

DESIGNING THE BALANCE BETWEEN SOUND AND TOUCH: METHODS FOR MULTIMODAL COMPOSITION

Claire Richards,^{abc} Roland Cahen,^b and Nicolas Misdariis^c

^aActronika SAS

^bLe Centre de Recherche en Design ENSCI Les Ateliers/ENS Paris-Saclay

^cSTMS-IRCAM-CNRS-SU-Ministère de la Culture

clairicha@gmail.com, roland.cahen@ensci.com, misdariis@ircam.fr

ABSTRACT

In this paper, we discuss key research questions generated from a collaborative workshop, during which our aim was to explore the potential of a wearable device to produce novel audio-haptic sensory experiences. The main intention of this research is to enable users with any type of hearing profile to appreciate a body-centered listening experience. The multimodal harness, developed by the co-authors, integrates nine voice coil actuators into a wearable structure, stimulating both the auditory and tactile senses via extra-tympanic sound conduction and vibrotactile stimulation of the skin on the upper body (spine, clavicles, ribs). We need to create our own interfaces in order to judge its capacity to elicit three modes of vibratory sound perception: auditory, tactile and bi-modal. To this end, we used the Max environment to make several preliminary authoring tools, whose compositional features allow us to explore three main themes of sensory composition: multimodal music listening experiences, spatialization of audio-haptic signals on the body, and sensory equalization. Feedback from trial sessions, along with current constraints due to the wearable and interface design, give us direction for our future work: iterate to improve the sensory experience of the multimodal harness, and apply these tools experimentally in order to contribute to multi-sensory processing research.

1. INTRODUCTION

The experiences of grazing one's hand across a soft surface and listening to a concerto may not seem to have much in common. However, in both cases, the body is detecting and interpreting vibratory information, whether that vibration is a result of the friction produced by the movement of skin on a surface, by the movement of performers' bows on the strings of violins, or by touching the surface of a speaker playing music. In our research, we seek to illuminate the physical aspect of sound perception, by bringing the vibrations directly in contact with the body. In parallel to the development of our multimodal wearable device, we conducted a psychophysical study on extra-tympanic hearing thresholds of the human torso, with the aim to begin a definition of audio-haptic signal

design guidelines for the wearable [1]. In this study, instead of taking a psychophysical approach, we decided to start by applying the basic parameters obtained from the first study's results within a more creative developmental framework: an intensive workshop.

Our intention in organizing this workshop was not to produce data, but to gain more clarity on the research questions we can ask. In the next sections, we first present the project context of exploring a set of audio-haptic forms of expression, followed by basic technical details of the multimodal harness. Then, we describe functionalities of authoring tools we developed, trial feedback, and the research questions generated by the sensory experience they permit (or not). Finally, we discuss our perspectives on future development, for both experimental and compositional purposes.

2. CONTEXT AND MOTIVATIONS

The perceptual ranges and discriminatory limits of hearing and touch are compatible and overlapping. Understanding the body's ability to interpret vibrotactile information can be broken down into several key characteristics of the presented stimulus: its frequency, intensity, duration, waveform and position on the skin's surface. The elements of vibrotactile signal design have much in common with those of auditory signal design (pitch, amplitude, duration, timbre, distance), but the two senses do not process the sensory information in the same way [2]. Opposed to hearing, which is centralized to the head, we perceive vibrations all over our body [3]. The perception of vibrations in the frequency range between approximately 10-1000 Hz overlaps the detection capacities of the two senses.

Though there are distinct differences in processing capacities between the two sensory organs, researchers have found that auditory stimulation can affect tactile perception [4, 5], and vice versa [6, 7]. Beyond influencing the perception of presented auditory stimuli, tactile stimuli alone have been found to activate the auditory cortex, both in normal-hearing and congenitally deaf adults. Researchers have found this effect to be greater for deaf versus hearing participants [8]. The mutual facilitation of one sense on the perception of the other can be attributed to

processes of multimodal integration and cross-modal enhancement [9].

