Reflets: Combining and Revealing Spaces for Musical Performances

Florent Berthaut
University of Bristol
florent@hitmuri.net

Diego Martinez Plasencia
University of Bristol
diego.martinez-plasencia@bristol.ac.uk

Martin Hachet
INRIA Bordeaux - Sud-Ouest
martin.hachet@inria.fr

Sriram Subramanian
University of Bristol
sriram@cs.bris.ac.uk

ABSTRACT

We present *Reflets*, a mixed-reality environment for musical performances that allows for freely displaying virtual content on stage, such as 3D virtual musical interfaces or visual augmentations of instruments and performers. It relies on spectators and performers *revealing* virtual objects by slicing through them with body parts or objects, and on planar slightly reflective transparent panels that *combine* the stage and audience spaces. In this paper, we describe the approach and implementation challenges of *Reflets*. We then demonstrate that it matches the requirements of musical performances. It allows for placing virtual content anywhere on large stages, even overlapping with physical elements and provides a consistent rendering of this content for large numbers of spectators. It also preserves non-verbal communication between the audience and the performers, and is inherently engaging for the spectators. We finally show that *Reflets* opens musical performance opportunities such as augmented interaction between musicians and novel techniques for 3D sound shapes manipulation.

Author Keywords

reflets, augmented performance, reflection, optical combiners, revealing, mixed-reality

ACM Classification

H.5.1 [Information Interfaces and Presentation] Multimedia Information Systems — Artificial, augmented, and virtual realities
H.5.5 [Information Interfaces and Presentation] Sound and Music Computing
J.5 [Computer Applications] Arts and Humanities — Performing arts (e.g., dance, music)

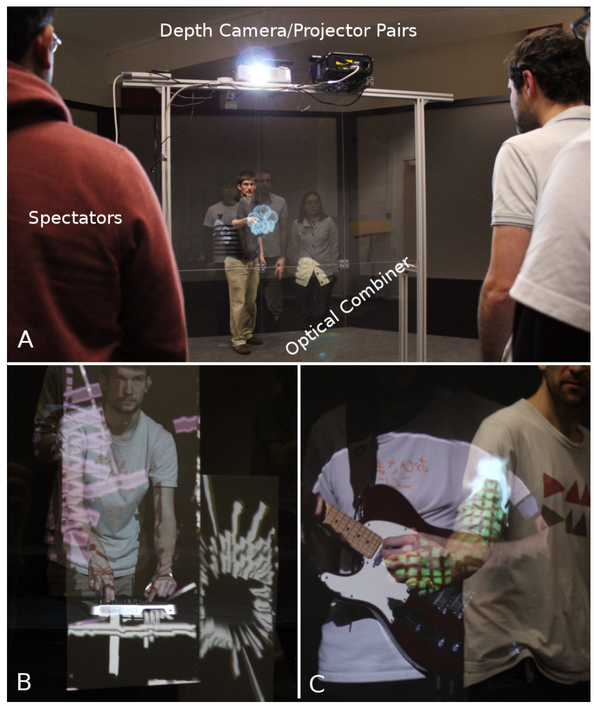

Figure 1: A) The blue sound shape manipulated by the performer is revealed by a spectator. B) Spectators reveal virtual components of a musical instrument using white boards. C) Two performers interact through the shared space.

1. INTRODUCTION

Musical performances are more and more enriched with 3D virtual content. This content may serve as an artistic visualisation, as seen in many recent DJ performance, but it can also be used for musical control, either as one component [9] of the performance, or as the instrument itself [2]. In order to display virtual content amongst physical scenery and performers on stage, Mixed-Reality (MR) display technologies are required. However, musical performances imply demanding constraints that are only partially fulfilled by existing MR technologies: 1) they involve a large number of users (both spectators and performers); 2) large display volumes (i.e. covering the whole stage); 3) they require visual communication between the spectators and performers; and sometimes 4) audience participation. While existing MR display technologies have been used to create a number of great musical performances, they do not match all these specific requirements, as we explain in more details in section 4. We believe that a MR display that would take all these constraints into account would both benefit existing performances and open many artistic opportunities.

