
Steve Benford

Mixed Reality Laboratory
School of Computer Science
The University of Nottingham
Jubilee Campus
Wollaton Road
Nottingham NG8 1BB UK
sdb@cs.nott.ac.uk
www.mrl.nott.ac.uk

The field of Human–Computer Interaction (HCI) has long been interested in how people interact with digital technologies, including—through the closely related field of Computer-Supported Cooperative Work (CSCW)—how they collaborate through and around these technologies. Although initially focused on office applications and work, the spread of digital technologies into nearly every aspect of our everyday lives has led these fields to increasingly focus on emerging leisure, entertainment, and cultural applications of digital technologies in areas such as games, museum installations, interactive artwork, and, of course, playing and listening to music. In its turn, the focus of studying and designing interfaces has also shifted from issues of usability and productivity to encompass new goals such as pleasure, creativity, expression, and aesthetics.

For more than a decade, research at Nottingham’s Mixed Reality Laboratory has explored the use of digital technologies in live performance. This has involved working with artists to create, tour, and study a series of theatrical experiences that mix fictional stories with real settings, virtual environments with physical sets and props, and interaction with computers with live encounters with actors and other participants. Various examples have shown how digital technologies can be embedded into extended theatrical performances, including *Can You See Me Now?*, a game of chase in which on-line players logged in over the Internet were chased through a 3-D virtual model of a city by actors who, equipped with handheld computers with GPS receivers, had to run through the actual city streets to catch them; *Uncle Roy all Around You*, in which on-line and “street” players collaborated to navigate a mixed real and virtual cityscape, encountering various actors, props, and settings on

Computer Music Journal, 34:4, pp. 49–61, Winter 2010
© 2010 Massachusetts Institute of Technology.

Performing Musical Interaction: Lessons from the Study of Extended Theatrical Performances

the way (Benford et al. 2004); and *Fairground: Thrill Laboratory*, which used bio-sensing technologies and wireless communications to transform the act of riding a rollercoaster into a public performance. An overview of several of these performances can be found in Benford et al. (2009). These experiences were also the subject of ethnographic studies in which observation of participants, including the public, actors, and technical crew, revealed the fine details of how the interactions were delivered and experienced.

Reflecting on these experiences and studies led to the development of various theories to account for the design and experience of performance interfaces. This article takes these theories, alongside others from HCI and CSCW, and considers how they might be relevant to the design of musical interfaces, identifying key issues and approaches that might inform an agenda for future work in this area. The argument unfolds by following a trail of ever-widening participation in a musical performance, from an initial focus on the issues that arise when just one musician interacts with their digital instrument, through consideration of ensemble playing, to different ways in which interfaces might address an audience, to the embedding of musical interfaces within an extended performance structure.

Interacting: The Musician and Their Instrument

The first thing to note is that there are many traditional forms of interaction with instruments (plucking, bowing, and strumming strings; pressing keys; striking drums; and so forth) that are not the primary focus of this article. Also out of scope are mainstream interfaces in which desktop displays, mice, keyboards, and similar devices are used to interact with musical software tools. Rather, the

focus of attention is on emerging forms of interface that might enable particularly interesting and alternative forms of musical performance.

From enhancing traditional instruments (Bevilacqua et al. 2006; Poepel and Overholt 2006), to creating modified digital instruments (Jordà et al. 2007; Warming Pedersen and Hornbaek 2009), to attaching sensors to their own bodies (e.g., Pamela Z; see Lewis 2007), performers have employed sensing-based interfaces to lend greater expression to their playing, allowing interaction with digital music through gestures and other bodily or facial movements. Such interaction via sensing-based interfaces is often indirect in the sense that the musician is not immediately physically connected to their instrument, and the sensors may even be invisible, potentially allowing the kind of untethered and unencumbered interaction that could, for example, support a more seamless integration of music with dance.

However, interacting with these kinds of sensor systems can be challenging, owing in large part to their invisible nature, which is often combined with a relatively high level of unreliability, at least when compared to the operation of traditional buttons, key, sliders, and wired devices. Bellotti and colleagues have articulated these challenges in terms of five questions for the designers of sensing-based interaction (Bellotti et al. 2002):

- (1) Address: How do I address one (or more) of many possible devices?
- (2) Attention: How do I know the system is ready and attending to my actions?
- (3) Action: How do I effect a meaningful action, control its extent, and possibly specify a target or targets for my action?
- (4) Alignment: How do I know the system is doing (has done) the right thing?
- (5) Accident: How do I avoid mistakes?

One response to these questions takes the form of a framework that encourages the designers of sensor-based interfaces to systematically explore a space of partially overlapping expected, sensed, and desired movements and actions, consciously seeking out misalignments between them (Benford et al. 2005).

Expected movements are those that the user might normally be expected to make owing to a combination of their prior expectations, any metaphor associated with the interface, its physical affordances and constraints, and also the ergonomics of their own bodies. In some cases, such as when augmenting traditional violins bows with sensors, the repertoire of expected movements might be relatively predictable (Bevilacqua et al. 2006), whereas with new instruments, it might be more emergent. Importantly, this framework also explicitly encourages designers to playfully envisage less expected, unusual, or even impossible movements of the interface alongside those that a user might normally be expected to make. How might it be possibly manipulated in bizarre and unlikely as well as normal ways?

