Real-Time Musical Haptics With Ultra-Wideband: A Study on Latency, Reliability, and Perception

Luca Turchet, *Senior Member, IEEE*, Christian Sassi, Davide Vecchia, and Gian Pietro Picco, *Senior Member, IEEE*

***Abstract*—Ultra wideband (UWB) radios are popular for accurate distance estimation between devices. However, UWB also offers low-power, fast, reliable wireless communication. We exploit it here in a *real-time musical haptics* system for live performances: a wearable, wirelessly activated via UWB by the performer's instrument, augments the audience *musical* experience with a *tactile* sensory layer. Two challenges are crucial to the experience quality: *i)* communication must be *reliable*, to prevent corruption of tactile signals, and *ii)* these must reach the audience *synchronously* with the instrument sounds. We perform micro-benchmarks of UWB links alone in a controlled setup, showing that the haptic signal can be delivered reliably over UWB *before* the instrument sound, thus enabling proper compensation delays to *perfectly* realign sound and tactile vibration. We confirm this holds on the end-to-end system including haptic components by characterizing four proof-of-concept prototypes combining different UWB-enabled instruments and wearables. Finally, we reconcile these objective measures with subjective ones via a user study focusing on *perception*, yielding very positive outcomes. Together, these results confirm the potential of UWB-based musical haptics for enhancing the audience experience at live performances in ways hitherto unexplored.**

***Index Terms*—Musical haptics, tactile music augmentation, ultra-wideband, Internet of Musical Things.**

I. Introduction

MUSICAL Haptics is a young field of research that investigates touch and proprioception in musical contexts [1], [2]. Several disciplines have contributed to this field, including haptic engineering, human-computer interaction, applied psychology, musical acoustics, aesthetics, and music performance [3]. Recently, a new area of research in the music technology domain has started to impact the musical haptics field, the Internet of Musical Things (IoMusT) [4].

The IoMusT relates to the extension of the Internet of Things paradigm to the musical domain, leveraging technologies such as networking, embedded audio systems, and artificial intelligence. The IoMusT vision has paved the way for novel forms of interactions between performers and audience members, mediated by the network and musical devices. Among these interactions, is the case where performers play on stage for a co-located audience, equipped with haptic devices connected via wired or wireless links to the musical instruments used by performers. These devices are intended to enhance the musical listening experience by adding a layer of vibrotactile sensations to the sound people hear [5], [6], [7].

Today, several wireless technologies are available to developers to create IoMusT musical haptics systems, including popular ones like Bluetooth, WiFi, and 5G. For instance, the recent Bluetooth-based Auracast system enables broadcast audio with a latency of only 20 ms [8]. Much less explored is the use of ultra-wideband (UWB) radios (§II), which offer short-range, fast, and reliable communication along with the capability to accurately estimate distance motivating their typical use in localization systems [9], [10]. However, to date UWB communications in musical contexts, let apart involving haptics, have been largely neglected by researchers.

To fill this gap, we focus on the Spark Microsystems SR1020 [11] UWB transceiver that, unlike popular ones like the Qorvo DW1000 [12], explicitly targets low-latency, high-rate, audio streaming. Nevertheless, our goal is *not* to compare UWB vs. the wireless technologies above from a pure communication perspective. Our experiments (§IV) actually show that UWB achieves a latency of *one order of magnitude smaller* than the 20 ms reported for Auracast; yet, this should be confirmed by a comparative study in the same experimental conditions, which is entirely out of the scope of this paper.

Our goal, instead, is to investigate whether UWB is suitable for musical haptics scenarios involving an audience co-located with performers delivering both audio and tactile signals in real-time, therefore considering the entire end-to-end chain from the generation of sound down to the delivery of the haptic signal. Fig. 1 illustrates the high-level view of our proposed system, which we implemented and evaluated under several dimension, including and beyond UWB communication.

The key challenge of this scenario, however, is that the two sensory contents, audio and haptics, are *synchronously* delivered to the ears and body of the audience member, respectively. No *perceivable* delays or anticipations should occur, disrupting the musical experience.

Addressing this challenge is non-trivial. The latency of sound between performer and listener depends solely on distance and the speed of sound in air, ≈343 m/s in normal conditions,

Received 20 June 2024; revised 16 October 2024 and 8 December 2024; accepted 2 January 2025. Date of publication 6 January 2025; date of current version 21 March 2025. This work was supported in part by the Italian Government via the NG-UWB project (MIUR PRIN 2017) and in part by European Union—Next Generation EU under the Italian National Recovery and Resilience Plan (NRRP), Mission 4, Component 2, Investment 1.3, CUP B53C22003970001, partnership on "Telecommunications of the Future" (PE00000001—program "RESTART"). This article was recommended for publication by Associate Editor J. Park and Editor-in-Chief D. Prattichizzo upon evaluation of the reviewers' comments. *(Corresponding author: Luca Turchet.)*

This work involved human subjects or animals in its research. Approval of all ethical and experimental procedures and protocols was granted by the Ethical Committee of University of Trento.

The authors are with the Department of Information Engineering and Computer Science, University of Trento, 38100 Trento, Italy (e-mail: luca.turchet@unitn.it).

Digital Object Identifier 10.1109/TOH.2025.3525959

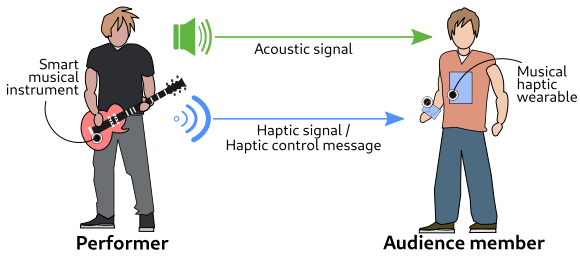

Fig. 1. A schematic diagram of the audio and haptic signal flows involved in an IoMusT-based musical haptics application.

and cannot be modified. On the other hand, the overall latency induced by the haptic system depends on its various components encompassing the generation of the haptic signal, its transmission over the UWB link, and its use in haptic devices at the listener.

Ascertaining whether the haptic latency is lower than the sound one, and can therefore be time-compensated at the receiver, entails an understanding of the contribution of each of the components above. A similar reasoning holds for reliability, where nonetheless the dominant factor is UWB communication; packet loss may degrade the haptic signal preventing its correct reconstruction.

We define a model (Section III) accounting for all these aspects and providing the reference framework for the rest of the paper, where we determine the practical feasibility of the envisioned system via real-world experiments. We begin by characterizing the latency and reliability of the UWB communication link (§IV) offered by the SR1020 [11] transceiver, which targets audio streaming applications, and also enables direct connection between audio input/output and hardware boards, greatly simplifying the development task. However, this radio is significantly less popular than other UWB ones and, to the best of our knowledge, its use is unreported in the literature. Therefore, our evaluation of its communication support for audio stream, via state-of-the-art equipment and in different communication conditions, is a contribution per se.

However, the UWB link is only one of the components contributing to the haptic latency. We evaluate the others by relying on accurate measurements of proof-of-concept prototypes (§V) composed by two types of smart musical instruments [13] wirelessly connected via UWB to two types of musical haptic wearable. The resulting four combinations, along with the different conditions above, enable us to characterize precisely the overall haptic latency, and pinpoint precisely the combination of components and communication conditions that may be problematic.