In this paper, we introduce *Reflets*, a MR environment for musical performances. As depicted on Figure 1, a large set of slightly reflective transparent panels is placed between the performers and spectators, with a depth camera/projector pair covering each side. Spectators and performers use their bodies or other objects to *reveal* virtual content (3D shapes, images, textures), on both the stage and audience sides. The reflections on the panels make spectators, performers and virtual elements appear together in a shared space. We present our approach in Section 3. We discuss the specific requirements for mixed-reality displays in the context of musical performances in Section 4 and explain how *Reflets* meets those requirements. Although *Reflets* can be used as a generic MR display for musical performances with virtual content, we also demonstrate that it opens specific performance opportunities, through a series of workshops with performers, and four scenarios, as described in Section 5.

2. RELATED WORK

Our approach relates to a number of previous research in 3D display technologies.

First, we use the audience as a display. This idea has been explored in public events, for example by giving display elements to spectators, either static such as coloured boards or active such as small LED displays that can be controlled at a distance. An interesting work is the "Phone as pixel" [8] project, which displays 2D animations from a collection of spatially distributed smart phones. Our system extends these ideas from 2D displays to volumetric displays, and does not require devices that users have to hold during the performance.

In [5], authors track markers on pieces of paper and use them to reveal 3D objects. This is similar to the *revealing* technique that we describe in Section 3.2, but our ability to use it in either side of the combiner allows for novel scenarios. For instance, several slices can be revealed simultaneously, one from each side. The slices can refer to real objects on the other side, revealing their internals. Our method also handles larger slicing objects, e.g. to reveal large virtual elements, and a greater range of contents can be used, including both 3D meshes and volumetric textures.

We also strongly rely on the guidelines provided in [6]. The authors describe the possibility of creating a shared space, and quickly suggest the use of a revealing technique for augmenting objects in a museum cabinet. In this paper, we extend these basic ideas in the context of musical performance, ensuring that they fit the specific constraints of this domain. The mirrors and revealing techniques are brought to a much larger scale and the possibilities for musical visualisation, 3D interaction, collaboration and audience participation are explored.

3. REFLETS

3.1 General approach

We use the characteristics of musical performances, such as a large number of potentially active spectators and performers, and a separation between stage and audience spaces, to design a MR environment that matches their specific requirements, as explained in Section 4.

The principle of *Reflets* is to combine: 1) the technique of *revealing* virtual content through intersection with physical elements and users; 2) slightly reflective transparent panels that act as an optical combiner. The revealed content is perceived placed in a unique shared space, due to the reflections, overlapped with both the stage and audience side.

An example use of *Reflets* is depicted on Figure 1.A. A 3D sound shape manipulated by the performer is revealed by a spectator, i.e. projected on him. From the audience side, the reflection of this projection appears on the stage space overlapped with the performer's hand. From his side, the performer sees the reflection of his hand overlapped with the audience space, allowing him to grab and manipulate the shape. Users on both sides therefore perceive the same space with virtual and physical elements, alternatively seen directly and as reflections. Both performers and spectators act as a moving screen, with which it is possible to augment physical elements or display virtual objects anywhere on stage, by revealing them from either side of the mirror. *Reflets* can be used as a generic MR display for any existing performance that involves virtual content, such as Immersive Virtual Musical Instruments, 3D visualisations of music, augmented reality audience participation, and so on. In addition, we show in Section 5 that *Reflets* opens new opportunities for musical performance.

In the following sections, we explain the two components of *Reflets* in more details.

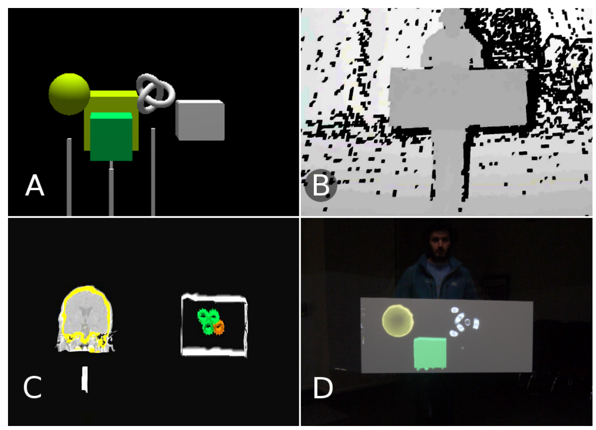

Figure 2: Virtual content at different positions can be revealed. (A) Two shapes at the back and three at the front. (B) A depth map of the physical space is captured, then per-pixel intersection with virtual objects is computed and rendered. (C) shows the intersection with the two shapes at the back. (D) shows the re-projection of the intersection with the three shapes at the front.