The perception of hearing and touch can therefore be influenced separately and simultaneously, by using the same source of stimulation: vibration. In this paper, we discuss the excitation of the two modalities by means of extra-tympanic conduction of sound and vibrotactile stimulation of the skin. The existing literature on extra-tympanic conduction provides an essential foundation towards understanding how we are able to hear sounds that do not stimulate the outer ear. However, notable past investigations of these mechanisms are primarily based on the anatomy of the skull [10]. Commonly described as bone conduction, this paper refers to contact-based sound transmission as “extra-tympanic”, because there is speculation that soft tissue is involved in transmission mechanisms, and researchers have found that sound can be conducted from other parts of the body [11].

Many researchers have therefore investigated the use of vibration technology to augment auditory experiences by creating interfaces for the body. A large portion of these interfaces were specifically created for the improvement of musical experiences, and within this category, many devices have been developed specifically for the back and upper body. The model human cochlea [12] is a tactile device that separates audio signals into frequency bands from 20-1000 Hz, displayed to the back via 16 voice-coil actuators, arranged on the back of an armchair. The *Hap-beat* and *Vibeat* both integrate motors that vibrate in time with music, worn as a necklace or bracelet. The *Sound-shirt* integrates 30 micro-actuators (type unspecified) in a thin, flexible garment created for hearing-impaired individuals to have a tactile experience of philharmonic music. Users feel lower frequencies in the lower abdomen, while higher frequencies are displayed across the upper body. This frequency-based organization of tactile stimulations draws comparison to the model human cochlea while it differs from the *Subpac*, a backpack-style vest that enhances bass frequencies via two large voice coil actuators located on the back.

Our research differentiates itself from the prior art through our approach to mixed-modality effect design. The resulting experience of sound, perceived either by means of extra-tympanic hearing, tactile stimulation of the skin or a mixture of both modalities, is accessible to all hearing profiles. Hearing loss comes in many shapes and sizes, differing based on the frequencies affected, severity, symmetry, and type of lesion (conductive, sensorineural or mixed) [13]. Someone with severe sensorineural hearing loss could appreciate tactile sensations, while an individual with conductive hearing loss could hear the wide range of acoustic frequencies displayed via extra-tympanic conduction. Beyond catering to diverse hearing profiles, our project addresses normal hearing profiles as well, and is inspired by the principles of universal design: conception whose aim is to include all users, regardless of ability.



Figure 1. Image of the multimodal harness. There are nine actuators integrated on the structure: two on the clavicles, two on the ribs, and five along the spine. The straps are velcro-adjustable, two of which go around the legs to secure the last module against the body. The multichannel audio-haptics card is connected to the harness, situated below the last module on the spine.

3. THE MULTIMODAL HARNESS

3.1. Conception and Functionality

Before beginning the design process¹, we defined two key constraints for the multimodal harness: adhesion of the module against the body, and user comfort. In order to achieve effective sound transmission and avoid signal attenuation, the vibrating module should be made of hard material, secured tightly against the body. Balancing tightness and comfort, two contradictory qualities, became a principal design challenge which we addressed with the flexibility of the modules and the ergonomic aspects of the harness structure (fig. 1).

The technical aspects of the module (fig. 2) allow them to adapt to the body’s form at the different points of stimulation (spine, clavicles, and ribs). The hard plastic parts in direct contact with the user are ergonomically shaped (wide and flat on the ribs, extruded on the lower spine, rounded on the clavicles) in order to maximize the ratio of comfort to auditory signal transmission. The harness structure (fig. 1) has adjustable Velcro straps attached to all corners of the modules, pulling them against the body. Nine of Actronika’s *proprietary voice coil actuators* are connected to the *Mk-1 20-channel audio-haptic card*. The Max environment interfaces with the USB Audio output on the Mk-1 board, which is detected as a sound card.

¹ A co-development organized with *Actronika*, industrial and research partner, and *Les Vertugadins*, a Parisian costume design studio.




Figure 2. Schema of the module, which integrates the actuator. The module snaps onto extended fabric so it can “breathe” up and down on its frame (B), attached via the PCB clip (C). The 3D-printed clip-on piece (A) attaches onto the indented edges of the actuator, and comes into contact with the body, transmitting the vibrations to the user.

Each of the nine actuators on the multimodal harness has its own channel on the multichannel audio-haptic card. In its current wired design, users’ movements are restricted to the length of the 9V power supply and audio signal transmission cables.