Nevertheless, the measurement above are all *system-level and objective*; they allow us to identify the worst-case latency and reliability, but do not inform us about whether their corresponding values bear an actual impact on the inherently *subjective* dimension of *user perception*. To this end, we perform a user study (§VI) in which we administer, in a *repeatable* fashion, audio-haptic stimuli akin to the ones in the previous experiments, and ask the subjects to rate their experience along several perception dimensions.

The outcomes of this multi-faceted study are positive and confirm the practical viability of our proposed system, as summarized in the discussion (§VII) that precedes our concluding remarks and outlook on research opportunities (§VIII).

To the best of our knowledge, this study represents the first attempt to exploit the potential of UWB communications to timely and reliably deliver haptic content in a live music setting. Conducting this line of inquiry is important to provide composers of tactile music and designers of musical haptic systems with knowledge about the strengths and limitations of the hitherto unexplored alternative offered by UWB.

II. Background and Related Work

We concisely report on the necessary background and state of the art related to the contribution of this paper.

Ultra-wideband radios: UWB has become a common choice in applications targeting accurate proximity detection and localization [9], [10], as witnessed by its recent integration in smartphones as well as other personal and IoT devices.

The defining characteristic of UWB is its large bandwidth, either $>500\,\mathrm{Hz}$ or greater than 20% of the center frequency. Modern transceivers are referred to as impulse radios, as they transmit information through very narrow pulses ($<2\,\mathrm{ns}$). UWB receivers exploit pulses to obtain precise timing information; reception timestamps can be used to estimate the time of flight (ToF) between devices, enabling distance estimation (*ranging*) and localization with decimeter-level error. This capability, unmatched by other radios [10], is exploited by several IoT applications involving wearable devices, e.g., museum tracking [14], contact tracing [15], motion analysis [16], or gesture recognition [17].

Nevertheless, UWB offers valuable features also for *communication*, significantly less explored in the literature. The large bandwidth translates to a high channel capacity and a higher data rate w.r.t. other low-power radios, up to 31 Mbps [18]. Moreover, as UWB systems operate in the unlicensed bands between $3.1\,\mathrm{GHz}$ and $10.6\,\mathrm{GHz}$, they do not interfere with the crowded $2.4\,\mathrm{GHz}$ spectrum. These unique features are crucial to our target scenario; however, to the best of our knowledge, the use of UWB to support musical haptics applications has never been investigated before.

Systems for musical haptics: In the past two decades, researchers focused on ways of augmenting the experience of listeners in the audience of a live performance by means of vibrotactile stimulation. Several systems have been proposed, broadly divided according to whether their vibrotactile actuator is controlled from audio signal or not [19]. Actuators in the first category (e.g., voice coils or solenoids) are typically capable of reproducing a wide range of frequencies spanning the tactile perception. Examples of their use in musical haptics are reported in [20], [21]. Actuators in the second category are instead controlled by direct current (DC), pulse width modulation signals, or other techniques (e.g., Eccentric Rotating Mass), via discrete signals or continuous signals at low rates, typically via a microcontroller. These actuators usually work at a single frequency in the range of 200–300 Hz, i.e., the frequency range at which human sensitivity to vibrations is highest. Examples can be found in [6], [22].

Existing musical haptics systems can be categorized based on their use. In *offline* systems, the enhanced listening experience might occur during listening to recordings [20], [22], [23], [24],

[25]. Instead, in *real-time* systems the vibrotactile stimuli are delivered at the very moment when the music is produced during a live performance [6], [26], [27].

Finally, connectivity also plays a key role. In *standalone* systems, a single device is capable of delivering audio and haptic signals at the same time [20], [22], [23]. Otherwise, the delivery of audio and haptic content is accomplished via a *wired* [25], [26], [27] or *wireless* [6] connection between the generating device and both the haptic actuators and loudspeakers.

Interestingly, in the latter case existing systems deliver haptic content only in terms of control messages for actuators rather than the actual haptic signal. In contrast, and based on the classification above, our work explores for the first time the real-time, wireless delivery of haptic content in the form of an audio stream for controlling voice-coil actuators.

Audio-tactile synchronization: The coupled perception of sound and vibration is a well-known phenomenon occurring during live music performances. Merchel et al. [28] investigate the influence of audio-induced vibrations on the perceived quality of a concert experience. The results show how vibrations play a significant role in music perception. A natural and realistic multisensory experience of the world results from the integration of inputs from all sensory systems into a single, unified perception. Experiencing live concerts is no exception. To achieve this binding among senses, it is essential to ensure the temporal coherence among the different sensory inputs.

Nevertheless, the synchrony between the auditory and tactile modalities is still an open research topic, with only few studies involving musical stimuli. Moreover, findings are not always in agreement. For instance, concerning non-musical sounds, some studies report that the synchrony between auditory and tactile stimuli when considering vibrations transmitted to the hand [29] falls in the range between -12 ms and 25 ms, the minus sign indicating the anticipation of the vibration w.r.t. the sounds. However, others report a different range approximately between -25 ms and 50 ms [30]. Similarly, when considering whole body vibrations reported ranges are $[-47, 63]$ ms or $[-58, 79]$ ms [31]. These differences are considered to depend on the type of stimulus and on the involved body area [28]. Regarding musical sounds, the point of subjective simultaneity has been reported to vary considerably—from a few milliseconds to more than one hundred—according to the involved musical instrument, which may lead to sounds with radically different attack or decay times [32], [33], [34].

From the few results available in the literature, we can draw the general conclusion that the kind of signal utilized has an impact on the detection of auditory-tactile asynchrony, with impulsive stimuli being more prone to perceived delays among modalities. On the other hand, Navarra et al. [35] have proven the existence of perceptual adaptation mechanisms able to widen the temporal window for auditory-tactile integration after prolonged exposure to asynchronous stimuli—as in the case of the live performance setting we target.

Haptic quality: The studies that have investigated the quality of haptic stimulation and its impact on the user experience mostly focused on the use of haptic stimuli with different frequency ranges (e.g., [36]). Only a handful of studies have assessed the impact on task performance of haptic signal deterioration due to packet losses (e.g., [37]), reporting that haptic discontinuity becomes noticeable for the hands at signal interruptions of about 60 ms. However, to the best of our knowledge no study has been conducted on haptic quality perception of signal deterioration in musical haptic settings.

III. Modeling Latency and Reliability

We begin our study by identifying the latency and reliability contributions of the components of the end-to-end UWB-based musical haptics system that takes a musical signal as an input at the performer and delivers a haptic signal in output on a wearable device at the audience.

The input may be created by a single acoustic, electric, or digital musical instrument, or by a device that mixes the contribution of multiple instruments. As for the haptic signal generated as output, we consider actuators controlled by audio signals (e.g., voice coils or solenoids). Furthermore, in our system characterization we assume that a single haptic signal is transmitted. These assumptions parallel the prototypes developed for this study (§V), but similar considerations can be easily made for haptic systems that transmit over multiple channels, or use actuators controllable by DC or pulse width modulation in conjunction with a microcontroller (e.g., eccentric rotating mass devices).

Latency: We model the overall haptic latency path (Fig. 2) from a musician (sender) to an audience member (receiver) as

$$\Lambda_h = \lambda_{instr} + \lambda_{UWB} + \lambda_{dev} \tag{1}$$

The first component, λ_{instr}, is the delay incurred at the musical instrument for the generation of the haptic signal, and can be further decomposed as

$$\lambda_{instr} = \lambda_{ADC} + \lambda_{outbuf} + \lambda_{trans} \tag{2}$$

accounting for the delays:

- λ_{ADC}, due to the acquisition of the signal to be sent via an analog to digital converter (ADC) present in a Digital Signal Processing (DSP) unit embedded into, or external to, the musical instrument;
- λ_{outbuf}, due to the acquisition of the digital signal in the audio buffer of the DSP unit;
- λ_{trans}, due to the algorithms that create the haptic signal based on the audio signal in input.