3.2 Revealing technique

The first component of *Reflets* is the *revealing* technique. It allows for displaying virtual objects placed at any position in the physical space by intersecting them with a physical element, e.g. object or body part. The intersection, or slice, of the object is then projected back on the physical element, resulting in part or all of the virtual object appearing in the physical space, revealed by the user. If the physical element used matches the surface of the virtual object, then this technique gives the same result as projection mapping. Otherwise a slice with both the inside and the edge of the object is displayed. The hardware side of this *revealing* technique consists of a Depth Camera (Asus Xtion) and Projector pair. They are calibrated so that the acquired depth image and the projection are both correctly aligned with the physical world. This means that the image of a physical object captured can be exactly reprojected on the object itself. These pairs are placed so that they capture the physical spaces occupied by the spectators and the performers. The software side is 3D scenegraph based on Ogre and OpenGL [1]. It allows for building scenes with primitives or mesh files, and controlling all the parameters of the scene with audio features and MIDI and OpenSoundControl messages. However, the main part of this application is the rendering method. It is composed of three steps, as depicted on Figure 2. We first capture the surfaces from the physical space, e.g. audience, performers and objects, using the depth cameras and render them as meshes to textures that will serve as depth maps. The second step involves slicing the virtual elements using these depth masks. For each virtual object that we want to display, we render it from the projector point of view, and test the intersection of the object per pixel with the depth masks. Sliced meshes can be rendered in different ways. The edges can be highlighted with a color or texture. The inside can be displayed as a solid color, a gradient, or even a volumetric texture (e.g. an MRI scan), as depicted on Figure 2.C. The third step of our method is the re-projection of the slices in the physical world on the intersecting objects, as depicted on Figure 2.D.

Figure 3: A) Shared space created by the optical combiner. B) and C) The three users can all be perceived from both sides. D) Virtual content (here a cog) can be revealed from both sides and appears in the shared space at the correct depth for all users.

3.3 Optical Combiner

The second component of *Reflets* is a set of optical combiners, more specifically slightly reflective transparent surfaces. A 2m*1.5m grid of flat 8mm acrylic panels is placed between the musicians and the spectators as depicted on Figure 1. Users on both sides can perceive their reflection on the panels, and these reflections appear overlapped with the space on the other side. From both sides, the two spaces delimited by the panels are therefore perceived as combined in a single shared space, as explained in [6] and shown on Figure 3. Because the panels are flat, the reflections do not depend on the position of the observer. Therefore, if a reflected object overlaps a physical object on the other side, these two objects will be perceived as overlapped wherever the observer stands in front of the panels. This ensures that all spectators and performers perceive the same shared space, with reflections and physical objects combined. The originality of *Reflets* is to add the *revealing* technique, by covering both sides with depth camera / projector pairs, as shown on Figure 1.A. The virtual elements revealed by users or physical elements on one side are reflected on the panels and therefore appear overlapped with the other side. A virtual object placed in the shared space can therefore be revealed from both sides. Both slices are seen together (due to the reflections) completing the rendering of the object (see Figure 3.D). As explained earlier, the revealed objects are perceived in the shared space at the same position by all performers and spectators. This technique can also be used asymmetrically, with a referent real object on one side and revealing happening on the other. Augmentation can then be revealed floating around the referent object or even inside it, as shown on Figure 1.B. The optical combiners in *Reflets* are only slightly reflective, so that by controlling the lighting carefully we can exactly decide which physical elements on each side will be perceived, both directly and in reflections, and therefore which elements will compose the shared space.

[1] https://launchpad.net/rouages

4. MIXED REALITY DISPLAYS FOR MUSICAL PERFORMANCES

In this section, we discuss the specific requirements for mixed-reality displays in the context of musical performances and explain how *Reflets* takes them into account. We also discuss the issue of engaging the audience with the system.

