Active listening to a virtual orchestra through an expressive gestural interface: The Orchestra Explorer

Antonio Camurri
InfoMus Lab – Casa Paganini
University of Genova
Piazza Santa Maria in Passione, 34
16123 Genova, Italy
+39 010 2758252
antonio.camurri@unige.it

Corrado Canepa
InfoMus Lab – Casa Paganini
University of Genova
Piazza Santa Maria in Passione, 34
16123 Genova, Italy
+39 010 2758252
corrado@infomus.org

Gualtiero Volpe
InfoMus Lab – Casa Paganini
University of Genova
Piazza Santa Maria in Passione, 34
16123 Genova, Italy
+39 010 2758252
gualtiero.volpe@unige.it

ABSTRACT
In this paper, we present a new system, the Orchestra Explorer, enabling a novel paradigm for active fruition of sound and music content. The Orchestra Explorer allows users to physically navigate inside a virtual orchestra, to actively explore the music piece the orchestra is playing, to modify and mold the sound and music content in real-time through their expressive full-body movement and gesture. An implementation of the Orchestra Explorer was developed and presented in the framework of the science exhibition "Cimenti di Invenzione e Armonia", held at Casa Paganini, Genova, from October 2006 to January 2007.

Keywords
Active listening of music, expressive interfaces, full-body motion analysis and expressive gesture processing, multimodal interactive systems for music and performing arts applications.

1. INTRODUCTION
Nowadays, listening to sound and music is usually still a passive, non-interactive experience. Quoting John Sloboda "In highly industrialized societies, we listen to more music, but we make less" [1]. Even modern devices do not allow for interactive user participation. This can be considered a degradation of traditional listening experience, in which the public can interact in many ways with performers to modify the expressive features of a piece.

In this paper, we introduce a new system, the Orchestra Explorer, enabling a novel paradigm for active experience of sound and music content. With *active experience* and *active listening* we mean that listeners are enabled to interactively operate on music content, by modifying and molding it in real-time while listening. Active listening is the basic concept for a novel generation of interactive music systems [2], which are particularly addressed to a general public of beginners, naïve and inexperienced users, rather than to professional musicians and composers.

The Orchestra Explorer allows users to physically navigate inside a virtual orchestra, to actively explore the music piece the orchestra is playing, to modify and mold in real-time the music performance through expressive full-body movement and gesture. Concretely, the virtual orchestra is spread over a physical surface. By walking and moving on the surface, the user discovers each single instrument and can operate through her expressive gestures on the music piece the instrument is playing.

The Orchestra Explorer is not a simple reproduction or remixing of multiple audio tracks, nor a full automatic conducting system. It is something in between these two extremes. On the one hand, it provides the active listener with means for operating on the sound and music content that are not available in simple (passive) reproduction or remixing of multiple audio tracks. On the other hand, it does not provide the full control of the performance as in traditional conducting systems (e.g., see [3]). This approach is motivated by our aim of developing a new paradigm which, while actively engaging listeners, is at the same time different from traditional metaphors such as conducting, and enabled by recent research on expressive multimodal interfaces. The user really becomes an explorer of sound and music, i.e., she discovers the content step by step; she gradually understands how music performance works; she learns how to operate on the content.

Section 2 introduces the overall system architecture; Section 3 describes the model we designed for interacting with space; Section 4 and 5 illustrate the implementation of an Orchestra Explorer installation developed for the science exhibit "Cimenti di Invenzione e Armonia". The conclusions summarize some issues that emerged from such installation.

2. SYSTEM ARCHITECTURE
Figure 1 shows the overall architecture of the Orchestra Explorer. Since in this paper we wish to put into evidence how localization and expressive gesture features operate in controlling real-time processing of the input audio stream, we distinguished two major components of the system: the audio component in the lower part and the control component in the upper part of the figure. The system also operates on a video stream and provides a video output, but this is secondary with respect to the audio output.

