared score. In other words, the ability to localise comments to musical objects supports the work of group music making.

Reflections and Summary

The forms of shared annotations outlined in this article embody the communication that has been observed emerging around a shared and distributed musical score – essentially these emergent uses of annotations are the likely to be the kinds of things that people will want to be able to communicate about in distributed music making. It is interesting that similar patterns of communication emerge on the web and on mobile, and in the range of different contexts from controlled studies to public exhibitions. Only the size of the drawing area makes significant changes to the use of annotation with Daisyfield's large drawing area being used more to co-ordinate the task of distributed music making than Daisyphone's. However, it is worth noting that Daisy* are very simple Networked Music systems which do not provide the rich and immersive group music experience of tele-present or co-located systems, partly in order to focus on what happens when rich human-interaction channels are not available. Where other communication channels are available, the role of shared annotations may decrease, and may return to being

annotation for later review rather than annotation for real-time interaction. It would be interesting to explore how shared annotation could usefully be deployed in richer environments.

In designing future web based music making systems, shared annotation mechanisms which support task management, presence awareness, social interaction, and, importantly, localisation, seem to be important. A conventional text chat system would struggle to support all these in a fluid and engaging way, especially localisation of annotations as text may appear in a separate text window to the score itself. This has implications for systems which Barbosa refers to as Networked Music systems, both Synchronous and Asychronous. Tools such as Ocarina provide cues to identity of creators, but do not explicitly support collaborative editing of distributed scores. Moreover, although they provide a 'like' tag to allow people to comment on the quality of each others' compositions, there is little social interaction supported. Similarly, interactive tabletop interfaces such as reacTable can be linked together to create distributed music making environments, providing a Local Networked Music experience, but they provide no annotation mechanisms even though tabletops provide an intuitive drawing area. It would be interesting to explore future versions of such systems which combine their playful music making interfaces with richer shared editing and annotation mechanisms. Supporting shared annotation in and around the score essentially embeds the creative process in the score, creating a rich and engaging multi-modal experience. However, the persistent nature of annotations in Daisy* may reduce their utility in longer term collaborations – some form of erasing mechanism by users, or over time, may be necessary to reduce screen clutter and confusion. Also, Daisy* has primarily been used to support the process of shared composition, so it would be interesting to explore their use in more performative forms of interaction (cf. [26]) where mutual engagement between participants is more improvisational than compositional.

Distributed music making relies on task management, knowing who is collaborating with you, and fun social interaction to glue it together. Daisy* provide elements of all of these in addition to the ability to localise annotations using a simple shared graphical annotation mechanism. Future work will need to explore the role of persistence of the annotations and the score itself, as well as richer and more performative music making experiences.

Acknowledgment

References and Notes

1. G. Weinberg, "Local Performance Networks: musical interdependency through gestures and controllers," Organised Sound, Vol. 10, No. 3 (December 2005), 255-265 (2005).
2. J. T. Titon, *World of Music* (NY USA: Schrimer Books, 1996).
3. J.-B. Thiebaut, P. G. T. Healey, and N. Bryan-Kinns, "Drawing Electroacoustic Music," In Proceedings of International Computer Music Conference (ICMC 2008), Belfast (2008).
4. S. Nabavian, and N. Bryan-Kinns, "Analysing Group Creativity: A Distributed Cognitive Study of Joint Music Composition," In Proceedings of Cognitive Science 2006, 1856-1861 (2006).
5. A. A. Sawchuk, E. Chew, R. Zimmermann, C. Papadopoulos, and C. Kyriakakis, "From Remote Media Immersion to Distributed Immersive Performance," In Proceedings of ETP '03, Berkeley, California, USA (2003).
6. G. Föllmer (translated by Martin Watson), "Lines of net music," Contemporary Music Review, Vol. 24, No. 6, 439 – 444 (2005).

7. T. Blaine, and S. Fels, "Contexts of collaborative musical experiences," In Proceedings of the 2003 Conference on New Interfaces for Musical Expression (NIME-03), Montreal, Canada, 129–134 (2003).

8. G. Föllmer, "Electronic, aesthetic and social factors in Net music," Organised Sound, Vol. 10, No. 3, 185–192 (2005).

9. S. Jordà, M. Kaltenbrunner, G. Geiger, and R. Bencina, "The reacTable,"
In Proceedings of the ICMC2005, Barcelona, Spain (2005).

10. G. Wang, "Designing Smule's iPhone Ocarina," In Proceedings of the International Conference on New Interfaces for Musical Expression, Pittsburgh (2009).

11. A. Barbosa, *Displaced Soundscapes: Computer-Supported Cooperative Work for Music Applications* (Spain: Universitat Pompeu Fabra, PhD Thesis, 2006).

12. R. Mills, "Dislocated Sound: A Survey of Improvisation in Networked Audio Platforms," In Proceedings of NIME 2010, Sydney, Australia, 186-191 (2010).

13. S. Fels, F. and Vogt, "Tooka: Explorations of two person instruments," In Proceedings of the 2002 Conference on New Instruments for Musical Expression (NIME-02), Dublin, Ireland (2002).

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

15. Daisyphone.net

16. P. Marks, "Will Jamming be the New Texting?," New Scientist, Vol. 180, No. 2418, 25 (25 Oct 2003).

17. P. Burk, (2000) "Jammin' on the Web - a new Client/Server Architecture for Multi-User Musical Performance," Presented at ICMC 2000.

18. *Surface Tension*, Dana Centre, Science Museum, London, UK, http://www.danacentre.org.uk/events/2009/06/23/506 (23 June 2009).

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

20. M. Csikszentmihalyi, *Flow: The Psychology of Optimal Experience* (Harper Collins, 1991).

21. K. Sawyer, *Group creativity: Music, theater, collaboration* (NJ USA: Lawrence Erlbaum Associates, 2003).

22. N. Bryan-Kinns, and P. G. T. Healey, "Decay in Collaborative Music Making," In Proceedings of NIME 2006, Paris, France, 114-117 (2006).

23. C. Small, *Musicking: The Meanings of Performing and Listening (Music Culture)* (Wesleyan University Press, 1998).

24. N. Bryan-Kinns, and P. G. T. Healey, "Exploring Mutual Engagement in Creative Collaborations," In Proceedings of Creativity and Cognition 2007, Washington, USA, 223-232 (2007).

25. J. S. Olson, G. M. Olson, M. Storrøsten, and M. Carter, "Groupwork close up: a comparison of the group design process with and without a simple group editor," ACM Transactions on Information Systems, Vol. 11, No. 4, 321-348 (1993).

26. J. G. Sheridan, and N. Bryan-Kinns, "Designing for Performative Tangible Interaction," International Journal of Arts and Technology, special Issue on Tangible and Embedded Interaction, Vol. 1, No. 3-4, 288-308 (2008).

Biographical Information: Dr. Nick Bryan-Kinns is a Senior Lecturer in Graphical User Interface Design and lead of the Interactional Sound and Music Group at Queen Mary, University of London. His research focuses on exploring mutual engagement in creative collaborations and has been a panel member for the NSF CreativeIT programme, chair of the ACM Creativity and Cognition 2009 conference, co-editor of special issues of journals such as CoDesign.