4.1 Consistency for a large number of users

The first requirement is the ability to handle a large number of users. Many mixed-reality displays focus on one or few users, to ensure that each perceives the virtual content consistently aligned with the physical space. However, musical performances typically involve a large number of users, both on the spectators and on the performers' sides. This requirement rules out technologies such as 3D stereoscopic projections, as used in [9]. The virtual elements, although perceived at different depths, would not appear at the same position relative to the physical world for users outside a sweet spot, except when placed on the surface of the screen.

Reflets can provide a consistent view of virtual objects for any number of users. The virtual objects are displayed in the physical space at their correct position, including depth, in the scene, because they are revealed by and projected on physical elements. Thanks to the properties of flat optical combiners, as described in section 3.3, the displayed virtual objects then appear in the shared space, and can be perceived consistently by users standing at any position on both sides of the panels.

4.2 Full volume display

The second constraint is the scale of the augmentation. Performances may cover large volumes due to performers' movements and stage scenery. In addition, within these volumes, virtual elements might need to be placed anywhere, even inside physical elements such as in [3] or overlapping performers. This issue heavily restrains which display technology can be used. For example volumetric displays are often of a limited size due to the complexity of the technology (moving mechanical parts, number and range of voxels, ...). Projection mapping and semi-transparent screens are often used to provide large scale rendering of virtual content. However, these technologies require that physical elements such as screens are placed at the same position on stage as the virtual elements. The virtual content then appears only on the projection surfaces. It prevents the visualisation of performers reaching through 3D sound shapes in the case of 3D instruments, or virtual content overlapping physical instruments. The Pepper's Ghost technique, with a 45 degrees angled half-silvered mirror projecting the reflection of a screen in mid-air, allows for overlapping projections and physical objects. However, the virtual content remains limited to the reflected surface, it can not cover the full volume of the performance.

On the other hand, with *Reflets* and thanks to the transparent reflective panels and revealing technique, virtual content can be revealed from the audience side anywhere on stage, even inside physical elements. The captured and projected volume on each side can easily be scaled by adding more Depth Camera / Projector pairs around it. The width of the screen is an important variable as one needs to ensure that all spectators and performers are able see the entire shared space through it. It can be extended simply by adding more panels between the performers and spectators. The virtual volume that can be displayed finally depends on the number of physical elements on both sides to reveal them. While this means that the more spectators, the bigger the displayed volume, a small number of spectators can also cover large areas with the use of physical props such as white movable panels, inflatable shapes and so on.

4.3 Visual communication

Visual communication between performers and spectators is an important part of the performance. Members of the audience need to be able to clearly perceive performers' gestures and intentions [4]. In the other direction, some performers such as DJs need to perceive the reaction of the audience to adapt their performance. This constraint implies that the display technologies shouldn't impair direct visual communication. Video augmented reality, with spectators using their smartphones or shared screens to perceive the virtual content superimposed over the stage, forces them to look at the screen instead of the physical performance. Visual contact between performers and the audience is then lost. Moreover, the performance looses visual quality, as it has to be digitized and rendered on a small screen.

In *Reflets*, although the audience and performers are completely separated by acrylic panels, these are mostly transparent, preserving the visual contact. In addition, the lighting of both spectators and performers can be finely controlled by placing virtual shapes with solid colours which act as spotlights when users reveal them. *Reflets* also preserves the physicality of the performance as spectators directly perceive the performers, not a digitized version.

4.4 Audience Participation

Audience participation is not a requirement of musical performances. However, engaging spectators process of displaying the virtual content can add to their experience and open possibilities such as, as described in [7], allowing them to control some parts of the music.

The participation of the audience is at the core of *Reflets*, since the system relies on them to reveal virtual objects of the performances. However, it can take some time for spectators to engage with it. We ran three public demonstration sessions with 10 participants each to study this issue. Each group of participants watched short performances based on *Reflets*. They were then asked to comment on the system individually on paper and collectively as a discussion. The conclusion of this study is that spectators can benefit from guides before or during the performance. Solutions are either to have all or some of the spectators go through a tutorial before the performance, or to build the performance with a progression on the complexity of the spectator's actions. Some performers can also play the role of participants showing the others how to interact with the system. This encourages members of the audience transitioning from the role of spectators to that of participants, as described in [1].