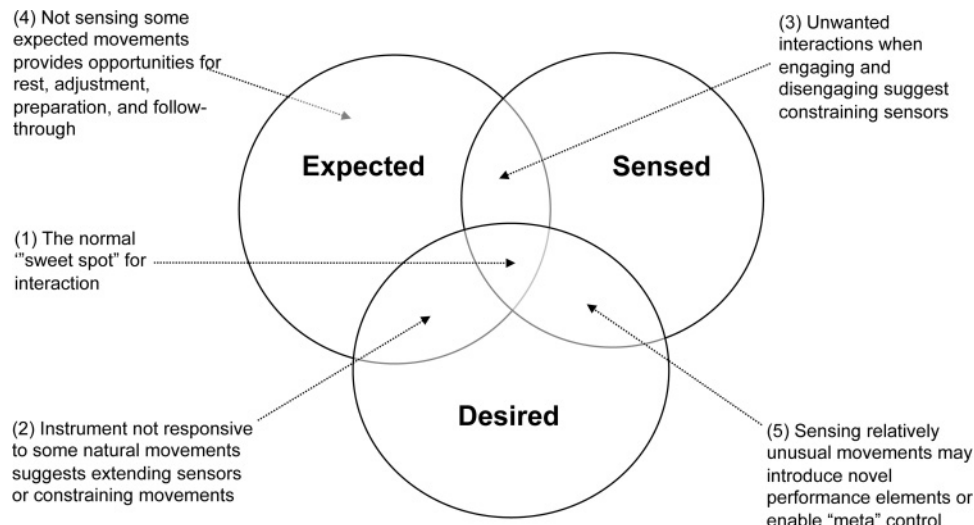
Sensed movements are those that can actually be measured by the interface's sensing systems. Here, designers are asked to chart out the range, degrees of freedom, responsiveness, and accuracy of each available sensor, and then to consider these in combination. They are also encouraged to explicitly identify any movements that might not be sensed for some reason, such as limited coverage of the sensors, their responsiveness, interference, and the impact of other environmental conditions.

Finally, *desired* actions are those that result in appropriate, useful, or otherwise engaging functionality. Again, interface designers are asked to consider what kinds of movements and actions might be undesirable.

The key point is that these three facets of interaction with sensors frequently only partially overlap and so can be considered in terms of the Venn diagram shown in Figure 1. The authors of the framework argue that many interface designs limit their view to the "sweet spot" where expected, sensed, and desired movements overlap, but that the other areas of only partial or even non-overlap are also interesting to designers, both to identify potential problems with the interface and as a source of inspiration for new opportunities.

Even a brief exploration of this design space reveals some interesting issues for the design of musical interfaces. In this case, the sweet spot—labeled (1) in Figure 1—for designing the interface to an instrument is that point where the musician's natural

Figure 1. *Expected, sensed, and desired interactions for musical performance.*



expectations and most fluid physical movements naturally map onto the capabilities of the sensor system so as to produce a desirable response from the instrument. Unsurprisingly, this confluence of expected, sensed, and desired actions tends to be the dominant focus for much interface design: It makes obvious sense as a design strategy to build on the most common or natural movements, work within the capabilities of sensors, and produce pleasing results. However, the designers of musical interaction might benefit from also exploring other less obvious parts of the design space.

Movements that would be expected and ideally desirable, but that cannot be sensed—labeled (2) in Figure 1—are a potential problem, suggesting that the range of the instrument is limited by the constraints of the sensing technology; in other words, the instrument is not responsive enough. Designers may therefore wish to explore whether the range or other capabilities of the sensors can be extended somehow, or if this is not possible, they may need to design the interface to communicate these constraints to the musician, perhaps by introducing physical constraints or visual markers. Such strategies can avoid frustration as the musician tries to trigger sounds that they would naturally expect to play but in fact cannot, especially in situations where they are not intimately familiar with the instrument through many hours of practice.

Conversely, some movements that are expected and that can be sensed may actually have undesirable musical consequences—labeled (3) in Figure 1. A common problem in this space occurs when a musician first approaches an instrument and engages with it prior to beginning playing (e.g., donning wearable sensors at the start of a performance), which can trigger clumsy, unmusical interactions. Similar problems can occur when they set the instrument down again or hand it to another performer. This is also the territory of glitches caused by inaccurate or jittery sensor systems that need to be smoothed out, or, if this is not possible, the musician’s expectations of fine-grained control may need to be relaxed in favor of them anticipating a less predictable or more ambiguous response from the instrument. In these cases, designers need to constrain or program the system to ignore some kinds of expected movement.

Turning to another part of the design space, there are some movements that are to be expected, and cannot be sensed, and where it would not be desirable to trigger musical interaction—labeled (4) in Figure 1. As with a golfer’s swing, the actual moment of interaction may be preceded by a preparatory movement and succeeded by a follow-through movement, both of which are vital to the successful performance of the overall interaction but do not directly trigger any sound. Similar movement may

be essential to allowing the musician to temporarily disengage from the instrument so as to reposition or rest before reengaging again (the equivalent of shifting position on a traditional keyboard or fret board). More broadly, other gestures that naturally occur “around the instrument” but that are not sensed and thus do not trigger music can lend a degree of physical expression to a performance, revealing the performer’s emotional engagement as well as their skill and control over the instrument. (Think here of the movements of a pianist’s hands when they are not actually striking a key.) Rosen (2002) describes how a performer’s gestures at the piano influence spectators’ appreciation of the skill and emotion involved in the performance of a piece of music, whereas Sudnow (1978) describes how seemingly extraneous gestures become part of the practice of productions at the keyboard. Wanderley et al. (2005) conducted an exploratory study of what they term the “ancillary gestures” that are made by clarinetists. By analyzing video recordings alongside data gathered from movement sensors, they found that ancillary gestures are an integral aspect of musical performance; that they tend to be consistent for a given performer across multiple performances; that performers can be grouped based upon the parts of the body that they tend to move (e.g., knees vs. waist); and that there are two dominant trends of movement in relation to groupings of notes: regular and consistent rhythmic movements versus flourishes at the endings of phrases. In a similar vein, previous HCI research has discussed the role of performative gestures in playing electronic instruments, using the term “expressive latitude” to refer to performance gestures that are not directly sensed by the instrument (Bowers and Hellström 2000).