4. DESIGN WORKSHOP

The two-week workshop, organized by the co-authors, took place at IRCAM in Paris, France in August 2021. During the workshop, five participants (P.01, P.02, etc.) tested our tools based on their own interest or direct involvement in the project: they were neither specifically selected nor remunerated for their participation. During every trial of the multimodal harness, users wear ear plugs and a noise-canceling headset in order to mask the external noise created by the modules’ vibrations and focus on their hearing via extra-tympanic conduction.

The main objective of the design workshop was to create first drafts of audio-haptic signal design interfaces. Creating these tools would allow us to refine our assumptions about the sensory experiences afforded by the multimodal harness, informed by our previous psychophysical study, haptic illusions of signal movement [14], and existing research on vibrotactile musical experiences [2, 12].

We hypothesize that the vibratory activation of tactile, auditory and bi-modal perception enables myriad possibilities for signal design: sensations that can travel in all directions to create illusions of movement, sound perception by vibratory transmission through musculoskeletal structures, and novel effects based on the multimodal integration of both sensations. While it is possible to approximately localize tactile stimuli on the body’s surface [15], auditory perception by means of extra-tympanic conduc-

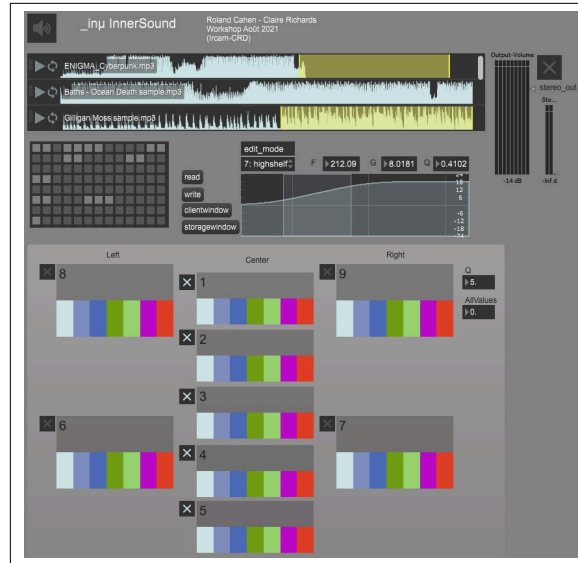


Figure 3. Inner.Music as seen in Max environment. Any sound file can be uploaded into the playlist object at the top of the interface. By default we propose a short list of rhythmic music with powerful bass content: electronic, disco-funk and pop. Left and right channels are routed to the left and right lateral actuators on the clavicles and ribs and mixed together on the vertical positions of the spine. A graphic equalizer and a mute button is available for each actuator, with different colored bars for 7 frequency bands spanning 50-4000 Hz. The user can save their equalizer settings in the grid below the playlist object, recorded and saved for future use. To the right of the preset grid, a high shelf filter object, applied to all channels, attenuates the bass frequencies with respect to the high frequencies due to the resonant frequency of the actuator (70 Hz).

tion on the torso is essentially mono, since all sounds are entering the inner ear from the same pathway: the cervical spine. In other words, spatial audio information is lost, but we hypothesize that the localization of tactile stimuli could remain strong enough to transmit impressions of spatialized auditory source positions.

5. AUTHORING TOOLS FOR THE MULTIMODAL HARNESS

Each of the following authoring tools offers a different possibility for designing audio-haptic effects. In this section, we describe each of the authoring tools in three parts: the interface design, feedback from trial sessions, and finally key questions and limitations related to interface functionalities and perceptive experiences.

5.1. Inner.Music: Multimodal Music Player

Inner.Music is a music player, designed to test the perception of filtered source files transmitted to the multimodal




Figure 4. Photo of P.01, taken during the workshop. Since he wasn't satisfied with the quality of sound perception, he placed his fingers directly on the modules to enhance his tactile perception.

harness. This tool features a body point-based equalizer tool, which allows the designer to calibrate the signals to each position of stimulation.

5.1.1. Interface Design and Composition Features

Inner.Music allows the designer to self set, adjust and balance the volumes of each actuator's filter bands, while listening to music. The interface is also a preset manager: individual sensory profiles, created by saving the nine filter band settings in presets, can vary each person's sensory preferences (see fig. 3).