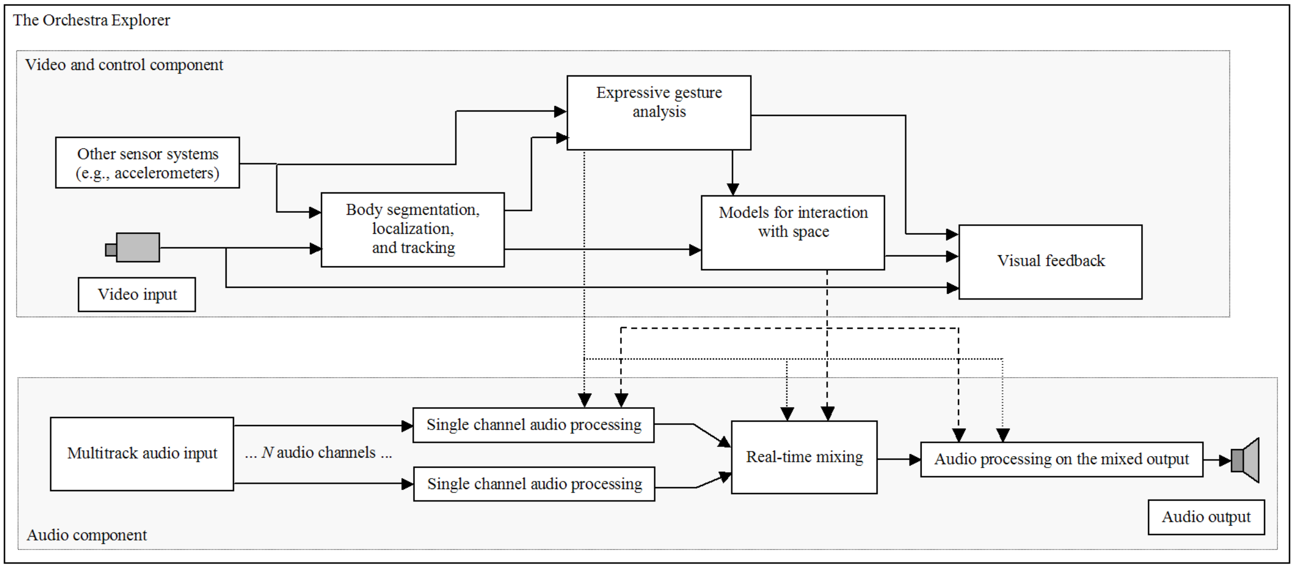

Figure 1. The overall system architecture of the Orchestra Explorer

The input of the audio component is multi-channel audio. In a typical Orchestra Explorer configuration two audio channels (stereo recordings) are associated to each instrument of the virtual orchestra. Recordings of whole sections of an orchestra can also be used. Initially, each audio channel is processed separately, keeping however synchronization among tracks. One or more sound processing modules operate on every single audio channel. The sound processing modules receive values of parameters from the control component. The processed sound tracks are then mixed in real-time. The mixing process is also controlled by the control component. Finally, further effects and sound processing techniques can be applied to the mixed output, again controlled in real-time by the control component. Typical output is stereo or quadrifonic audio, reproduced in the space the user is navigating.

Inputs of the control component are images from a videocamera and possibly data from sensors (e.g., accelerometers). Images are processed by means of computer vision techniques in order to segment the body of the user from the background and to track in real-time her (x, y) position on the surface she is walking on (body segmentation, localization, and tracking module). The module for interaction with space (models for interaction with space) receives as input such (x, y) position, whereas the expressive gesture analysis module receives the user position, processed images (e.g., the body silhouette), and possible filtered data from sensors.