Finally, it may be useful to limit interaction with the instrument to take place in one or more local “hotspots” on a broader stage of movement, for example restricting musical interactivity to certain areas of the stage during a wider dance production in which dancers may only occasionally wish to trigger music as a special effect. In short, there are several compelling reasons why designers should deliberately build in to a musical interface opportunities for movements that are to be expected,

but that are not sensed and so do not trigger music.

A final interesting area of the overall design space occurs where it may be desirable to deliberately sense relatively unusual movements—labeled (5) in Figure 1. Such movements might allow innovative musicians to experiment with novel or extreme ways of playing the instrument in which they must push themselves into unusual positions and actions to create particular sounds, which could lend an interesting dynamic to their performance. More mundanely, it may be useful to have a class of relatively unusual gestures that trigger meta-level control of the instrument, changing its tonal and other parameters, without having to resort to “out of band” controls such as the foot-pedals and additional switches and buttons that are routinely used with electric guitars and keyboards. In other words, the performer might be able to use the same underlying interaction mechanism to both play and configure the instrument, although the two sets of gestures may have to be clearly separable to avoid confusion, with playing gestures being more expected and configuration ones less so.

In summary, although we have avoided an exhaustive presentation of the *expected-sensed-desired* framework, this discussion demonstrates that designing a sensor-based musical instrument involves both opportunities and challenges that arise from having only partial overlaps between expected movements, those that can be sensed, and desired outcomes. In particular, the framework may help designers identify several key aspects of performative musical interaction that need to be considered including supporting gestures around the instrument that are not sensed; supporting meta-control of the instrument as well as the direct generation of music; and considering how the instrument is picked up, set down, and handed over to others.


Collaborating: The Ensemble and Their Instruments

So far, we have discussed interaction between a single musician and a single instrument. However, for much of the time music is played in ensembles,

Figure 2. Multichannel collaboration around an instrument. Adapted from Dix (1997).

leading us to the topic of how multiple musicians interact with multiple instruments. The field of CSCW has many things to say about the design of collaborative interfaces that are frequently summarized under a “space–time” classification, which, as a high-level generalization, divides collaborative interfaces into those supporting face-to-face interactions (at the same time and in the same place); remote interactions (same time but different places); continuous tasks (different times but the same place); and communication and coordination (different times and different places; Johansen 1988). Although all of these broad modes of collaboration are relevant to music-making in general, especially when we extend our view to collaborative composition and various forms of music distribution and sharing, we restrict our discussion here to those aspects that affect live performance, i.e., that take place at the “same time.”

The starting point for our exploration of live collaboration considers the situation in which several musicians share a common musical instrument. Although this can happen with traditional instruments (e.g., two pianists playing a duet on the same piano), it is a relatively rare occurrence. CSCW, however, has considered various technologies to support collaboration around and through a shared interface including tabletop displays and other tangible interfaces in which everyday objects can be used to interact with shared surfaces (Ishii and Ullmer 1997), and also wall displays and “roomware” (Streitz et al. 1999), many of which fall under the general heading of Single-Display Groupware (Stewart, Bederson, and Druin 1999). Such technologies are now finding their way into musical instruments, for example the *reactTable*, a modular synthesizer with a multi-touch tangible interface that was used by the musician Björk during her 2007 world tour (Jordà et al. 2007). Blaine and Fels (2003) have conducted an extensive review of how such co-located displays, especially when deployed in public environments, can support social music-making by novice rather than virtuoso musicians. Through an analysis of eleven examples, they identify key factors in the design of such instruments including their degree of focus towards the audience or performers, the location of the interface, the media involved, the




level of scale in terms of the number of players, the nature of individual interfaces, the extent to which these enable physical interactions, whether players have identical interfaces, the musical range of the instrument, the extent to which interactions are directed, the learning curve for the instrument, and, following on from this, the pathway to more expert performance (Blaine and Fels 2003).

These new forms of musical instruments in which multiple musicians share a single display raise significant interaction–design challenges, not least of which is the multi-channel nature of collaboration which, as Dix (1997) articulates, involves both direct coordination between musicians who can see one another and communicate, and feedthrough in which each musician’s interactions with the interface are indirectly passed onto the other via the interface itself. Both of these channels—direct coordination and feedthrough—must be considered in the design of a musical instrument, as shown in Figure 2. On the one hand, how do its shape, size, layout, and placement in the local performance setting afford direct coordination between musicians? Can they easily see one another, share talk and gestures, and witness each other’s interactions? On the other, how does the instrument itself provide feedthrough that reflects the players’ actions upon it? Are these actions highlighted in any ways and identified with the different musicians?