The second component,

$$\lambda_{UWB} = \lambda_{pack} + \lambda_{link} + \lambda_{unpack} \tag{3}$$

is the overall communication delay necessary to the signal transmission and reception, composed by the delays

- λ_{pack}, due to the acquisition and packetization of the digital signal at the UWB transmitter;
- λ_{link}, the propagation time across the UWB wireless link;
- λ_{unpack}, due to the depacketization of the received signal via the UWB receiver.

Finally, the last delay component

$$\lambda_{dev} = \lambda_{inbuf} + \lambda_{proc} + \lambda_{DAC} + \lambda_{act} \tag{4}$$

occurs at the haptic device for the processing and actuation necessary to its enactment, inducing the delays

- λ_{inbuf}, due to the acquisition of the depacketized digital haptic signal into the buffer of the DSP unit of the device;
- λ_{proc}, due to the algorithms that process the incoming haptic signal, e.g., to route it to the actuators of an haptic devices using different amplitudes;

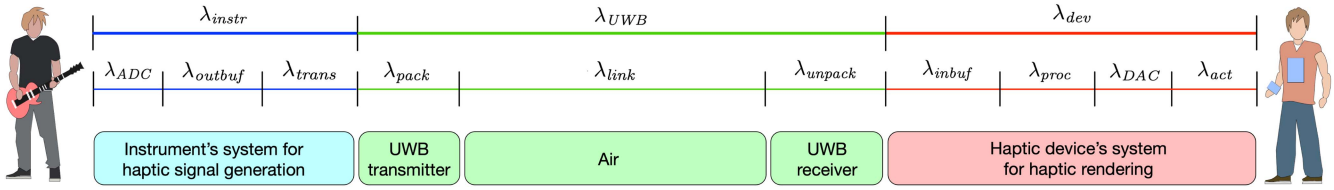

Fig. 2. Schematic representation of the components contributing to the overall latency in a UWB-based musical haptics system.

- λ_{DAC}, due to the delivery of the received signal to the actuators via a digital to analog (DAC) converter;
- λ_{act}, due to the mechanical activation of the actuator.

Reliability: In principle, the factors contributing to the reliability R parallel those identified for latency:

$$R = \rho_{instr} \cdot \rho_{UWB} \cdot \rho_{dev} \tag{5}$$

However, in practice, the components on the instrument and haptic device (ADC, DSP, actuator) can be considered perfectly reliable w.r.t. the possibility of packet loss. Therefore, the latter is the only contributor to reliability, $R = \rho_{UWB}$, and can be defined in terms of the packet reception rate (PRR), i.e., the ratio of the number packets received over the total number of packets sent across the link.

Requirements: Based on the modeling above, we can formalize the requirements of our UWB-based musical haptic system.

From a latency standpoint, the goal is to keep the difference in absolute value between the latency Λ_a of the audio signal through air and the overall haptic latency Λ_h below a given perceptual threshold τ_Λ:

$$|\Lambda_a - \Lambda_h| < \tau_\Lambda \tag{6}$$

Another related desirable property, greatly simplifying design and implementation, is for the *variation* (jitter) of latency Λ_h to be very low, i.e., yielding a nearly constant latency.

As for reliability, the goal is to guarantee perfect communication of the haptic signal with zero packet losses, yielding $R = 1$. Nevertheless, this is not generally possible. Moreover, what matters from a user perception perspective is not the packet loss per se, rather the impact it bears on a proper reconstruction of the haptic signal. Therefore, the goal is to keep the difference between the original haptic signal $\mathcal{H}_{instr}$ generated at the instrument and the haptic signal $\mathcal{H}_{dev}$ experienced by the audience member via the haptic device below a given perceptual threshold τ_R:

$$\mathcal{H}_{instr} - \mathcal{H}_{dev} < \tau_R \tag{7}$$

Unfortunately, defining τ_Λ and τ_R is a challenge per se. We have already discussed (§II) how an agreed-upon value for τ_Λ is still missing from the literature, although clearly the lower the better. Similar considerations hold for τ_R [38] that, in our context, may be also affected by crossmodal perceptual mechanisms due to the simultaneous presence of the concurrent audio stream.

Given this state of affairs, we resort to an experimental approach. We first characterize the latency and reliability of the raw UWB communication link when used to transmit audio signals (§IV). Next, we characterize the remaining components of latency based on real-world measurements from full-fledged prototypes for musical haptics (§V). Finally, we ascertain the impact on perception of latency and reliability by means of a user study (§VI), effectively providing an empirical answer to the above formal requirements.

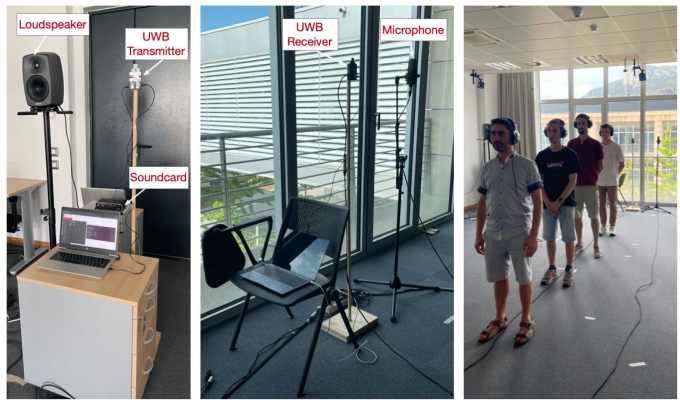

Fig. 3. Transmitter (left) and receiver (middle) apparati; NLOS setup (right).

IV. Characterizing the UWB Link

We now analyze the UWB communication link used to transmit the audio signal triggering the haptic response at the receiver. We first illustrate the experimental setup (§IV-A) and the procedure we used for measurement (§IV-B), significantly complicated by the need to accurately ascertain both sound and communication characteristics in an integrated fashion. We then present separately the results we gathered when considering unicast (§IV-C) and broadcast (§IV-D) communication.

A. Experimental Setup

The measurements were conducted in an acoustically insulated laboratory at our premises.

Device and people placement: The source and receiver of the haptic signal are positioned at a height of 1.5 m. Three setups are investigated, mimicking different real-life situations. In the first one (LOS5) the source and receiver are in line-of-sight (LOS) at a distance of 5 m, mimicking e.g., receivers placed in the first row of seats at during a concert. The second (LOS10) is a variation with a longer distance of 10 m. Finally, a third condition (NLOS) investigates, in the same setup as LOS10, the impact of 4 people placed between source and receiver, the first being placed at 5 m. This setup yields non-line-of-sight conditions common, e.g., when receivers are seated one or more rows behind the first. The subjects were asked to stand still and silent during the measurement process (Fig. 3).

UWB transceiver: Source and receiver host evaluation boards equipped with the Spark Microsystems SR1020 UWB transceiver [11]. This radio supports the 6–9.25 GHz frequency band and dynamically shapes the output spectrum to obtain a $-41.3\,\mathrm{dbm/MHz}$ output power, ensuring compliance with most regulatory bodies worldwide. The receiver is non-coherent, which reduces system complexity and yields

order-of-magnitude improvement in terms of per-bit energy consumption over Bluetooth Low Energy (BLE), commonly used in smart wearables. The communication range in the configuration we used is up to 50 m, based on the manufacturer information and confirmed by our experiments.