The expressive gesture analysis module extracts expressive features from the incoming images and data. Expressive features provide a high-level qualitative description of the gestures a user is performing. For example, Quantity of Motion (QoM) measures the amount of detected motion (related to the energy of a gesture), Contraction Index (CI) measures how much the body is contracted or expanded, Directness Index (DI) measures how much a gesture trajectory is straight, Stability Index (SI) measures how much a posture is stable. A detailed description of the expressive features we extracted and used in several applications can be found in [4][5]. Starting from such expressive features, higher-level aspects of gesture can be investigated. For example, in a recent study [6] we used a set of expressive features of gesture trajectories for classifying expressive gesture according to the Time and Space dimensions of Rudolf Laban's Theory of Effort [7][8] (i.e., as Direct or Flexible, Quick or Sustained). In the same way, the movement of an active listener in the Orchestra Explorer space can be characterized as Direct or Flexible, Quick or Sustained. Starting from that, further high-level characterization of user's behavior is possible. For example, a flexible, non-fluid, sustained movement can indicate hesitation, whereas direct, quick, and fluid movement can indicate decision and determination.

The expressive gesture analysis module is organized according to the multi-layered framework for expressive gesture processing [9] we developed during the EU-IST Project MEGA (Multisensory Expressive Gesture Applications) and further refined in the EU-IST Tai-Chi Project (Tangible Acoustic Interfaces for Computer – Human Interaction).

The module that models interaction with space (discussed in details in Section 3) computes and superimposes onto the physical space a collection of 2D potential functions. At least one potential function is associated to each input audio channel for controlling real-time mixing. However, further potential functions can be created for controlling effects and sound processing modules. While the listener navigates the physical space, the values of such potential functions in correspondence of her (x, y) position are computed. The current values of the potentials associated with real-time mixing are used as weights for the sound level of the corresponding channel. The current values of the potentials associated with parameters of effects or other sound processing modules operate on the corresponding modules. The expressive gesture analysis module can dynamically modify the current value of the parameters of the potential functions, so that their profile can be molded by the expressive gestures the user performs while exploring the virtual orchestra.

The control component can therefore operate on the audio component in two different ways (see the different kinds of arrows in Figure 1):

(i) Through a direct mapping of expressive gesture features onto sound processing parameters (dotted arrows in Figure 1). For example, a higher value of CI can be associated to a larger bandwidth of a band-pass filter, a higher value of QoM to a stronger reverb component in the sound.

(ii) Through the mapping of the current values of the potential functions onto sound processing parameters (dashed arrows in Figure 1). E.g., a specific effect can be activated in a certain area of the space only, or hesitating movement can cause the instruments to be more spread over the orchestra space.

3. MODELS FOR INTERACTION WITH SPACE

This Section describes in more details the 2D potential functions we developed for modeling interaction with space. Let us consider for example the 2D potential functions used for controlling real-time mixing (remember however that a similar potential function can also be used for real-time control of audio effects). Typically, a potential function is associated to each instrument of the virtual orchestra. The potential function associated to the k-th instrument $P_k(x, y)$ is given by the weighted sum of two components, an exponential one $E_k(x, y)$ and a logarithmic one $L_k(x, y)$:

$$P_k(x,y) = w_{1k}E_k(x,y) + w_{2k}L_k(x,y) \qquad k = 1 \ldots N$$

where N is the number of instruments.

3.1 Exponential component

The exponential component is computed as a non-normalized bi-dimensional Gaussian function, that is:

$$E_k(x,y) = E_{Mk}\exp\left\{-\frac{1}{2(1-\rho_k^2)}\left[\left(\frac{x-\mu_{xk}}{\sigma_{xk}}\right)^2 - 2\rho_k\frac{x-\mu_{xk}}{\sigma_{xk}}\frac{y-\mu_{yk}}{\sigma_{yk}} + \left(\frac{y-\mu_{yk}}{\sigma_{yk}}\right)^2\right]\right\}$$

for $k = 1\ldots N$.

The Gaussian is placed on the 2D surface through the μ_{xk} and μ_{yk} parameters. The point with coordinates (μ_{xk}, μ_{yk}) corresponds to the center of the Gaussian, i.e., to its peak value. Note that such parameters can be changed in real-time, i.e., the instruments can be potentially moved around while the user is interacting with the virtual orchestra. For example, it is possible to make an instrument move according to user's gestures (e.g., by indicating toward the right, the user could move toward the right the instrument she is listening to).