A situation perhaps more familiar to many musicians is one in which each person brings their own instrument. With electronic instruments, these can then also be networked together using protocols such as MIDI to create further possible channels of communication. Figure 3 shows a typical situation in which two musicians have networked their instruments such that feedthrough passes through

Figure 3. Feedthrough and direct coordination around networked instruments.



each instrument and over the network to the other. Although this is a common performance situation, it is notable that the current generation of commercial electronic instruments does not generally support feedthrough: There is generally a lack of representation on their displays of other instruments, musicians that are on the network, and the actions they are performing (e.g., of the notes they are playing or the settings of their instruments). One immediate lesson from CSCW then is that future instruments might incorporate this kind of feedthrough to better support collaboration.


This argument may not at first be fully convincing in purely face-to-face situations; after all, groups of musicians appear to be able to use electronic instruments reasonably well with just channels of direct coordination (e.g., using glances, nods, and talk to coordinate their actions across a shared stage). The introduction of remote interactions, however, changes the situation greatly. Now we must consider a fellow musician who is not physically co-present, and so there is no immediate channel of direct coordination available. A common solution is to use separate video and audio links to restore this channel so that the musicians can now see one another over a remote link. However, a host of CSCW systems and related studies have highlighted the benefits of also providing a feedthrough channel by embodying users within the digital space of the interface itself. For example, shared drawing tools have effectively enhanced remote collaboration through telepointers that convey the presence, identity, and activity of remote others, and virtual worlds routinely use avatars to embody their users within a 3-D digital space (Benford et al. 1995).

Feedthrough may be especially useful for dealing with the effects of network latency, which are especially challenging for distributed musical performance (Zimmermann et al. 2008). Studies of both distributed drawing tools and simple networked 3-D games have shown that enhancing users' visual embodiments with indications of the current level of network delay can improve their performance of collaborative tasks and have led to proposals for a variety of "decorators" that can be attached to embodiments (Stach et al. 2007).

Beyond remote interactions, there is also the need to support orchestration work—the process of shaping and guiding a performance, often invisibly from behind the scenes. Observational studies of a sequence of interactive theatrical performances have repeatedly stressed the importance of orchestration work and the need to support this with dedicated interfaces to support monitoring, intervening, and communicating (e.g., Koleva et al. 2001; Crabtree et al. 2004). Of course, musical performance also routinely involves orchestration work, carried out by technical groups including a sound crew, lighting crew, and others; for large performances, each crew can involve many individuals who are themselves distributed across multiple locations. These individuals also require channels of communication if they are to coordinate with the musicians and with each other, including channels of direct coordination as well as feedthrough channels through which they can remotely monitor (and even intervene in) the state of individual instruments.

Drawing these various threads together, we can see that ensemble playing can involve quite complex ecologies of instruments, musicians, and technicians. Figure 4 shows a general case that combines different forms of collaborative interfaces (single-display groupware and networked individual instruments) in both face-to-face and remote modes. The key lesson from CSCW is that designers must consider the use of both channels of direct coordination and feedthrough in each case, whether they are needed, and if so, how they can be best supported. Moreover, while direct coordination is already quite well supported through side channels such as walkie-talkies, monitor systems, or simply the careful arrangement of the

Figure 4. An example ecology of distributed instruments, musicians, and technicians.



performance space, feedthrough channels in which information about the use of instruments flows through the network connections between them alongside MIDI performance data to be displayed on the interfaces of instruments, appears to be far less well supported—probably much less so than in many other non-musical collaborative applications where it has proved to be a very successful.

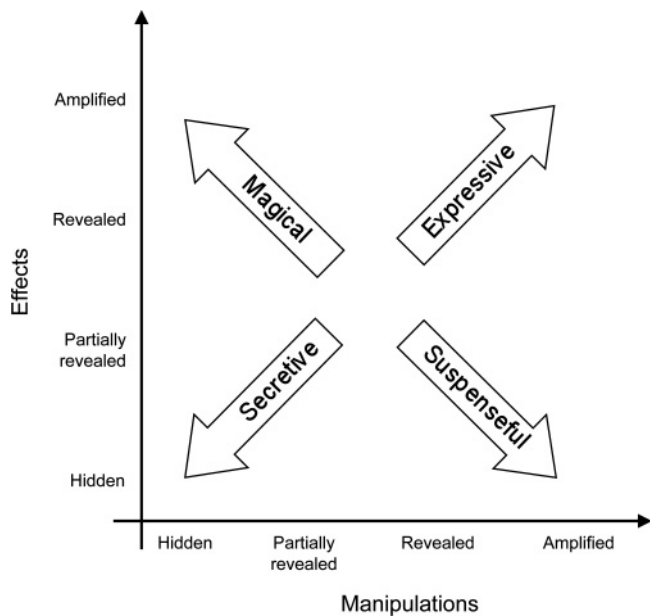
Spectating: Addressing the Audience

A performance also involves an audience, and so we now turn to the question of how this audience experiences the musicians' interactions with their instruments and with each other, an issue referred to as "transparency" by Blaine and Fels (2003). As computers have increasingly spread out of the workplace into public settings such as cafes, bars, and city streets, so our interactions with them have taken on many aspects of being a public performance. Some of these performances are consciously part of new forms of theatrical event, staged by artists who are drawn to the potential of mobile, wearable, embedded, and other increasingly ubiquitous interfaces to interweave digital media with the everyday world around us, which then become a rich backdrop for new theatrical experiences. Others are more implicit everyday performances, such as when mobile phone

calls are consciously "performed" with a local audience in mind. In either case, the field of HCI has become increasingly concerned with how interface designers can address the performative nature of public interaction or, put another way, how they can create interfaces that reflect the needs of nearby spectators (deliberate or accidental) as well as of their direct users.