Firmware for audio streaming via UWB: Interestingly, unlike other UWB transceivers, the SR1020 privileges low-latency, high-rate streaming applications rather than accurate ranging. As a consequence, as part of the developer software development kit (SDK) the manufacturer provides a wireless and audio core, which we used to develop the firmware hosted on our boards. The SDK greatly simplifies development, in that it takes as input *directly* an audio stream and handles automatically its delivery to the receiver over the UWB channel, without no additional intervention required by the programmer. Specifically, the SDK includes a neighbor discovery protocol to establish a connection between transmitter and receiver, and a TDMA schedule that accommodates repeating epochs, each further divided in timeslots (200 μs in our application) where the actual packet transmission of audio data occurs.

Each packet contains an 80 B audio payload; 78 B are dedicated to encapsulating the audio data, while the remaining 2 B are allocated for the packet. The SR1020 supports on-off keying (OOK) and pulse-position modulation (PPM); we used PPM as it provides a higher data rate. However, little information on the final data rate is provided by the manufacturer, therefore we estimated it with dedicated experiments. We observed values in the range 8.38–9.58 Mbps with the above timeslot duration and the 80 B packet size, automatically augmented by the transceiver with a few bits necessary to forward error correction (FEC), increasing reliability.

Audio devices: A loudspeaker (8030 CP by Genelec) is placed on a stand and aligned both horizontally and vertically with the UWB transmitter; a condenser microphone (C414 by AKG) is similarly aligned with the UWB receiver. Audio cables connect the inputs of both UWB transmitter and loudspeaker with the outputs of a sound card (UMC202HD by Behringer). Similarly, other audio cables connect the output of both UWB receiver and microphone with the inputs of the soundcard.

B. Measurement Procedure

Latency: A laptop connected to the soundcard measures latency via a Pure Data (Pd) application we developed to create the audio and haptic signals to be transmitted and record the corresponding received signals. For the audio signal to be transmitted via the loudspeaker, we created a. wav file lasting 10 minutes and containing a short sound with a sharp attack (i.e., the kick of a drum) repeated every 2 s. Across two repetitions, this led to a total of 600 sounds used as measurement points for latency. As for the haptic signal to be delivered by the UWB transmitter, we applied a bandpass filter to the audio signal to adapt it to the vibrotactile perception range and the range supported by the actuators involved in our prototypes (§V), i.e., $[30, 1.000]$ Hz. Pd was configured to work with a sampling rate of 48 kHz and recorded three tracks: *i)* the audio signal generated, *ii)* the signal from the microphone, and *iii)* the signal from the UWB receiver.

This setup is able to measure the latency λ_{UWB} of the communication link plus the delay introduced by the bandpass filter, λ_{trans}; this is however negligible ($\ll 1$ ms) w.r.t. λ_{UWB} and ignored hereafter. The main goal of latency measurements in this section is to compute the difference, measured at the receiver in the three conditions outlined earlier (§IV-A), between the latency Λ_a of the sound generated at the source and the latency of the haptic signal experienced at the receiver *before* being input to the haptic device that, once we ignore the contribution of λ_{trans}, coincides with the latency λ_{UWB} to traverse the UWB link (Fig. 2). In practice, our setup measures the quantity $\Lambda_a - \lambda_{UWB}$ which *directly* informs us whether the latency induced by UWB is compatible with our envisioned system, i.e., whether this difference is *positive* and therefore the sound arrives *after* the tactile signal.

Finally, to measure the audio and the haptic signals, we created in Python an onset detector based on a simple thresholding scheme and computed the difference between the corresponding onsets in the three recorded tracks.

Reliability: Other two laptops, connected to the UWB transmitter and receiver instrumented with our firmware, measure reliability in terms of the common communication-level PRR metric (Section III). Nevertheless, this metric is significantly trickier to ascertain in our context: paradoxically, the reason lies in the unique capability of the SR1020 boards to directly take raw audio input and send it to a receiver over the UWB link. This is a tremendous asset when developing applications; however, the SDK provided by the manufacturer treats this functionality as a black box. As a consequence, although a log of the packets *received* can be obtained, there is *no access* to the information necessary to compute the PRR. We cannot access the packet sequence number, which would allow us to determine directly packet losses by identifying gaps in the sequence. Similarly, we cannot access the timestamps associated to packet transmissions and receptions; a tight, per-packet synchronization is impossible to achieve, which is key to align these operations at sender and receiver and identify failed transmissions.

We estimate the *raw* link reliability of the UWB link as follows. First, we observe that the SR1020 receiver, when configured in unicast mode, *automatically* acknowledges each received packet. Unicast communication, in which a source transmits the signal to a single receiver, is not our target application case; we want to support multiple receivers at once, which is naturally and efficiently supported via broadcast communication and would not scale with unicast. However, the presence of acknowledgments enables measurements that are precluded in broadcast. Specifically, we exploit the automatic acknowledgments in unicast as a means for coarse synchronization between the sender and receiver logs in the case of a successful reception. We cannot, however, exploit acknowledgments directly as a measure of reliability for the audio data, which are carried by a different type of packet.

To mitigate this aspect, we consider only the packet traces following the first reporting of an acknowledgment by the sender, effectively marking the beginning of the experiment. Traces are further split in 1-second windows; in each 20-minute experiment, 1200 windows are therefore logged, each containing a few thousands of packet transmissions, in their 200 μs slot, plus the necessary printing operations. Inside each window, the sender (receiver) keeps track of the packets sent (received) until the 1 s interval expires and the process restarts for a new

TABLE I
UNICAST COMMUNICATION: LATENCY AND RELIABILITY

Condition	ρ_{UWB} PRR (%)	$\Lambda_a - \lambda_{UWB}$ (ms)	Λ_a (ms)	λ_{UWB} (ms)
LOS5	92.69	4.34 ± 0.024	14.57	10.23
LOS10	99.81	18.85 ± 0.073	29.15	10.30
NLOS	99.89	18.84 ± 0.066	29.15	10.31

TABLE II
BROADCAST COMMUNICATION: LATENCY

Condition	$\Lambda_a - \lambda_{UWB}$ (ms)	Λ_a (ms)	λ_{UWB} (ms)
LOS5	12.73 ± 0.069	14.57	1.84
LOS10	27.11 ± 0.063	29.15	2.04
NLOS	27.12 ± 0.065	29.15	2.04

window. The final PRR is then simply the aggregation of the one recorded in each window. Crucially, although the SR1020 uses the acknowledgment to retransmit a missed packet, we are entirely oblivious to whether a packet is a retransmission or not; the PRR is is simply computed over *all* the packets transmitted. For the same reason, we cannot measure the actual unicast packet *delivery* rate (PDR) at the application, which is nonetheless irrelevant when estimating the reliability of broadcast communication—our ultimate goal.

Indeed, the SR1020 provides no automatic acknowledgment in broadcast, as this would require knowing in advance all potential recipients. This means that we cannot use the technique above; given the limitations of the SDK, it is therefore essentially impossible to accurately determine the PRR. On the other hand, the unicast PRR serves as a good proxy given that, at the PHY level, the communication mechanism is the same as broadcast. As a consequence of these considerations, in our results (§IV-D) we do not report the reliability of broadcast communication as this is accurately approximated by the one we measured in unicast with the procedure above.