Parameters σ_{xk} and σ_{yk} control how much the Gaussian is spread over the space along the x and y directions. Small and similar values for σ_{xk} and σ_{yk} produce a Gaussian having a peak situated in a very localized area of the space, that is the single instruments are very separated and do not overlap too much (a kind of spatial staccato). Large and similar values make instead the single instruments to be very spread over the space and overlapping each other (a kind of spatial legato). Very different values of σ_{xk} and σ_{yk} cause the Gaussian to be elongated along the x or the y direction.

For example, Figure 2 shows the projections on the (x, y) plane of two Gaussian potentials localized along a diagonal of the surface (dark red color corresponding to high values of the potential function). In case (a) σ_{xk} and σ_{yk} are both set to 0.5, in case (b) they are both set to 0.1, in case (c) they are both set to 0.25. As for music listening, in case (a) the user will listen to the two instruments separately and in the space in between the two instruments no sound will be produced. In case (b) the two instruments are highly overlapped, i.e., there is a wide area in between the two peaks where the user will listen to both instruments. In case (c), there is only a small area where the instruments are overlapping, that is, while moving from one instrument to the other the user will listen to the first one, then to both of them for a short time, and finally to the second one.

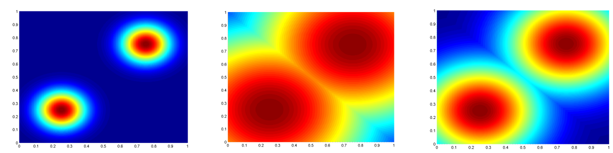

(a) (b) (c)

Figure 2. Effect of different values of parameters σ_{xk} and σ_{yk} on 2D Gaussian potential functions

σ_{xk} and σ_{yk} can also be modified in real-time according to user's gestures. For example, the way of approaching an instrument can change the way the instrument is spread over the space (e.g., a direct and decided gesture may be associated to a narrow area, whereas a flexible, hesitant movement may produce a wider one).

Parameter ρ_k controls the orientation of the 2D Gaussian. By jointly setting σ_{xk}, σ_{yk}, and ρ_k we can control the elongation of the Gaussian along a given direction. Figure 3 shows for example a Gaussian, localized in the centre of the space, obtained with σ_{xk} = 0.7, σ_{yk} = 0.2, and ρ_k = -0.9. In such a way, it is possible to stretch an instrument along space, i.e., by walking along the stretching direction, a user will listen to increasing levels for the instrument, until she will reach the peak, and then decreasing levels after the peak. Stretching can be the result of an expressive gesture of the user. For example, a repetitive run forth and back along a given direction can cause an instrument to be stretched along that direction.

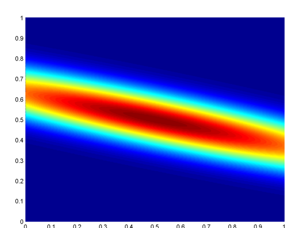

Figure 3. Joint effect of σ_{xk}, σ_{yk}, and ρ_k on the elongation and orientation of 2D Gaussian potential functions

Finally, E_{Mk} controls the amplitude of the Gaussian, which is related to the level of the associated instrument. Usually, E_{Mk} is set so that the maximum value of the overall potential function is 1. However, it could also dynamically change in real-time, i.e., some instruments could gradually fade out, while others fade in.

3.2 Logarithmic component

The logarithmic component is computed as follows:

$$L_k(x,y) = \begin{cases} \dfrac{L_{Mk}}{\log(1+a_k d_{Mk})}\log\left[1+a_k\left(d_{Mk}-d_k(x,y)\right)\right] & \text{if } d_k(x,y) \le d_{Mk} \\ 0 & \text{otherwise} \end{cases}$$

for $k = 1\ldots N$.