A recent taxonomy of spectator interfaces classified various public and performance interfaces, including some musical examples, in terms of the extent to which they hide, reveal, or even amplify a performer's manipulations of the interface compared with the extent to which they hide, reveal, or amplify the effects of these manipulations (Reeves et al. 2005). Manipulations here include both direct inputs to the interface that trigger interactions, but also the kinds of expressive gestures around the interface that were previously discussed. In turn, effects include the direct outputs that are displayed by the interface, but also the visible effects of these in the performer. A classification of many different interfaces, from everyday mobile phones, laptops, kiosks, and projection interfaces, to modified instruments and installations revealed that different styles of public interface can adopt radically different strategies with regard to hiding or revealing combinations of manipulations and effects to nearby spectators. However, these can be grouped under four high-level design strategies for spectator interfaces, as depicted in Figure 5.

Figure 5. Strategies for designing spectator interfaces.



The horizontal axis shows whether the manipulations of a given interface are hidden, partially revealed, revealed, or even amplified in some way (e.g., by making the controls particularly large or otherwise drawing attention to them). The vertical axis does the same for effects. One popular design strategy, *expressive*, aims to make both manipulations and effects as visible as possible, emphasizing the connection between the two and hence the virtuosity of the performer in being able to control and manipulate the interface. This is an approach that is well suited to the design of musical instruments, and the reviewed examples included Toshio Iwai's Piano that used dynamic lighting to augment his interactions with the keyboard (Wilson 2002), and also Pamela Z's use of body-worn sensors to enable her to use large, effectively amplified, gestures to produce music (Lewis 2007).

The opposing strategy is *secretive*, in which both manipulations and effects are hidden from spectators, in the extreme case by hiding the interface behind barriers such as curtains, or in kiosks that are deployed in public places but that wish to maintain a sense of privacy for the user (e.g., passport photo kiosks). A related facet of privacy considered by subsequent work is the separation

of official spectators (e.g., who hold tickets) from others, traditionally accomplished through the more permanent barriers of the auditorium walls. Recent writing in HCI has considered how the boundaries or "frame" of an interactive performance become increasingly blurred in public settings and makes a distinction between the audience who are within the performance frame (i.e., who are aware that a performance is taking place and are able to interpret the action appropriately) versus possibly unwitting bystanders who may be passing through the performance space and may be less able to interpret what is taking place (Benford et al. 2006). New modes of performance such as "Flash Mobs," which might potentially include performing music, may involve apparently spontaneous outbreaks of performance and thus deliberately blur the performance frame, introducing a degree of ambiguity as to who are the performers, audience, and bystanders.

Our third strategy, *magical*, involves revealing effects but hiding away the manipulations that caused them. This strategy is clearly relevant to the public performance of illusions as part of stage magic, but more generally, enables a performer to hide clumsy interactions that might detract from the overall aesthetic of the performance. A recent study of a computer-vision system being used to support a stage magic trick showed how magical interfaces may even deliberately exploit the multi-channel nature of interaction around and through shared interfaces as discussed earlier, by deliberately misdirecting the spectator's attention from direct to feedthrough channels or vice versa, so as to reveal manipulations that support the fiction of the illusion while hiding others that are concerned with the way in which it is actually achieved (Marshall, Benford, and Pridmore 2010). In terms of musical instruments, the magical strategy might be applied to interactions that are concerned with the "meta-level" control of instrument settings, rather than the immediate production of musical sound, often accomplished by the use of foot-switches, pedals, and similar devices that are not always immediately visible to the audience. Another example lies in the design of "augmented instruments," traditional instruments that are extended with (often invisible) sensors so

that apparently conventional musical interactions produce surprising additional effects. In practice, the design of musical instruments might benefit from applying both the expressive and magical strategies in tandem, revealing those manipulations that convey the performer's skill in playing (making gestures, plucking strings, and so forth) while hiding other more mundane interactions that might control less-interesting parameters of the instrument.

Finally, the strategy of creating *suspenseful* spectator interfaces involves revealing manipulations to the spectator while hiding their effects. At first glance, this may be the most counterintuitive strategy of the four: Why would an audience not be able to see or hear the output of an instrument or other display? However, it may be a useful strategy in some situations, especially where audience members are waiting in line to experience an installation and can benefit from seeing in advance how others use it (so that they know what to do when their turn comes), but without seeing the "payoff" until it is actually their turn. An example discussed in Reeves, Fraser, and Benford (2005) involves a museum augmented-reality display in the form of a telescope that overlaid video images on a collection of bottles displayed on a pedestal. Bystanders could see from the actions of others that they had to approach the telescope, look through the eyepiece, and rotate the display, but they only saw the payoff (the video augmentations) when they took their turn. Musically, this strategy is perhaps best suited to the design of public sound installations in exploratoria, museums, galleries, and similar settings.

In summary, interfaces, including musical instruments, must be designed with spectators in mind, and previous work in HCI has identified a range of complementary strategies for approaching this task in terms of hiding, partially revealing, revealing, or even amplifying different combinations of manipulations and effects, and it has also discussed related issues concerning the framing of performances in public spaces.