C. Unicast Communication

Our main goal here is to determine the reliability of the UWB link in terms of PRR, as this serves as an accurate proxy for the reliability when using broadcast communication, our target scenario. Nevertheless, by also comparing the latency of unicast vs. broadcast we gain a better understanding of the inner working of the SR1020 UWB transceiver we chose.

Results: Table I shows the reliability and latency we measured. Concerning the former, the PRR value reported is the average of both *full* recordings, i.e., is computed across a total of *more than 5 million packets*, yielding a strong statistical significance. We observe how reliability is very good, except for the shorter, 5-meter distance in LOS5. This is actually surprising, especially considering that popular UWB transceivers like the Qorvo DW1000 [12] have near-perfect reliability in all the conditions we tested, including LOS5. On the other hand, our test have been carried out in an environment with many reflections from the close walls, causing collisions and therefore packet loss; we verified that we do not observe the same drop in reliability at 5 m when testing in a different, less challenging environment. In any case, this aspect is further motivation, if at all needed, that a user study investigating whether this impacts user perception is key (§VI).

As for latency, we report the mean and standard deviation of our measurements of $\Lambda_a - \lambda_{UWB}$ in each condition. The first, important observation is that this latter quantity *is always positive*, therefore confirming that our envisioned application is indeed possible—at least from a communication standpoint, given that here we do not consider the contribution of the other Λ_h components in (1) we measure later on actual prototypes (§V). Moreover, latency is also stable, with a standard deviation of only tens of microseconds. From this measurement, and knowledge of Λ_a based on the speed of sound in air for each distance, we can indirectly derive the latency λ_{UWB}, which is nearly constant at slightly more than 10 ms. This should not come as a surprise, as the only difference between LOS5 and LOS10 is an increase in distance by 5 m, which translates in a mere 16.7 ns. On the other hand, the latency difference $\Lambda_a - \lambda_{UWB}$ is quite small for LOS5, possibly making it challenging to realign the sound and haptic signals once the latency λ_{instr} induced by the instrument is factored in (§V). We verify now whether this is the case also with broadcast communication.

D. Broadcast Communication

Our focus is on live performances, e.g., a concert, in which the audio-haptic signal generated by the musical instrument and the haptic transmitter would reach the whole audience, not just a single person. This can be accomplished via broadcast communication, supported by the audio core of the SR1020 device and analyzed here.

Results: Table II shows the latency we measured. We do not report the PRR due to the impossibility to measure it (§IV-B); the values for unicast (Table I) are nonetheless a good proxy.

A comparison with Table I shows that broadcast exhibits significantly lower latencies w.r.t. unicast. Indeed, the SR1020 SDK in unicast alternates packet transmission with the reception of its acknowledgment, which significantly reduces throughput and increases latency, effectively delaying the packets containing audio data. Broadcast communication, instead, does not rely on acknowledgments and can therefore explore the full potential of UWB transmissions.

Consequently, the difference $\Lambda_a - \lambda_{UWB}$ significantly *increases* in broadcast, providing a greater margin for compensation in our target scenario. For LOS5, the value of this difference more than 3x w.r.t. unicast; for the other two cases, we observe more than a 50% increase. The values of λ_{UWB}, computed from the measured difference and the same values of Λ_a for unicast, show that the reason is that the corresponding latency is *one order of magnitude* smaller in broadcast. Finally, the other trends concerning the different conditions are essentially confirmed.

V. PROOF-OF-CONCEPT PROTOTYPES

The characterization we just provided focuses on the communication link. Nevertheless, as discussed earlier (Section III), latency is affected also by the delays induced by the generation of the haptic signal at the source and of the actuation at the receiver. These need to be accounted for to verify the practical

TABLE III
CHARACTERISTICS OF THE PROOF-OF-CONCEPT PROTOTYPES

Prototype	Sound source		Haptic receiver	
	acoustic (Fig. 5a)	electric (Fig. 5b)	armband (Fig. 6a)	vest (Fig. 6b)
1	✓		✓	
2		✓	✓	
3	✓			✓
4		✓		✓

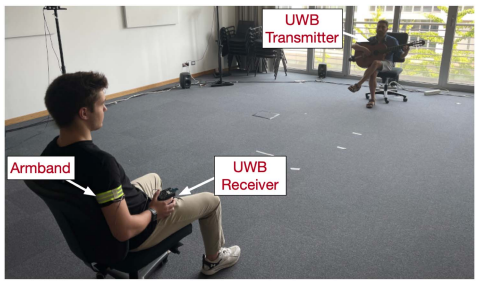

Fig. 4. Prototype 1 in action: UWB-enabled haptic interaction between a performer and an audience member.

feasibility of our envisioned system and, if so, to determine the time compensation necessary to realign the stimulus of the haptic device with the sound propagating through air.

We investigate this aspect via four proof-of-concept prototypes. At the performer, they augment a common instrument—an acoustic or electric guitar—with the ability to generate and transmit the haptic signal. At the listener, they provide a wearable device—an armband or a vest—capable of receiving the haptic signal and perform the corresponding haptic actuation. Table III shows at-a-glance how we combined the two components at the performer and the two at the listener (§V-A) in the four prototypes. In all cases, the generation of the haptic signal at the instrument is performed digitally. As an example, Fig. 4 shows Prototype 1 in action, where a performer playing the smart acoustic guitar in a multi-sensory interaction with an audience member wearing musical haptic armband.

We detail the characteristics of these components, followed by a rationale of their choice w.r.t. latency and a quantitative assessment of their impact on it, based on real-world measurements enabled by our prototypes.

A. Smart Musical Instruments and Haptic Wearables

All prototypes are equipped with the SR1020 UWB board (§IV-A) both at the instrument (source) and wearable device (receiver) configured to operate with broadcast communication, as this reflects our target application scenario. For all components, both smart musical instruments (Fig. 5) and haptic wearables (Fig. 6), Pd was configured to run at a sampling rate of 48 kHz and audio buffer size of 64 samples. The Pd task is assigned highest priority by the Bela board operating system, when used.

Smart acoustic guitar: This instrument (Fig. 5(a)) is composed by: *i)* a conventional classical guitar enhanced with a contact microphone, *ii)* a small loudspeaker, *iii)* the Bela-mini low-latency audio processing board, including a hard real-time Linux kernel [39], *iv)* an embedded battery, *v)* a short 3.5 mm jack mono audio cable connecting the first output channel of the Bela board to the embedded loudspeaker, *vi)* a short 3.5 mm jack mono audio cable connecting the second output channel of the Bela board to the UWB transmitter. The haptic generation engine is coded in Pd and simply provides a band-pass filtered

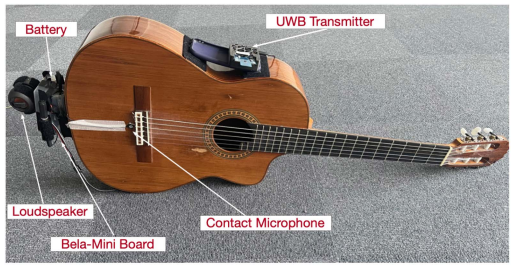

(a) Smart acoustic guitar (Prototype 1 and 3).

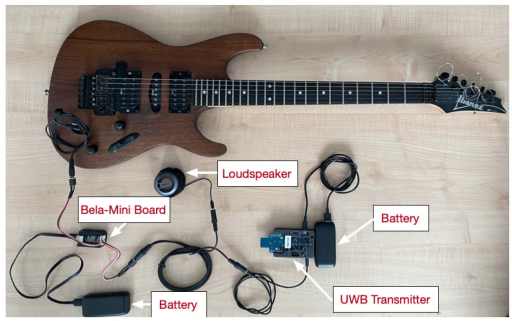

(b) Smart electric guitar (Prototype 2 and 4).