$d_k(x, y)$ is the (Euclidean) distance of the current position of the user from the point (μ_{xk}, μ_{yk}) where the logarithmic component is centered. That is:

$$d_k(x,y) = \sqrt{(x-\mu_{xk})^2 + (y-\mu_{yk})^2} \quad k = 1 \ldots N$$

The point (μ_{xk}, μ_{yk}), where the logarithmic component is localized, is typically the same point where the Gaussian component is also localized.

d_{Mk} is the maximum allowed distance from the point (μ_{xk}, μ_{yk}). For distances larger than d_{Mk} the logarithmic component is set to 0. By default d_{Mk} is set to the maximum distance that a user can reach while interacting with the virtual orchestra. That is, if S is the set of points (x, y) that can be reached by the user during interaction, d_{Mk} is computed as:

$$d_{Mk} = \max_{(x,y)\in S} d_k(x,y)$$

In the very usual case of a rectangular surface, d_{Mk} can be easily computed as the maximum distance of the central point (μ_{xk}, μ_{yk}) form the four corners of the rectangle.

However, the default choice for d_{Mk} often needs to be modified because of the overall musical result of the Orchestra Explorer. Taking d_{Mk} as the maximum reachable distance would cause an instrument to be more or less perceivable in almost the whole space. This is usually not the purpose of the installation. Smaller values for d_{Mk} are thus often selected in a way, which is analogous to what we already discussed for σ_{kx} and σ_{ky}, i.e., small values for d_{Mk} generate localized and separated instruments, while high values produce spread and overlapping instruments.

a_k controls the slope of the logarithmic component. The slope is related to how much gradual the fade in and the fade out of an instrument is while a user enters the area occupied by the 2D potential function. A small value for a_k produces a smooth slope, i.e., a more gradual fade in/fade out of the instrument. The higher is the value of a_k, the more sudden is the attack of the instrument. Similarly to the other parameters, a_k can also be changed in real time, so that the suddenness of the attack of an instrument can be related to the user's expressive motion and gesture. For example, repeated impulsive movements can cause an instrument to become more sudden in its attack[1].

Finally, L_{Mk} controls the amplitude of the logarithmic component, which is related to the level of the associated instrument. As for the Gaussian component, L_{Mk} is usually set so that the maximum value of the overall potential function is 1.

Figure 4a shows the overall potential function $P_k(x, y)$. The profile of $P_k(x, y)$, obtained for $y = 0.5$ is shown in Figure 4b (where the continuous line corresponds to the overall potential function, the dashed line to the Gaussian component, and the dotted line to the logarithmic component). The overall potential function is given by the weighted sum of the Gaussian and the logarithmic components. The overall potential function depicted in Figure 4 is obtained by setting the two weights to 0.5. The two components are both centered in the middle of the space.

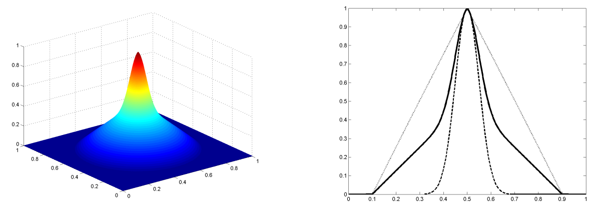

(a) (b)

Figure 4. The overall potential function $P_k(x, y)$ (a) and its profile for y = 0.5 (b)

The logarithmic component provides a smooth fade in/out while the user enters or exits the area covered by the potential function. The Gaussian component produces a strong peak in proximity of the central point of the potential. In such a way the user will gradually listen to an increasing level of the instrument while entering the area and moving toward the center. She will then listen to it playing loudly, once she reaches the center of the area. This mechanism provides a stronger perceptual awareness of what is happening.