Our initial intention for Inner.Music was to create a tool that adjusts each actuator's spectrum in order to obtain a homogeneous audio-haptic sensation across all stimulated positions. However, we understood during the trial sessions that the ranges and nature of the two modalities' perceptions vary across positions, and according to different types of effect design or scenarios of use. We therefore chose to exploit these variations, creating an interface that allows the user to tweak their own sensory experience instead of imposing a generic predetermined sensory calibration.

5.1.2. Trial Feedback

While toying with the Inner.Music interface, participants employed various strategies: balancing the signal to achieve a perception of homogeneity across the positions, attributing levels and spectrum bands to each position in order to focus their perception on either the tactile or auditory modality, or muting positions that they did not appreciate. For example, some favored a stronger tactile sensation of the bass in the lower back, and routed the trebles to the upper back and clavicles to hear via extra-tympanic conduction. A common remark was that the music is audible,

but appears to be coming from a distant place within their body, and in order to hear more details, they have to focus their attention. P.03 said, "It's possible to hear the sound, but to feel practically nothing." P.05 was especially sensitive to extra-tympanic auditory perception, saying: "I can hear the sound as if I were listening with headphones - maybe even better". On the other hand, P.01 was disappointed by the quality of the auditory experience and resorted to actively touching the modules to enhance his tactile perception (fig. 4). These differences in sound perception could be attributed to morphological differences affecting the transmission of the mechanical waves within the body [11]. A user's expectations may also impact how they judge their sensory experience: they might compare the quality of their auditory experience to other systems that were specifically designed for the ears, not the body.

5.1.3. Key Questions and Limitations

The strong morphological differences between each individual can be compared in a similar way to how differences in the structure of the outer ear can affect one's hearing. The Head Related Transfer Function (HRTF) is a physical transfer function that describes how a given sound coming from a specific point will reach the ear, in terms of spectral characteristics [16]. In comparison, we wonder if we could develop a Body Related Transfer Function (BRTF), a functionality capable of adjusting the multimodal harness's response to each individual in order for them to perceive the same signal, based on their morphological differences (height, weight, musculoskeletal characteristics).

5.2. Mp2p: Body Spatialization Interface

We created the Mp2p (Mono point-to-point) interface as a creative tool for spatialization of multimodal signals on the harness, by manipulating the signals' positions and trajectories.

5.2.1. Interface Design and Composition Features

Within the design space, simple mono signals can be played on any single actuator or grouped set of actuators and panned² between them according to a set 1D path. Spatialization of the multimodal signals in the Mp2p system works based on configuration presets which organize the points of stimulation along different paths. The signal can be routed to each independent point, one channel after another, or to several points at once, by grouping channel outputs together. In fig. 5, we see three possible configuration presets: these allow the designer to create different illusions of movement (sporadic in the first preset, side-to-side in the middle preset, top-to-bottom in the right preset). The actuators' positions on the lateral and vertical axes of the body allow for directional and alternating effects. We defined a set of configurations, providing a

² Linear panning with Max [mc.mixdown @pancontrolmode 2] object gave the best results.




Figure 5. Diagrams of three possible Mp2p spatialization path configurations. On the top of the figure, arrays of colored dots represent the signal path: each different color is a different channel. Below each signal path is a graphic representation of the nine positions on the harness structure: the upper horizontal line is the clavicles, the lower is the ribs, and the vertical line is the spine. In the left-most configuration, each position has its own independent channel. The middle configuration uses only two channels (light and dark green), with two positions paired to each channel - the five spine points are inactive. The right-most configuration groups the positions among four channels: the clavicles and top spine point (dark green), the second spine point (light green), the middle spine and ribs (cyan) and the lowest spine point (navy).

diverse selection of routing structures for the spatialized effects. The designer can manipulate the signal's position, duration, waveform (sine, saw tooth or noise), frequency (50 Hz to 4 kHz), and linear amplitude (0. to 1. float multiplier), and play their composition looped or one-shot (see fig. 6). By manipulating these variables, the sensory effects can be perceived as having continuous or discrete inter-position transitions: the signal can move gradually, or switch abruptly along its path.