Trajectories Through an Entire Performance

In this final section, we further expand our perspective to consider the manner in which interaction

with a digital instrument can be successfully embedded into the wider context of a musical performance. HCI is increasingly turning its attention to the question of how designers can understand and create entire user experiences (Law et al. 2008). One response has been to draw on expertise from performance studies to develop a theoretical account of the nature of "mixture reality performance," i.e., of the emerging genre of performances that combine physical and virtual spaces in various ways to create mixed-reality stages and that also combine elements of live performance with interaction with digital technologies (Benford et al. 2009). Mixed-reality performances reported in the literature appear to be extremely complex, combining multiple physical and virtual spaces, multiple timescales, different performative roles, and diverse interfaces into complex hybrid structures. However, it has been proposed that they can be understood using the overarching concept of "trajectories." Inspired by recent writing about the nature and history of lines and of the importance of continuous rather than discrete structures in many disciplines (Ingold 2007), a trajectory is intended to capture the idea of artists and performers trying to construct coherent and more or less continuous journeys—threads of experience—through an extended performance that are then negotiated with participants, each of whom who may follow their own path, and where these paths then meet and separate as part of a complex social tapestry. Specifically, it has been proposed that a mixed-reality performance can be described in terms of three fundamental kinds of trajectory: canonical trajectories, participant trajectories, and historic trajectories.

Artists create *canonical trajectories* that express one or more ideal journeys through a performance. In a sense, canonical trajectories capture the design of the underlying narrative that guides the performance, although this is broadened to include all aspects of the experience from ticketing and admissions, framing and engaging with interfaces (as discussed earlier), to the structure of the digital media, to the ending of the performance. Multiple canonical trajectories can be created for a given performance expressing the routes taken by different roles or the choices that any one participant may make.

Participants in the performance then follow *participant trajectories* that inscribe their actual journeys through the work. Each participant creates an individual participant trajectory that describes a specific experience, and a participant will tend to create a new participant trajectory each time that person revisits the work. Interactivity, in which participants make their own choices about how to act, drives their trajectories to diverge from underlying canonical trajectories, while the opposing force of orchestration as discussed earlier tends to steer participant trajectories back toward canonical trajectories. An unfolding performance therefore involves a continual and productive tension between interactive and orchestration, leading to continuously diverging and re-converging canonical and participant trajectories. In a similar way, the convergence and divergence of multiple participant trajectories expresses the social dynamics of a particular performance, reflecting moments at which different participants are brought together to share aspects of an experience, as well as important moments of contemplative isolation in which they are deliberately separated.

Historic trajectories provide the ability to record and replay a performance by constructing particular historic views of what took place. This involves selecting and recombining segments from among different participant trajectories that have been recorded by the underlying system. In the simplest case, this may involve replaying a given participant trajectory to recreate a particular individual's experience as it took place. However, it might also include mixing elements from multiple participants' trajectories to create new fictional views of history, including mixing elements of different recorded journeys by a single participant to create an idealized view of their experience (e.g., showing one's overall history in a game by just selecting the best attempt at each level). Finally, historic trajectories might then be reused as canonical trajectories in future experiences so that previous participants can act as guides for future participants.

Although the concept of trajectories is intended to capture a sense of a continuous journey through a performance, this ideal of continuity is in fact often threatened by various transitions and significant

moments in the structure of a performance that require careful design to maintain an overall sense of coherence. Key transitional moments in mixed-reality performances may include the following: *beginnings* (carefully designing how the performance is framed and how participants are admitted, briefed, and otherwise engaged); *endings* (how the performance ends, often including the use of various physical and digital mementos that are given or sent to participants after the event to provoke reflection and discussion); *role and interface transitions* (supporting people changing roles, for example, moving from bystanding to spectating or from spectating to performing, and also how they pick up, put down, or hand over interfaces, which can be challenging where invisible sensing technologies are involved as discussed earlier); *seams* (coping with the practical constraints of underlying digital technologies, especially the limited coverage and accuracy of wireless communications and invisible sensing systems, which may threaten the smooth running of the performance); *access to physical resources* (unlike their purely digital counterparts, physical resources such as props, physical interfaces, and also real spaces, cannot readily be replicated; as a result, designers must pay careful attention to how access to these is scheduled and managed so as to avoid potentially disruptive contention, for example, when several participants arrive at a key physical location at the same time and have to wait, which may detract from their engagement with their individual experience); and finally *episodes* (some performances, for example slow games that are delivered over long time periods using mobile phones, involve highly episodic modes of engagement in which frequently disengagement and subsequent reengagement need to be carefully managed).

This emerging theory of trajectories captures many of the concepts that we discussed in earlier sections and tries to wrap them into an overarching conceptual framework through which we might analyze and ultimately design new kinds of performances that make extensive use of digital technologies. In particular, trajectories express the idea that an overall performance can combine multiple interfaces and roles (e.g., performer, audience,

and bystander), with people moving between them at different times. The idea of transitions relates to picking up and putting down musical interfaces, which in turn relates to Bellotti's five questions discussed earlier. The distinction between canonical and participant trajectories reflects the degree of direction, orchestration, or alternatively improvisation that might be possible. Up to now, these concepts have been driven from studies of theatrical rather than musical performances, but this is not to say that there are not clear resonances with the latter. When we consider the potential application of trajectories to musical performance, we must be careful about how to describe various performance roles. In the following, we consider that a performance is first "designed" by some combination of composers, directors, set designers, stage managers, and others who plan how it is intended to unfold and who arrange the combination of digital and physical materials (sounds, scores, interfaces, auditoria, and so forth) that constitute its canonical trajectories. These shape the actions of performers, who may often be trained musicians giving a deliberate performance, but might also potentially be more everyday users of musical interfaces that are displayed in public settings, for example in sound installations in galleries and exploratoria. These people interact with the various technologies involved to create their own participant trajectories, often in the presence of watching spectators, who might comprise both audience and bystanders, as discussed previously. Musical performance is often recorded of course, which is where historic trajectories come in to play, ranging from a simple recording of a live performance to more complex overdubbing and mixing that can be described in terms of the synthesis of a historic trajectory from many recorded participant trajectories. These recordings can then be sampled and replayed in future live performances, which may in turn be recorded, and so forth (a popular trend in modern music).