4. THE INSTALLATION AT THE SCIENCE EXHIBITION "CIMENTI DI INVENZIONE E ARMONIA"

We implemented several instances of the Orchestra Explorer. The main one was developed for the science exhibition "Cimenti di Invenzione e Armonia", held at Casa Paganini, Genova, Italy, from October 2006 to January 2007. The exhibition was part of "Festival della Scienza", a huge international science festival held in Genova every year.

The Orchestra Explorer was installed on the stage of the 250-seats auditorium at Casa Paganini, an international center of excellence for research on sound, music, and new media, where InfoMus Lab has its main site. The installation covered a surface of about 9 m × 3.5 m. A single vidocamera observed the whole surface from the top, about 7 m high, and at a distance of about 10 m from the stage (we did not use sensors). Four loudspeakers were placed at the four corners of the stage for audio output. A white screen covered the back of the stage for the whole 9 m width. A videoprojector projected on such screen the video feedback. Lights were set in order to enhance the feeling of immersion for the users and to have a homogenous lighting of the stage.

The music piece "Borderline", by M. Canepa, L. Cresta, and A. Sacco, was selected for the installation. "Borderline" is an original piece of film music, which was never performed in public. It does not have a strong characterization and it is well suited as sound track for an exploration task. Since the users do not know it and it does not easily recall other pieces of music, "Borderline" helps the users in paying attention to what they are listening to. "Borderline" consists of 13 stereo audio tracks and includes the following music instruments: harp, cello, horn, flute, double bass, oboe, bassoon, percussions, piano, violins, and alto pizzicato. The instruments were placed over the stage and were associated to 13 Gaussian/logarithmic potential functions.

[1] Note that here we are referring to a kind of "spatial" attack, i.e., the suddenness of the fade in/out while entering in the area of an instrument, which is of course different from the usual "time" attack of sound.

A music stand was placed at the center of the stage where the conductor usually takes place. On the music stand a handmade map of the 13 instruments was placed. This helps the user in finding her way while exploring the virtual orchestra.

When the user steps on the stage, nothing happens until she goes in front of the music stand. Only at that point, the user can start interacting with the orchestra. Contraction Index (CI) is computed separately for the right and left side of the user's body. Expanded movement on a side makes the instruments on that side audible. For example, if the user stretches her right arm she can listen to the instruments on her right. At this stage we control separately all the 26 mono audio channels. The 26 audio channels were pre-processed so that by just summing them, the final mix can be obtained without further operations. CI on the right proportionally controls the level of the right channel of the instruments on the right, while CI on the left proportionally controls the level of the left channel of the instruments on the left. When the right part of the body is fully expanded the user listen to the mono version of the instruments on the right coming from the loudspeakers on the right. Only when the whole body is completely expanded, the user listens to the stereo reproduction of the whole orchestra. This mechanism together with the map the user has in front of her makes the user aware of the instruments she can listen to and of the position where the instruments are approximately located.

The user can now start moving around. While she moves, real-time mixing is performed and the level of each instrument in the mixed output is controlled by the associated potential function. Moreover, expressive movement features are extracted and used for direct control of parameters of sound processing blocks. For example, Quantity of Motion and Contraction Index can control reverberation, so that contracted and slow movements make the sound drier.

Visual feedback is also related to expressive features. Quantity of Motion controls the length and thickness of bright spikes in the output image and the length of a kind of shadow that follows the user's silhouette. However, visual feedback is only a secondary aspect of this installation and it was included mainly for creating a more pleasant environment for the user.

Figure 5 shows the running installation. Figure 5a shows a visitor stretching his arms in the position of the conductor. Figure 5b shows a visitor exploring the orchestra. This shot is taken from a position near to the location of the videocamera, therefore the point of view is similar to the point of view of the videocamera.

Over 2500 visitors visited "Cimenti di Invenzione e Armonia". Feedback was in general very positive. For example, visitors often interacted with the system for a time longer than the duration of the music piece. A brief training session, including explanations and demonstration by the staff, usually helps users in experiencing the installation. In order to have a reliable assessment, we are currently carrying out formal evaluation of usability and pleasantness with a sample of subjects. We are also considering analyzing trajectories of subjects e.g., for individuating possible frequently recurrent patterns in the exploration of the space.