5.2.2. Trial Feedback

The Mp2p interface is slightly more complex than the other two tools, so after learning about the basic parameters, participants generally chose to tweak our predefined patterns rather than create their own from scratch. P.03 was particularly impressed by the sensation created by the signal moving up and down the spine. If the signal moved quickly enough along the points, he felt a continuous sensation, like a "hand running up and down [his] back." The slower the movement of the signal along the points, the easier it was to localize from which position it came. The illusion we tried to create of a front-to-back shifting sensation between the front actuators (clavicles and ribs) and back (spine) was less convincing, perhaps due to the slightly central position of the actuators on the ribs, and because the sensation of something crossing one's body is a rather alien concept. The duration of the experience as a whole also impacted the sensory experience: after trying the patterns, P.02 said:

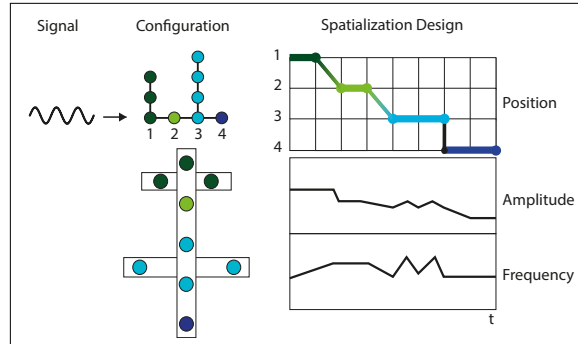


Figure 6. Graphic representation of the Mp2p interface functionality. A sine-wave signal follows a 1D path according to the chosen configuration preset (see fig. 5). To the right of the figure is the spatialization design, with three curves for position, frequency and amplitude. By modifying the curve, the designer can fine-tune the signal's duration at each group of positions, and its variations in amplitude and frequency across the composition.

I need more time to become familiar with the sensations I'm feeling on and in my body, since they are so new. Just a few minutes is not enough for my brain to really understand what is happening. I think with more time I'd be able to appreciate the sensations a bit more.

Participants made more comments on their "outer" body (tactile perception) than their "inner" body (extra-tympanic sound perception) during the spatialization experience. For most participants, the signal was only audible when displayed to certain positions (upper back and clavicles, i.e. those closest to the inner ear). The bi-modal aspect of the spatialization experience may therefore depend not only on a signal's position, intensity and spectral characteristics, but also on the user's conscious attention on either modality.

5.2.3. Key Questions and Limitations

When we started to design Mp2p, we tried to use an existing 3D spatializer tool: a vector-based amplitude panning (VBAP) algorithm [17]. However, our configuration did not permit this because the listening point (the head) was off-center, i.e. not based on diffusion points (speakers) placed around the origin (listener). If we use a classic calculation of spatial decay ($gain = 1/distance^2$) according to airborne sound transmission, the signal disappears as soon as we move the point of diffusion away from the actuators' exact positions. Realistically, we would need to create a more specific spatialization model which takes into account the body's responses to extra-tympanic conduction and tactile stimulation of the skin, as described in section 4.1.3. The listening point (i.e. listener's position) would be the head, while the haptic sensations would be distributed: they would not refer to any central reference or origin point, since tactile perception is decentralized and perceived all over the surface of the skin (i.e. each