In other words, the concepts of canonical, participant, and historic trajectories can be seen in musical as well as theatrical performances, and some of the discussions of how they may converge and diverge may therefore be of value when designing musical experiences—so too should be an analysis of the

various forms of transitions that need to be considered during a performance, including the challenges of handing over instruments and designing for the constraints of sensing interfaces that we discussed earlier, as well as the importance of framing. It is interesting to speculate whether other important transitions that are seen in theatrical experiences such as dealing with episodes of engagement and access to physical resources also have to be considered in designing new forms of extended musical performance.

Conclusion

There is an emerging body of work within the fields of Human Computer Interaction and Computer-Supported Cooperative Work on mixed-reality performance, i.e., the use of computers to support new forms of theatre that extend interaction out to city streets and other public settings. Studies of these works have informed a growing understanding of what it means to perform with a computer interface, either as part of a deliberately staged theatrical event or as part of the performance of everyday interactions. This article has attempted to distil some of the key concepts from this body of work and consider how they might potentially apply to the design of musical performances.

Beginning with how a musician interacts with an instrument, especially one that employs invisible sensing technologies, we have argued that a systematic comparison of expected, sensed, and desired actions can generate new design possibilities in areas such as allowing for expressive gesture around an instrument, building in opportunities for rest and repositioning, enabling unusual performance effects, supporting meta-control of an instrument's settings, and recognizing the potential difficulties of gracefully picking the instrument up and setting it down again.

Widening our perspective to consider an ensemble of musicians playing multiple—possibly networked—instruments, we have argued for the importance of recognizing the complex multi-channeled nature of interaction, requiring designers to consider different opportunities for

direct coordination but also for Dix's concept of feedthrough (1997). Rich and effective feedthrough is less well supported with the current generation of electronic instruments, and there may be potential in exploring richer forms of embodiment here.

Next, we introduced the audience, considering how interaction with an instrument might be designed with spectators in mind. By comparing the ways in which instruments might hide, reveal, or even amplify different combinations of manipulations of an interfaces with the resulting effects of these manipulations, we revealed four broad design strategies: expressive, secretive, magical, and suspenseful. Although the "expressive" strategy seems well suited to designing digital instruments (and can already be seen in many), it is worth considering whether the other strategies might also be relevant too.

Finally, we considered the embedding of all of these aspects of musical interaction into an overarching structure of performance. Here, we drew on an emerging theory of trajectories through mixed-reality performance to consider how the relationships between pre-composed, live, and recorded actions could be expressed through so-called canonical, participant, and historic trajectories, and also how such trajectories must negotiate various key transitional moments if an overall sense of coherence is to be maintained.

This article has focused on how the design of digital instruments might incorporate concepts from HCI. Although this is ideally a useful and thought-provoking exercise, it is of course also the case that HCI will have much to learn from the design of digital musical instruments. An important future step is to try to apply these ideas in practice, which will no doubt challenge these concepts and ultimately lead to their refinement and or extension in important ways.

Acknowledgments

I gratefully acknowledge the support of the Engineering and Physical Sciences Research Council (EPSRC) through their funding of the Challenge of Widespread Ubiquitous Computing project (grant

EP/F03038X/1) and of the Research Councils UK (RCUK) through their funding of the Horizon Digital Economy Hub (grant EP/G065802/1). I would like to thank Michael Gurevich for his insightful comments and suggestions for revising the initial draft.

References

- Bellotti, V., et al. 2002. "Making Sense of Sensing Systems: Five Questions for Designers and Researchers." *Proceedings of the 2002 SIGCHI Conference on Human Factors in Computing Systems*. New York: Association for Computing Machinery, pp. 415–422.
- Benford, S., et al. 1995. "User Embodiment in Collaborative Virtual Environments." *Proceedings of the 1995 SIGCHI Conference on Human Factors in Computing Systems*. New York: Association for Computing Machinery, pp. 741–750.
- Benford, S., et al. 2004. "Uncle Roy All Around You: Implicating the City in a Location-Based Performance." *Proceedings of Advances in Computer Entertainment*. New York: Association for Computing Machinery. Available on-line at performancestudies.pl/dydaktyka/files/ace2004.pdf.
- Benford, S., et al. 2005. "Expected, Sensed, and Desired: A Framework For Designing Sensing-Based Interaction." *ACM Transactions on Computer-Human Interaction* 12(1):3–30.
- Benford, S., et al. 2006. "The Frame of the Game: Blurring the Boundary between Fiction and Reality in Mobile Experiences." *Proceedings of the 2006 SIGCHI Conference on Human Factors in Computing Systems*. New York: Association for Computing Machinery, pp. 427–436.
- Benford, S., et al. 2009. "From Interaction to Trajectories: Designing Coherent Journeys Through User Experiences." *Proceedings of the 2009 SIGCHI Conference on Human Factors in Computing Systems*. New York: Association for Computing Machinery, pp. 709–718.
- Bevilacqua, F., et al. 2006. "The Augmented Violin Project: Research, Composition and Performance Report." *Proceedings of the 2006 International Conference on New Interfaces for Musical Expression*. Paris: IRCAM, pp. 402–406.
- Blaine, T., and S. Fels. 2003. "Collaborative Musical Interfaces for Novices." *Journal of New Music Research* 32(4):411–428.
- Bowers, J., and S. O. Hellström. 2000. "Simple Interfaces to Complex Sound in Improvised Music." *Proceedings of the 2000 SIGCHI Conference on Human Factors in*