The system demonstrated to be highly flexible. It proved to be robust for people of different sizes (e.g., children and adults) and dressed with different clothes. Some repositioning of the potential functions and retuning is needed in case of porting to another space of considerably different size.

5. IMPLEMENTATION: THE EYESWEB OPEN PLATFORM AND THE EYESWEB EXPRESSIVE GESTURE PROCESSING LIBRARY

The instance of Orchestra Explorer we developed for the exhibit "Cimenti di Invenzione e Armonia" was implemented using a new version of our EyesWeb open platform [9][10]: EyesWeb XMI (for eXtended Multimodal Interaction). The EyesWeb open platform and related libraries are available for free on the EyesWeb website www.eyesweb.org.

With respect to its predecessors, EyesWeb XMI strongly enhances support to analysis and processing of synchronized streams at different sampling rates (e.g., audio, video, data from sensors). We exploited such support for the synchronized processing and reproduction of the 13 audio tracks in "Borderline". The whole Orchestra Explorer installation was implemented as a single EyesWeb application (patch), managing audio processing, video processing, extraction of expressive features from movement and gestures, real-time audio mixing and control of audio effects, and generation of visual feedback. Every single component of the application was implemented as an EyesWeb sub-patch. The whole application ran on a single workstation (Dell Precision 380, equipped with two CPUs Pentium 4 3.20 GHz, 1 GB RAM, Windows XP Professional).

Potential functions were implement as EyesWeb modules (blocks) in a new version of the EyesWeb Expressive Gesture Processing Library. This library also includes modules for extraction and analysis of expressive features from human full-body movement and gesture.

(a)

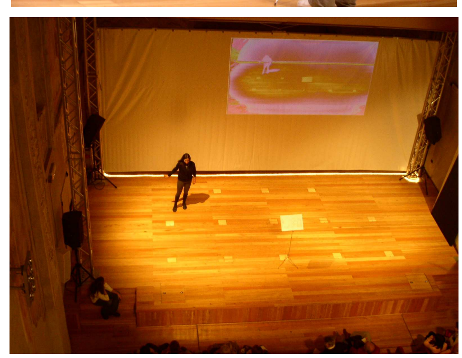

(b)

Figure 5. The Orchestra Explorer at the science exhibit "Cimenti di Invenzione e Armonia", Casa Paganini, Genova, Italy, October 2006 – January 2007.

6. CONCLUSIONS

From our experience with the Orchestra Explorer, especially at the science exhibit "Cimenti di Invenzione e Armonia", several issues emerged that need to be taken into account in future work.

The Orchestra Explorer deeply changes the way music is experienced. The exploration of a music piece along the space dimension causes a radical change in the way the same music piece is experienced along the time dimension. E.g., whereas the duration of "Borderline" is about 3' 30'', users usually spent much more time in interacting with the Orchestra Explorer. In another instance of the Orchestra Explorer we noticed that an excerpt of 30'' duration, was fruited for more than 10' without becoming annoying.

The Orchestra Explorer may have a strong impact in applications to music education. It provides a unique way to go deeply inside a complex music piece, by allowing the user to experience at the same time the whole piece and the single instrument (that can usually be accessed only by professionals). The inexperienced user can learn, for example, how to distinguish the timber of single instruments and how the single instruments contribute to the whole music piece. He can also learn that music instruments do not play all at the same time (in the Orchestra Explorer when an instrument does not play, entering its area does not produce any sound).