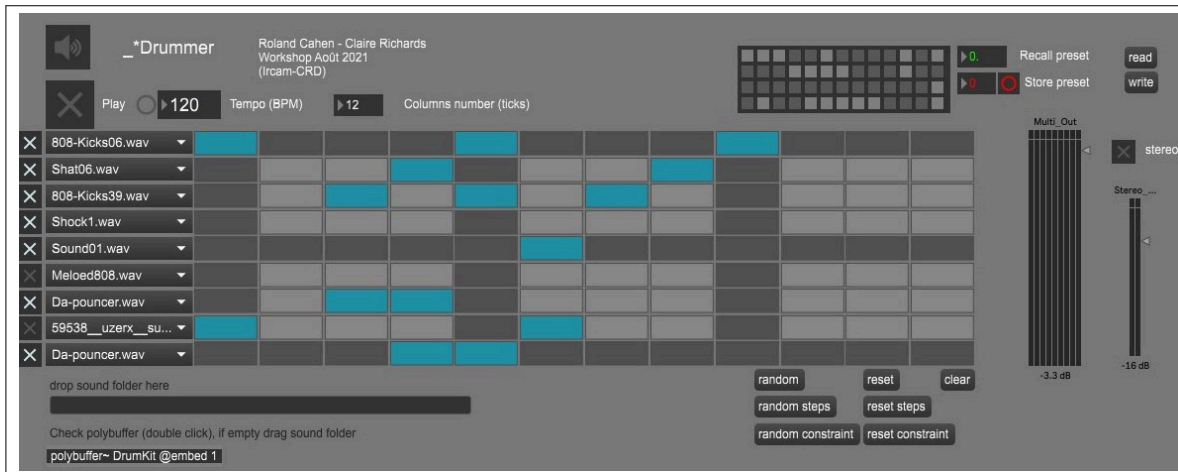


Figure 7. The Drummer interface seen in the Max environment. The sequencer tracks each correspond to one actuator position on the harness. Above the sequencer, there is a preset grid for saving specific drumming patterns.

actuator’s position would have its own sensory reference).

One key limitation of Mp2p is the static nature of the configuration preset system. The signal path configurations are saved within presets and are static within a single composition. In other words, the designer cannot use multiple configuration presets in one spatialization time frame, limiting each composition to a single path that the signal can follow. One possible solution is to create a state-based programming interface which would include a configuration for each state, and a sequencer that allows the designer to edit each successive state and concatenate them with discrete or progressive transitions between states. In its current state, Mp2p does not integrate the equalizer profile function of Inner.Music, so in order to calibrate the signal intensity for the different sensitivities of each point of stimulation, designers must manually tweak and adjust the amplitude levels along the duration of the displayed signal. It is time-consuming to find the right nuances of intensity and frequency values for a given point in order to obtain the intended perceptive results. A final limitation of this tool is that the signal is mono: we cannot create "harmonies" of different sounds and movements. In section 6, we discuss some steps to address this constraint.

5.3. Drummer: Body-based Drum Machine

While we discussed musical metaphors during the workshop, we decided to create a basic nine-track step sequencer and drum machine.

5.3.1. Interface Design and Composition Features

Each track on the interface corresponds to one actuator and plays one selected sound file, imported from the user’s library. The designer creates their drumming pattern by activating buttons on the sequencer, creating repetitive patterns of stimulation in which each actuator/position per-

forms a sequence of an instrument (see fig. 7). The patterns can be saved as presets.

5.3.2. Trial Feedback

Participants described using the Drummer tool as a fun experience: the interface is simple and user-friendly, patterns can be quickly designed, and the ability to assign different instruments to the different points of stimulation is unique to this tool. Participants tended to move around while testing it, in sync with the rhythm. We heard several remarks about enhanced awareness of one’s own body, and the presence of the wearable device. P.02 said, “I feel like the different instruments are playing directly on my body, or like I am the instrument!” She added that, since she heard the different sounds at the same time as she felt the stimuli on the different positions, the sound seemed to be coming directly from those points. This comment suggests that, even though the spatial audio information is lost in extra-tympanic conduction, the localized perception of tactile stimuli can contribute to impressions of auditory source positions. This comment also branched off into a discussion about sensory effects for physical rehabilitation, perhaps to help target the user’s attention on a specific part of their body.

5.3.3. Key Questions and Limitations

In the Drummer interface, the opportunities to create musical patterns remain basic: the designer can select the output channels, choose sound files, alter the tempo and number of columns, but cannot change relative signal intensity according to point of stimulation. The interface functionalities could be enriched with multichannel or MIDI sequences, an equalized distribution of the signals so that their intensities can be tailored to each point of stimulation, or a more dynamic variation of each track’s intensity.




Figure 8. Photo taken during the workshop. P.02 bends forward to change her perception of the transmitted vibrations.