During the performance, it is also important to guide spectators in the revealing process. This can be done through static physical or virtual guides such as physical frames, signs on the floor, or virtual stands of objects visible because they are intersecting the floor. Another possibility is to use animated guides, such as pulsating spheres intersecting the spectators who can follow their movements to discover virtual content placed at their center. Our study also yielded a number of insights regarding the virtual content. The first one relates to depth of virtual objects. While revealing content horizontally and vertically is straightforward, spectators seemed to need more time to find their depth and boundaries on the axis perpendicular to the panels. To make the exploration easier, it is therefore better to exaggerate the depth of virtual content. Finally, large white objects, such as inflatable shapes or cardboard panels, can be given to the audience during the performance. As depicted on Figure 1.B, they can be used to reveal larger parts of the virtual content and placed on the floor in order to compose a virtual screen.

5. PERFORMANCE OPPORTUNITIES

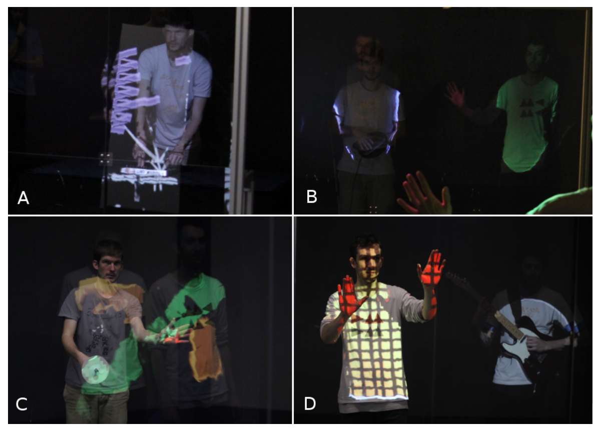

Figure 4: Performance scenarios as seen from the audience. A) Layers, B) Hidden Drums, C) Sound Paths, D) Grab and Morph

Although *Reflets* can be used as a generic Mixed-Reality display technology for augmented performances, we also believe that it opens novel musical performance opportunities. In order to verify this hypothesis, we ran a series of workshops with 8 performers of various backgrounds, including musicians, circus artists, visual artists and dancers. We organised both collective and individual sessions. We first presented the system and how it could be used to as a mixed-reality performance display. Then we asked the performers to design scenarios that would rely on the specificity of the system during both brainstorming sessions and hands-on sessions to implement test ideas. They proposed a number of scenarios and many insights on the opportunities opened by *Reflets*. We analysed these ideas and defined potential directions for the design of new performances. We then refined four of the proposed musical performance scenarios with the performers so that they illustrate these directions.

Performers discussed several scenarios in which *Reflets* could be used to allow the audience to explore virtual content surrounding or overlapping the performers. This content can be static scenery, e.g. in the case of theater plays, or active elements of the performance, e.g. augmentations of a dancer. We illustrate these augmentation and visualisation possibilities with the first scenario, *Layers*, depicted on Figure 4.A and Figure 1.B. Similarly to the Rouages project [3], the audience can reveal the various mechanisms of a Digital Musical Instrument, i.e. sensors extensions, sound processes and mappings, with their bodies or physical props. Contrary to Rouages, the augmentations can be placed inside the instrument, under the physical table, but also anywhere on stage around the musician and perceived with a correct alignment and perspective by all spectators.

Performers also discussed the use of the revealing technique as a way to control sound, by sending the position of the revealed slices within virtual objects via OpenSoundControl messages to external applications. Moreover, one of the performers highlighted a very important distinction. On one hand, 3D shapes can be used as discrete controls activated when revealed or as continuous ones, e.g. 3D sliders, even following complex 3D paths. Revealing volumetric textures can also be used as a metaphor for the exploration of soundscapes. In these cases, the focus of the revealing control is on the shapes themselves. But the focus can also be put on the movements inside large virtual volumes, the revealing technique then highlighting gestural control. The second scenario, *Hidden Drums*, depicted on Figure 4.B, illustrates the use of the revealing technique for discrete control. In this improvised performance, the musician plays virtual drums by revealing them and vibrating them with a handheld electronic drum pad. These vibrations are mapped to parameters of granular synthesis and a variable delay effect, allowing for both textures and rhythms to be played. The third scenario, *SoundPaths*, illustrates the use of the revealing technique as a continuous musical control at the shape level. As shown on Figure 4.C, the performer uses a WiiRemote and 3D position tracking to draw 3D sound paths that appear on stage, revealed by the spectators. While drawing, the sound sample is played so that there is a correspondence between the position in the sample and the position in the path. When the path is then revealed by the performer at one position, the sound is played as a loop from that position. The vertical and horizontal positions of the path respectively affect the frequency of a band-pass filter applied to the sound and the length of the loop. The last scenario, *Grab and Morph* illustrates the revealing technique highlighting gestural control. As depicted on Figure 4.D, it involves two performers, a guitar player on one side of the panels and a gestural performer on the other side. The gestural performer evolves in a large virtual box within which his movements, e.g. position and distance between hands, are mapped to various parameters of effects applied to the guitar sound.