Fig. 5. Smart musical instruments.

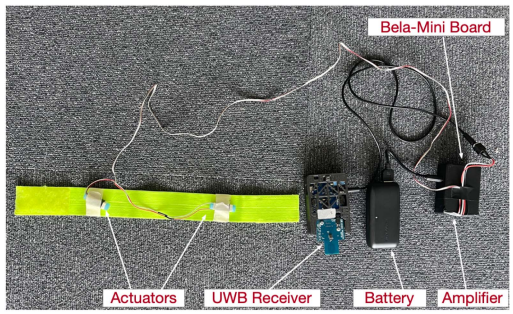

(a) Musical haptic armband (Prototype 1 and 2).

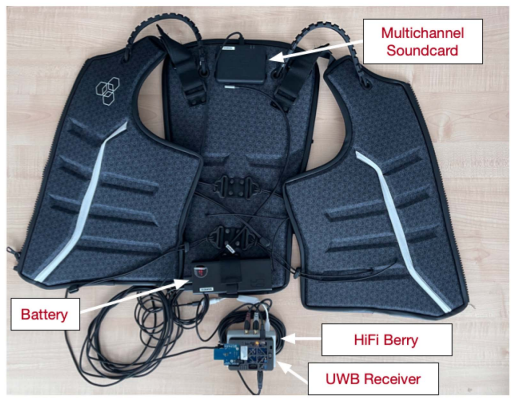

(b) Musical haptic vest (Prototype 3 and 4).

Fig. 6. Musical haptic wearables.

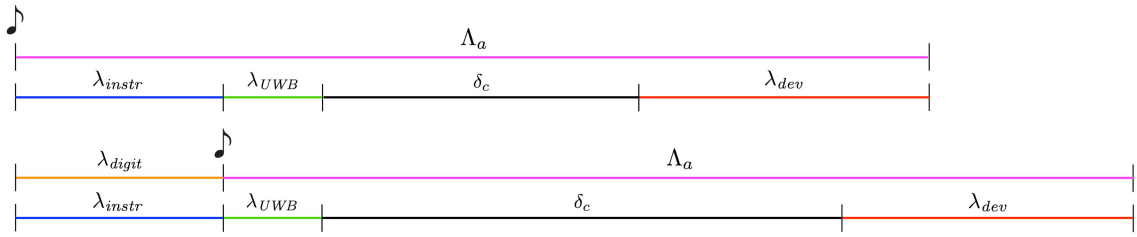

Fig. 7. Latency components in acoustic (top, Prototype 1 and 3) vs. electric (bottom, Prototype 2 and 4) smart guitars. The delay δ_c is a configuration parameter we set (Table IV) to time synchronize the audio and haptics signals at the receiver. Note the large value of δ_c in the case of the electric guitar.

version of the input audio signal in the tactile and actuators range (i.e., [30,1000] Hz).

Smart electric guitar: This instrument (Fig. 5(b)) contains the same components of its acoustic counterpart, except for the first one; instead of acquiring sound from the contact microphone it uses directly the piezoelectric microphone system of the electric guitar. Sound is then delivered by a loudspeaker connected to the Bela board. The audio engine running on the latter is coded in Pd and includes a module processing strings sounds with a chorus effect in chain with a delay and a reverb effect. The haptic generation engine, also coded in Pd, provides a band-pass filtered version, in the tactile and actuators range [30,1000] Hz, of the digital audio signal.

Musical haptic armband: This wearable (Fig. 6(a)) is an elastic armband encompassing *i)* two small voice-coil actuators, Tactuator MM3C by Tactile Labs [40], *ii)* a class D stereo amplifier, MAX9744 by Adafruit, *iii)* a Bela-mini board, *iv)* a short 3.5 mm jack audio cable connecting the UWB receiver to the input of the Bela-mini, *v)* a short 3.5 mm jack audio cable connecting the output of the Bela-mini to the input of the amplifier, *vi)* a short 3.5 mm stereo jack audio cable connecting the amplifier to the two actuators, modified to duplicate on both channels the mono signal arriving from the UWB receiver, *vii)* a battery. The haptic rendering engine is coded in Pd and simply routes the incoming signal to the two motors, with the compensation delay necessary to realign the haptic signal with sound (§V-C).

Musical haptic vest: We use the commercially available Skinetic vest by Aktronica (Fig. 6(b)). This includes 20 small voice-coil actuators driven by a dedicated multichannel sound card. This system is connected via USB to a Raspberry PI 4 (RPI) enhanced with the Hi-Fi Berry audio shield. The audio input of the shield is connected to the audio output of the UWB receiver. The RPI runs a Linux distribution with the PREEMPT-RT hard real-time kernel [41]. The haptic rendering engine, coded in Pd, focuses on routing the incoming signal to the 20 motors with the compensation delay (§V-C).

B. Rationale of Choices

The features of the components above are chosen to cover different points of the design space, yielding a different impact on the overall haptic latency Λ_h. In particular, the choice between acoustic and electric guitar bears an impact on latency constraints, more stringent in the former case.

Fig. 7 illustrates why, starting from the instant in which a string is struck. In an acoustic guitar, the sound is directly generated by the instrument and reaches the ears of the listener after a delay Λ_a. An electric guitar, instead, incurs a delay λ_{digit} before sound is emitted, due to the need to digitize sound and send it to the loudspeakers. In our envisioned system, this is actually an advantage, as we send to the receiver this digital audio stream as the haptic signal, after a filtering step with negligible latency (§IV-B); therefore, no additional latency $\lambda_{digit} \approx \lambda_{instr}$ is incurred w.r.t. sound propagation, as the digitization process occurs *before* it. In contrast, to produce a similar haptic signal in the acoustic guitar, *digitization inevitably occurs in parallel with sound propagation*. This reduces the margin available for compensating delays via a system-induced delay δ_c and may even prevent it altogether in case the latency λ_{instr}, once combined with the other haptic latency components in (1), yields $\Lambda_h > \Lambda_a$.

TABLE IV
PROOF-OF-CONCEPT PROTOTYPES: LATENCY AND COMPENSATION

Prototype	Condition	Λ_h λ_{instr} (ms)	Λ_h λ_{UWB} (ms)	Λ_h λ_{dev} (ms)	δ_c (ms)	Λ_a (ms)
1	LOS5	3.66	1.84	7.66	1.41	14.57
	LOS10		2.04		15.79	29.15
	NLOS		2.04		15.80	29.15
2	LOS5	—	1.84	7.66	5.07	14.57
	LOS10		2.04		19.45	29.15
	NLOS		2.04		19.46	29.15
3	LOS5	3.66	1.84	18.86	**−9.79**	14.57
	LOS10		2.04		3.59	29.15
	NLOS		2.04		3.60	29.15
4	LOS5	—	1.84	18.86	**−6.13**	14.57
	LOS10		2.04		7.25	29.15
	NLOS		2.04		7.26	29.15

Among these components, we have already assessed λ_{UWB} (§IV-D); as for λ_{dev}, this is determined by the haptic wearable. Our choice of an armband and a vest allows us to investigate very different conditions; the vest introduces a latency λ_{dev} 2.5x higher w.r.t. the armband. Again, this is crucial for compensation and the overall feasibility of our approach, investigated quantitatively next.