One could also imagine guiding the user's movements by distributing different sensations to different parts of the body. However, body movements and changes in posture interfere with audio-haptic perception. The pressure of the modules against the body changes based on the way the user stretches their joints and limbs (see fig. 8). For example, when the user bends downwards, or stretches their arms in front of them, they pull the modules closer against their spine. The added pressure on the points of stimulation, along with the difference in the curvature of the spine, facilitates good contact of the modules and thus good transmission of the vibratory signals. However, if the user raises their arms above their head, the upper back arches inwards, creating more distance between the modules and the skin and reducing the vibratory transmission. This aspect of use of the multimodal harness could be seen as a limitation but also as an advantage: though the body's movements cause variations in audio-haptic perception, those variations are part of the physical interaction between the user and the wearable, and could be a source of perceptual cues for movement-related use cases. Modifying accents, rhythm or instrument positions according to the user's movements would be interesting to explore, but this implies the integration of motion and pressure sensors which is currently out of the project's scope.

6. DISCUSSION

During the two weeks of the workshop, we were able to address each theme of audio-haptic composition (multimodal music, sensory equalization, and signal spatialization). However, we experienced developmental setbacks due to time constraints and a lack of knowledge about user perception. Evidence of these current limitations can be found in several aspects of the interface functionalities: a sensory equalizer function is present in only one of the authoring tools (Inner.Music), the spatialized audio-haptic

signals are mono and limited to basic wave types, and we are restrained to one-dimensional position configurations. Although they contain some similar elements, the authoring tools are distinct from one another, and the resulting sensory experiences are limited to each interface's distinct functionalities.

These constraints are temporary: we aim to streamline and optimize our authoring tools. For example, in future iterations, features could be consolidated into one interface: the sensory equalizer in Inner.Music, the curves in Mp2p for modulation of position, frequency and amplitude, and the position-specific sound selector in Drummer. For example, in a unified interface, the designer could use the sensory equalizer to first determine a baseline audio-haptic calibration that suits their perceptual preferences. They could then use Drummer's function to select specific sound files to display at each position, and then use Mp2p's curves for spatialization pattern design to map the signal display at each position, along with more fine-tuned variations in amplitude and frequency across the duration of the composition. Inspired by existing tools for audio spatialization [18], we have already begun to optimize the mono-source design constraint of Mp2p in a separate interface (called "MC-Curv"). Instead of one single spatialized source, the designer could edit and spatialize up to four sources on the body at the same time, creating harmonies of movement across the stimulation points instead of single-path sensations.

Regarding the design of the multimodal harness, we hope to improve many aspects in the next iteration: sound transmission quality, comfort, and ease of adjustability. The second iteration of the harness involves changes for both the wearable structure and the module. The actuators will be integrated in an orthogonal position relative to the surface of the body, instead of parallel. Though they will protrude from the surface of the device, this will allow the voice-coil actuator's lengthwise-directed vibrations to penetrate the surface of the body rather than rubbing horizontally across the skin. We hypothesize the change in orientation will increase sound perception in the mid/high frequency range, bringing more clarity to certain instrumental and vocal sounds, without impacting the quality of tactile stimulation.

7. CONCLUSIONS

Since there is no generic go-to tool for audio-haptic composition, we faced the same developmental challenges as past composers and researchers: we wanted to create and study multisensory effects, but we didn't have the tools to do so. In [2], the authors complain of the "awkwardness" of designing the vibrotactile sensory effects using Protools, a standard Digital Audio Workstation. Nine years later, in [14], the authors used Premiere Pro to create "vibetracks" for a film, by placing clips of different sine waves one-by-one at precise points along the video track. Today, researchers acknowledge that tactile composition is still "largely unexplored", that composers need to "take into

account haptic perceptual effects" and that they develop their own design tools for this purpose [19]. Resorting to these painstaking methods drastically slows down the design process, and makes it difficult for progress to happen in the domain.

Another obstacle to overcome is the general lack of guidelines about variations in both auditory and tactile sensitivity across different sites on the body, and across individuals. Established ranges of sensitivity according to each position of stimulation could help provide compositional suggestions, or presets according to each modality. To this end, in an upcoming study we will follow up on our past psychophysical research [1] and evaluate participants' responses regarding both tactile and auditory perception while using the multimodal harness.

The authors call for smoother workflow design for audio-haptic composition, accessible to the research community. We hope that the products of our workshop can foster discussions about what features and guidelines might contribute to an ideal audio-haptic composition interface.