Finally, many of the proposed scenarios involved the use of the optical combiner as a way to enable interaction between performers, or between spectators and performers. Because users on both sides appear in the shared space, they can touch and reach through each other, opening numerous performance opportunities. In the *Hidden Drums* scenario, in order to be active, the sound shapes played by the musician need to be revealed on both sides by the musician and the spectators. From their side, the audience is therefore able to influence the performance by interactively "building" the instrument that the performer plays. In the *Grab and Morph* scenario, the guitar player and gestural performer appear in the shared space as if they were side by side. As shown on Figure 1.C, the gestural performer can reach inside the guitar using his reflection to grab a short sample of the sound, and replay it through controlled effects. A musical dialog can therefore emerge from their interactions through the shared space.

6. CONCLUSION

In this paper, we presented *Reflets*, a mixed-reality environment for musical performances. It matches their specific requirements such as high number of users, large display volume and preserved visual communication. It is also inherently engaging for the audience. Finally, thanks to its unique components and approach, it opens new performance opportunities. We see two main future research directions for *Reflets*. First, the system can be extended to support more than two combined spaces, either with multiple pairs of spaces that will enable modular performances, or with multiple audience spaces around the stage space. Second, by using a more complex mirror arrangement, new opportunities for musical collaboration within digital orchestras could be created, enabling interactions between all musicians. Additional interaction techniques and devices will be needed in order to allow for fine manipulation of others' instruments or sounds at a distance.

7. ACKNOWLEDGMENTS

The authors would like to thank all the performers who participated in the workshops, and especially Alex Jones, Nick Jannaway and Sebastien Valade for their help. This project was funded through the Marie Curie FP7 framework (Grant Agreement PIEF-GA-2012-330770).

8. REFERENCES

[1] S. Benford, G. Giannachi, B. Koleva, and T. Rodden. From interaction to trajectories: designing coherent journeys through user experiences. In *Proceedings of CHI*, pages 709–718, 2009.

[2] F. Berthaut, M. Desainte-Catherine, and M. Hachet. Drile : an immersive environment for hierarchical live-looping. In *Proceedings of NIME*, pages 192–197, 2010.

[3] F. Berthaut, M. Marshall, S. Subramanian, and M. Hachet. Rouages: Revealing the Mechanisms of Digital Musical Instruments to the Audience. In *Proceedings of NIME*, 2013.

[4] B. Bongers. Exploring novel ways of interaction in musical performance. In *C&C '99: Proceedings of the 3rd conference on Creativity & cognition*, pages 76–81, New York, NY, USA, 1999. ACM.

[5] A. Cassinelli and M. Ishikawa. Volume slicing display. In *SIGGRAPH ASIA Art Gallery & Emerging Technologies*, page 88, 2009.

[6] D. Martinez Plasencia, F. Berthaut, A. Karnik, and S. Subramanian. Through the combining glass. In *Proceedings of ACM UIST*, pages 341–350. ACM, 2014.

[7] D. Mazzanti, V. Zappi, D. Caldwell, and A. Brogni. Augmented stage for participatory performances. In *Proceedings of NIME*, pages 29–34, 2014.

[8] J. Schwarz, D. Klionsky, C. Harrison, P. Dietz, and A. Wilson. Phone as a pixel: Enabling ad-hoc, large-scale displays using mobile devices. In *Proceedings of CHI*, CHI '12, pages 2235–2238, 2012.

[9] V. Zappi, D. Mazzanti, A. Brogni, and D. Caldwell. Design and evaluation of a hybrid reality performance. In *Proceedings of NIME 2011*, pages 355–360.