C. Impact on Latency and Overall System Feasibility

We exploit our proof-of-concept prototypes to derive real-world measurements and gain a concrete understanding of the tradeoffs at stake. Table IV shows our results in the three scenarios under consideration (§IV-A). The latency introduced by the smart guitars and the haptic wearables, respectively λ_{instr} and λ_{dev}, are measured directly at the corresponding components

in our prototypes. As for the UWB latency, λ_{UWB}, we cannot directly measure it with the same setup. For this reason, we report the same values we derived earlier for broadcast (§IV-D) as the experimental conditions are exactly the same from the point of view of communication.

We can distill a few key considerations. First and foremost, our measurements confirm that the system we envision is feasible, from the point of view of latency. Indeed, in all 12 cases but 2, the haptic signal arrives *before* the sound; therefore, synchronicity between the two can be easily re-established at the receiver by introducing a compensating delay at the source, whose magnitude $\delta_c = \Lambda_a - \Lambda_h$ we also report.

At the same time, the two cases where the haptic signal arrives *after* the sound are informative of the conditions that push our system to the limit. First of all, these values occur on the shorter distance of 5 m (LOS5), where sound propagation is twice as fast w.r.t. the other two conditions and therefore leaves much less margin to Λ_h, whose components are instead essentially constant w.r.t. distance. Among these, the latency introduced by the haptic vest in Prototype 3 and 4 is by far the largest, enough to push δ_c to a negative value; indeed, when using the haptic armband in Prototype 1 and 2 in the same LOS5 condition, δ_c remains positive in both cases. Finally, Prototype 3 experiences a larger anticipation of sound; this is due to the use of smart acoustic guitar in which, unlike the smart electric guitar of Prototype 4, the digital signal used for haptics cannot be reused and needs to be generated, causing an additional latency λ_{instr} (Fig. 7).

Clearly, these are just proof-of-concept prototypes and could be further optimized. Nevertheless, when testing our prototypes we noticed that sound and haptics are always perceived in synchrony regardless of the combination of prototype components and scenario. We investigate this aspect more rigorously in the user study described next.

VI. User Study: From Latency and Reliability to Perception

The goal of this section is to determine whether and how the worst-case *system-level* objective metrics of negative offset δ_c (§V) and of reliability (§IV) we observed have an impact on the subjective *user perception*. Indeed, in addition to latency, reliability is another potentially detrimental strong contributor, in that packet losses occurring on the audio stream used to generate the haptic signal may corrupt the latter.

We investigate these aspects via a user study where, compared to the results we showed thus far, we perform experiments with *many subjects* and in a *repeatable* way. Ideally, this could be accomplished with a musician performing live in front of an audience. Nevertheless, this is inherently limited by the availability of enough SR1020 boards and, more importantly, by the fact that live performances are intrinsically at odds with accurate repeatability. Therefore, we exploit an experimental setup (§VI-A) in which we control precisely the latency and reliability of the audio-haptic signals, based on our previous experiments. These are delivered via wire to the earphones of the subjects, who are then required to rate their experience across the three perception dimensions providing the basis for our analysis (§VI-B).

This is the last step of our evaluation and the one directly affecting the practical viability of the proposed system. However, the relevance of the results we report go beyond our direct needs; for instance, the effect of network reliability and imperfect packet loss reconstruction on tactile perception of musical haptic signals has been scarcely investigated thus far.

A. Experimental Setup

The methodology is approved by the local ethics committee in accordance with the relevant ethical standards of the Declaration of Helsinki and in compliance with the EU GDPR.

Subjects: The experiment involved 17 subjects, 11 men and 6 women, with age between 21 and 43 years (mean=26.41, stddev=5.02). All were healthy adults and did not report any hearing or somatosensory problems. Each subject spent on average 30 minutes to complete the experiment.

Stimuli: The stimuli consist of 4 music pieces provided to the subject at both auditory and haptic level, under four conditions detailed next. The music pieces (Table V) are selected to cover a large variety of features of the music signal (e.g., tempo, harmony, instrumentation), as well as genre and conveyed emotion. Each music piece is trimmed to last 30 seconds; the resulting stimulus is repeated twice in randomized order, for a total of 32 trials. The amplitude of the auditory and haptic stimuli are manually adjusted to be coherent between each other, based on the results in [42]. These amplitudes were determined via a pilot test with 3 subjects not involved in the experiment. For each piece, they were asked to agree on the amplitude values for both audio and haptic signals yielding a comfortable (i.e., neither too low nor too high) and musically meaningful audio-haptic experience.

Conditions: We designed the conditions in which the stimuli are applied to properly investigate the impact of latency and reliability in a controlled, repeatable setup. Specifically, we compare an ideal setup with zero latency and perfect reliability with realistic cases artificially reproduced in our setup.

Concerning latency, we focus on the difference $\Lambda_h - \Lambda_a$, crucial to perception. As a baseline, we consider the ideal case in which this difference is zero, realized by *directly* feeding the subjects with perfectly synchronized sound and haptics delivered via wires. At the other extreme, we repeated a similar procedure with $\Lambda_h - \Lambda_a = 10$ ms, approximating the worst-case offset we observed with our prototypes (Table III).

As for reliability, the wire-based procedure above already yields the ideal case. To reproduce real, worst-case conditions we play the music pieces over the broadcast channel in the LOS5 scenario, whose reliability we indirectly measured at $\rho_{UWB} = 92.69\%$ (§IV), and used the resulting haptic signal, containing the imperfections due to packet losses, as input to the wearable device worn by the subjects. The impact of packet losses on the reconstruction of the audio signal at the receiver is automatically mitigated by the SR1020 SDK wireless audio core via a packet loss concealment (PLC) technique. Nevertheless, this generates spikes in the signal that, in turn, result in glitches perceivable at auditory level. Interestingly, these occur mostly at high frequencies, i.e., outside of the [30,1000] Hz band relevant to tactile perception and targeted by the bandpass filter utilized at the receiver to extract the haptic signal. As a consequence, the latter removes a significant portion of the glitches caused by packet losses.

TABLE V
USER STUDY: MUSIC PIECES AS STIMULI

Piece	Title	Author	Genre	Emotion
M1	The Vengeful One	Disturbed	heavy metal	aggressiveness
M2	Ob-La-Di, Ob-La-Da	The Beatles	pop-ska	happiness
M3	Orfeo and Euridice	Christoph Willibald Gluck	classical	relaxation
M4	I Can't Make You Love Me	Bonnie Raitt	pop	sadness

TABLE VI
USER STUDY: CONDITIONS

Condition	$\Lambda_h - \Lambda_a$	ρ_{UWB}
IDEAL	0	100%
DELAY	−10 ms	100%
LOSS	0	92.69%
REALISTIC	−10 ms	92.69%

Finally, for a fair comparison, the resulting files were normalized to match the maximum amplitude of the haptic signal with perfect reliability. Table VI shows the resulting combinations. IDEAL and REALISTIC represent the two extremes, with the latter essentially reproducing the worst-case scenario of Prototype 3. DELAY and LOSS enable us to investigate separately the impact of latency and reliability.

Procedure: The subjects were asked to wear a pair of headphones (Beyerdynamic DT-770 Pro 80 Ohm) and the wearable haptic vest (Fig. 6(b)). At the outset, they were briefed about the experiment and asked to sign a consent form. Further, they familiarized with the system via a session in which they provided audio-haptic stimulus in IDEAL condition but with a different music piece (*"Don't stop me now"* by Queen).