Acknowledgments

The authors extend a special thanks to Yann Boulet (les Vertugadins) for defining the fine details of the module and the harness structure, and Fanjun Jia (Actronika) for his contributions in 3D design. We also thank Emmanuel Flety from the *Pôle Ingénierie et Prototypes* (PIP) of the IRCAM STMS Lab, les Vertugadins for their continued collaboration, and Actronika for their essential advice and technical support. Finally, we thank all of our friends and colleagues who tried using the harness. Funding for this research comes from the *Association Nationale de la Recherche et de la Technologie* (ANRT) and Actronika, in the framework of a CIFRE contract.

8. REFERENCES

- [1] C. Richards et al., "Vibratory Detection Thresholds for the Spine, Clavicle and Sternum", in *Proc. IEEE World Haptics Conference*, Jul. 2021, p. 346.
- [2] E. Gunther and S. O'Modhrain, "Cutaneous grooves: Composing for the sense of touch", *Journal of New Music Research*, vol. 32, no. 4, pp. 369-381, 2003.
- [3] S. Merchel and M.E. Altinsoy, "Psychophysical Comparison of the Auditory and Vibrotactile Perception-Absolute Sensitivity", presented at the Intl. Workshop on Haptic and Audio Interaction Design, Lille, France, March 13-15, 2019.
- [4] J.P. Bresciani et al., "Feeling what you hear: auditory signals can modulate tactile tap perception", *Experimental Brain Research*, vol. 162, no. 2, pp. 172-180, 2005.
- [5] L. E. Crommett, A. Pérez-Bellido, and J.M. Yau, "Auditory adaptation improves tactile frequency perception", *Journal of Neurophysiology*, vol. 117, no. 3, pp. 1352-1362, 2017.
- [6] M. Schürmann et al., "Touch activates human auditory cortex", *Neuroimage*, vol. 30, no. 4, pp. 1325-1331, 2006.
- [7] H. Gillmeister and M. Eimer, "Tactile enhancement of auditory detection and perceived loudness", *Brain Research*, vol. 1160, pp. 58-68, 2007.
- [8] Auer Jr, E.T. et al, "Vibrotactile activation of the auditory cortices in deaf versus hearing adults", *Neuroreport*, vol. 18, no. 7, pp. 645-648, 2007.
- [9] S. Shimojo and L. Shams, "Sensory modalities are not separate modalities: plasticity and interactions", *Current Opinion in Neurobiology*, vol. 11, no. 4, pp. 505-509, 2001.
- [10] S. Stenfelt and R.L. Goode, "Bone-conducted sound: physiological and clinical aspects", *Otology and Neurotology*, vol. 46, no. 12, pp. 1245-1261, 2005.
- [11] C. Adelman et al., "Relation between body structure and hearing during soft tissue auditory stimulation", *BioMed Research International*, pp. 1-6, 2015.
- [12] A. Baijal et al., "Composing vibrotactile music: A multi-sensory experience with the emoti-chair", in *IEEE Haptics Symposium*, March 2012, pp. 509-515.
- [13] R.H. Margolis, and G.L. Saly, "Toward a standard description of hearing loss", *International Journal of Audiology*, vol. 46, no. 12, pp. 746-758, 2007.
- [14] F.A. Geldard and C.E. Sherrick, "The cutaneous 'rabbit': A perceptual illusion", *Science*, vol. 178, no. 4057, pp. 178-179, 1972.
- [15] R.W. Cholewiak, "The perception of tactile distance: Influences of body site, space, and time", *Perception*, vol. 28, no. 7, pp. 851-875, 1999.
- [16] F.L. Wightman and D.J. Kistler, "Headphone simulation of free-field listening. I: stimulus synthesis", *The Journal of the Acoustical Society of America*, vol. 85, no. 2, pp. 858-867, 1989.
- [17] V. Pulkki, "Virtual sound source positioning using vector base amplitude panning", *Journal of the Audio Engineering Society*, vol. 45, no. 6, pp. 456-466, 1997.
- [18] L. Pottier, "Holophon: Projet de spatialisation multi-sources pour une diffusion multi-hautparleurs", in *Journées d'informatique musicale*, Bordeaux, France, May 2000.
- [19] L. Turchet, T. West and M. Wanderley, "Touching the audience: musical haptic wearables for augmented and participatory live music performances", *Personal and Ubiquitous Computing*, vol. 25, pp. 749-769, 2021.