- Computing Systems Extended Abstracts*. New York: Association for Computing Machinery, pp. 125–126.
- Crabtree, A., et al. 2004. "Orchestrating a Mixed Reality Game 'On the Ground.'" *Proceedings of the 2004 SIGCHI Conference on Human Factors in Computing Systems*. New York: Association for Computing Machinery, pp. 391–398.
- Dix, A. 1997. "Challenges For Cooperative Work on the Web: An Analytical Approach." *Computer Supported Cooperative Work* 6(2/3):135–156.
- Ingold, T. 2007. *Lines: A Brief History*. New York: Routledge.
- Ishii, H., and B. Ullmer. 1997. "Tangible Bits: Towards Seamless Interfaces Between People, Bits and Atoms." *Proceedings of the 1997 SIGCHI Conference on Human Factors in Computing Systems*. New York: Association for Computing Machinery, pp. 234–241.
- Johansen, R. 1988. *GroupWare: Computer Support for Business Teams*. New York: Free Press.
- Jordà, S., et al. 2007. "The reacTable: Exploring the Synergy Between Live Music Performance and Tabletop Tangible Interfaces." *Proceedings of the 1st International Conference on Tangible and Embedded Interaction*. New York: Association for Computing Machinery, pp. 139–146.
- Koleva, B., et al. 2001. "Orchestrating a Mixed Reality Performance." *Proceedings of the 2001 SIGCHI Conference on Human Factors in Computing Systems*. New York: Association for Computing Machinery, pp. 38–45.
- Law, E., et al. 2008. "Towards a Shared Definition of User Experience." *Proceedings of the 2008 SIGCHI Conference on Human Factors in Computing Systems Extended Abstracts*. New York: Association for Computing Machinery, pp. 2395–2398.
- Lewis, G. 2007. "The Virtual Discourses of Pamela Z." *Journal of the Society for American Music* 1(1):57–77.
- Marshall, J., S. Benford, and T. Pridmore. 2010. "Deception and Magic in Collaborative Interaction." *Proceedings of the 2010 SIGCHI Conference on Human Factors in Computing Systems*. New York: Association for Computing Machinery, pp. 567–576.
- Poepel, C., and D. Overholt. 2006. "Recent Developments in Violin-Related Digital Musical Instruments: Where Are We and Where Are We Going?" *Proceedings of the 2006 International Conference on New Interfaces for Musical Expression*. Paris: IRCAM, pp. 390–395.
- Reeves, S., et al. 2005. "Designing the Spectator Experience." *Proceedings of the 2005 SIGCHI Conference on Human Factors in Computing Systems*. New York: Association for Computing Machinery, pp. 741–750.
- Reeves, S., M. Fraser, and S. Benford. 2005. "Engaging Augmented Reality." *Proceedings of the 2005 SIGCHI Conference on Human Factors in Computing Systems Extended Abstracts*. New York: Association for Computing Machinery, pp. 2256–2265.
- Rosen, C. 2002. *Piano Notes: The Hidden World of the Pianist*. New York: Free Press.
- Stach, T., et al. 2007. "Improving Recognition and Characterization in Groupware With Rich Embodiments." *Proceedings of the 2007 SIGCHI Conference on Human Factors in Computing Systems*. New York: Association for Computing Machinery, pp. 11–20.
- Stewart, J., B. Bederson, and A. Druin. 1999. "Single Display Groupware: a Model for Co-Present Collaboration." *Proceedings of the 1999 SIGCHI Conference on Human Factors in Computing Systems*. New York: Association for Computing Machinery, pp. 286–293.
- Streitz, N., et al. 1999. "i-LAND: An Interactive Landscape for Creativity and Innovation." *Proceedings of the 1999 SIGCHI Conference on Human Factors in Computing Systems*. New York: Association for Computing Machinery, pp. 120–127.
- Sudnow, D. 1978. *Ways of the Hand: The Organization of Improvised Conduct*. Cambridge, Massachusetts: Harvard University Press.
- Wanderley, M., et al. 2005. "The Musical Significance of Clarinetists' Ancillary Gestures: An Exploration of the Field." *Journal of New Music Research* 34(1):97–113.
- Warming Pedersen, E., and K. Hornbaek. 2009. "mix-iTUI: A Tangible Sequencer for Electronic Live Performances." *Proceedings of the 3rd International Conference on Tangible and Embedded Interaction*. New York: Association for Computing Machinery, pp. 223–230.
- Wilson, S. 2002. *Information Arts: Intersections of Art, Science, and Technology*. Cambridge, Massachusetts: MIT Press.
- Zimmermann, R., et al. 2008. "Distributed Musical Performances: Architecture and Stream Management." *ACM Transactions on Multimedia Computing, Communications and Applications* 4(22):Article 14.