The Orchestra Explorer involves technical and artistic aspects. In fact, positioning the potential functions on the surface, setting their parameters, choosing the mapping between gesture features and parameters of the potential functions encompasses artistic choices that concur in the overall musical result of the Orchestra Explorer. Designing an instance of the Orchestra Explorer is somewhat similar to creating a completely new music piece. The original music piece is disassembled and recombined in another way, along another dimension (space instead of time). In this perspective, the Orchestra Explored can be considered as a kind of new composition tool or meta-instrument that allows beginners to access to music composition. As such, it shares similar problems with other existing Digital Music Instruments and sound synthesis techniques. For example, the Orchestra Explorer can be controlled through a huge number of low-level parameters. Each of the potential functions we discussed has 9 parameters that can be modified in real-time. This means that the Orchestra Explorer at "Cimenti di Invenzione e Armonia" had $13 \times 9 = 117$ parameters that could potentially be controlled in real-time. Finding suitable and meaningful (e.g., from a perceptual point of view) high-level metaphors for managing such a complexity is a topic which is still in the research agenda of the field of gestural control of digital musical instruments (see for example [11][12]).

Finally, the Orchestra Explorer provides an example of the concept of active listening of music that could characterize future devices and paradigms for music experience. We are currently developing an Orchestra Explorer that can be experienced through a hand held mobile device. Besides the obviously needed changes in the interaction metaphors, such a mobile Orchestra Explorer would allow active listening of music at home or in other personal environments without the need of large surfaces or complex video tracking systems.

7. ACKNOWLEDGMENTS

We thank composer Marco Canepa for providing the music piece "Borderline", for his precious contribution in preparing the audio material, and for the useful discussions about the design and the development of the Orchestra Explorer. We also thank our colleagues at DIST – InfoMus Lab for their concrete support to this work. Finally, we thank Festival della Scienza and the visitors of the science exhibition "Cimenti di Invenzione e Armonia" whose often enthusiast feedback strongly encouraged us in going on with this research.

8. REFERENCES

[1] Else, L. The Power of Music - Show Me Emotion. *New Scientist magazine*, 29 November 2003, 40-42.

[2] Rowe, R. *Interactive music systems: Machine listening and composition*. MIT Press, Cambridge MA, 1993.

[3] Lee, E., Karrer, T., and Borchers, J. Toward a Framework for Interactive Systems to Conduct Digital Audio and Video Streams. *Computer Music Journal*, *30,1* (Spring 2006), 21-36.

[4] Camurri A., Mazzarino B., and Volpe G. Expressive interfaces. *Cognition, Technology & Work*, *6,1* (Feb. 2004), 15-22.

[5] Camurri A., Lagerlöf I., and Volpe G. Recognizing Emotion from Dance Movement: Comparison of Spectator Recognition and Automated Techniques. *International Journal of Human-Computer Studies*, *59,(1-2)* (July 2003), 213-225.

[6] Camurri A.. and Volpe G. Multimodal and cross-modal analysis of expressive gesture in tangible acoustic interfaces. In *Proceedings of the 15th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN2006)* (University of Hertfordshire, Hatfield, United Kingdom, September 2006), IEEE Press.

[7] Laban R., and Lawrence F.C. *Effort*. Macdonald & Evans Ltd., London, 1947.

[8] Laban R., *Modern Educational Dance*. Macdonald & Evans Ltd., London, 1963.

[9] Camurri A., De Poli G., Leman M., and Volpe G. Toward Communicating Expressiveness and Affect in Multimodal Interactive Systems for Performing Art and Cultural Applications, *IEEE Multimedia Magazine*, *12,1* (Jan. 2005), 43-53.

[10] Camurri A., Coletta P., Massari A., Mazzarino B., Peri M., Ricchetti M., Ricci A., and Volpe G. Toward real-time multimodal processing: EyesWeb 4.0. In *Proc. AISB 2004 Convention: Motion, Emotion and Cognition* (Univeristy of Leeds, 2004).

[11] Hunt, A.D., Paradis, M., and Wanderley, M., The importance of parameter mapping in electronic instrument design. *Journal of New Music Research, 32, 4,* (December 2003), 429-440.

[12] Wanderley, M.M., and Battier, M., (Eds.), *Trends in Gestural Control of Music.* IRCAM - Centre Georges Pompidou, Paris, 2000