Subsequently, they were provided with a graphical interface allowing them to trigger the trials comprised in the experiment. After each trial, the subjects were asked to assign a value on an 11-point Likert scale to the following questions:

- *To what extent the audio and haptic signals were synchronous?* [very asynchronous, very synchronous].
- *How was the quality of the vibration you felt?* [low quality, high quality].
- *How pleasant was the overall audio-haptic experience?* [very unpleasant, very pleasant]

B. Results

Fig. 8 shows mean and standard error of the subject ratings of audio-haptic synchronicity, vibration quality, and pleasantness across conditions and music pieces.

An ANOVA was performed on generalized linear mixed-effects models for each of the three dimensions above. Each model had the rating of the three dimensions above, condition (Table VI) and piece (Table V) as fixed factors, and the subject as a random one. For each model, the assumption on the normality of residuals was verified. Post-hoc tests were performed on the fitted model using pairwise comparisons adjusted with the Tukey correction.

Concerning synchronicity, a significant main effect was found only for the piece factor ($F(3, 521) = 18.69$, $p < 0.001$) and not for conditions. The post-hoc test showed that piece M3 was rated with lower synchronicity than M1 ($p < 0.001$), M2 ($p < 0.01$) and M4 ($p < 0.001$). Therefore, on average, the participants did not perceive any asynchrony between the auditory and haptic stimuli across the four conditions. In contrast, the perception of synchrony was found to be dependent on the type of music provided.

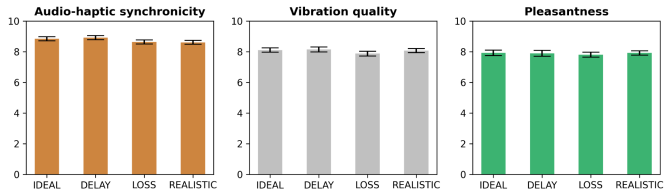

(a) With respect to the conditions in Table VI.

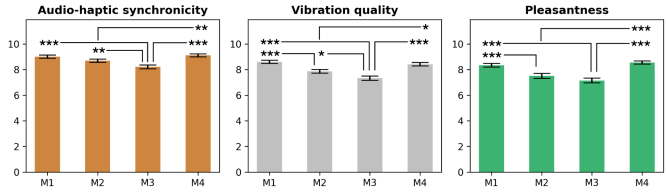

(b) With respect to the music pieces in Table V.

Fig. 8. Mean and standard error of the subjects' ratings of audio-haptic synchronicity, vibration quality, and pleasantness for the conditions in. Legend: $* = p < 0.05$, $** = p < 0.01$, $*** = p < 0.001$.

Regarding vibration quality, a significant main effect was again found only for the piece factor ($F(3, 521) = 28.37$, $p < 0.001$). The post-hoc test showed that *i)* M3 was rated with low vibration quality than M1 ($p < 0.001$), M2 ($p < 0.05$), and M4 ($p < 0.001$), and *ii)* M2 was rated with lower vibration quality than M1 and M4 (both $p < 0.05$). These results indicate that, on average, participants did not perceive any difference among the four conditions w.r.t. the quality of haptic stimulation, which was found dependent only on the musical piece.

As for pleasantness, results are in line with the above, with a significant main effect found only for the piece factor ($F(3, 521) = 28.58$, $p < 0.001$). The post-hoc test evidenced how *i)* M3 was rated to be less pleasant, at the audio-haptic level, than M1 and M4 (both $p < 0.001$), and *ii)* M2 was rated less pleasant than M1 ($p < 0.001$) and M4 ($p < 0.001$). Hence, consistently with the results for synchronicity and vibration quality, the pleasantness of the overall audio-haptic experience did not significantly vary among the four conditions, and was instead dependent on the type of music provided.

For all three dimensions investigated, we did not find any statistically significant interaction effect between the piece and condition factors.

We applied Pearson's tests to search for possible correlations among the three types of rankings. Significant correlations ($p < 0.001$) of medium strength were identified for synchronicity-vibration and for synchronicity-pleasantness

(both $r \approx 0.4$). A significant correlation ($p < 0.001$) of medium-high strength was identified for vibration-pleasantness ($r = 0.72$). These results suggest that the three dimensions investigated are not fully independent.

VII. Discussion and Outlook

Our measurements show that the haptic signal transmitted in broadcast over the UWB link is almost always faster than propagation of sound in air, and is also very stable, with only minimal variations. This enables the use of a compensation delay $\delta_c = \Lambda_a - \Lambda_h$ realigning the two sensory signals. As for the few cases where the haptic signal arrives *after* sound and compensation is not possible, the results of our user study (§VI) show that the -10 ms offset in our prototypes is not sufficient to perceive a synchronization mismatch between the two sensory modalities. This is in agreement with the findings of previous studies focusing on haptic signals delivered to the hands or the whole body [29], [30], [31], [32], [33], [34]. The value of δ_c depends on the time required for haptic signal generation and rendering at the specific instruments and wearable haptic devices used, respectively, for which we offered concrete examples. On the other hand, the experimental setup from which we derive our results considers stimuli from a haptic device that is *not* optimized for latency, as it involves the use of two soundcards. In a real product, the latency introduced by the haptic device would be much lower, likely bringing δ_c to a positive value and enabling compensation.

Another crucial dimension is reliability, whose impact we characterized both from both the system and user perception perspectives. Interestingly, our user study shows that even the worst-case, relatively high rate (7.31%) of packet loss we observed did not bear an impact on user perception, a result we argue is the interplay of two dimensions. On one hand, this can be attributed to the relatively poor tactile acuity of the upper body along with the likely presence of dominance mechanisms of the auditory system over the somatosensory one. On the other hand, the packet losses affected the haptic signal but without causing excessive deterioration. Indeed, while we observed imperfections in the received signal due to packet losses, the majority of these artifacts were easily removed via a bandpass filter, highlighting the importance of this on-board ability on musical haptic wearables.

Overall, the perception of latency and vibrations quality were found to be affected by the *type of music* played, rather than system-level quantities like latency and reliability. This finding suggests that audio-haptic synchronicity and vibration quality may be regulated by features specific of the musical signal, calling for research on this rather unexplored topic.

Of course, our study has limitations, essentially due to the desire to accurately measure the many interconnected aspects of our system in a *repeatable* way and the inherent logistic complexity this entails. As a consequence, our measurements were conducted in controlled conditions rather than in a real performance with a crowded audience. Moreover, constrained by the availability of an appropriate indoor space where to deploy our setup, we considered only distances of 5 m and 10 m. However, we argue that these distances are not only representative of real-world live performances in small venues, but also a *stress case* of sorts, at least from the point of view of latency. Indeed, over longer distances, e.g., in large concert halls, the haptic latency Λ_h remains nearly constant while the sound Λ_a increases significantly (Table II), providing even ampler margins for introducing a compensation delay δ_c.

In this respect, we did not exploit the capability commonly associated to UWB radios, i.e., the ability to accurately estimate distance. This could enable an *automatic, dynamic, and individual* determination of the compensation delay δ_c for each receiver, by computing the sound latency Λ_a as a function of the current distance and subtracting known haptic latency Λ_h, greatly simplifying system configuration while retaining the same quality. Exploring this opportunity in a study involving live performances is in our immediate research agenda.

VIII. Conclusion

We investigated the use of UWB radios, hitherto unexplored, to support situated live music performances enriched with a tactile sensory layer. Taken together, the objective system-level results from latency and reliability measurements, and the subjective perception results from the user study, indicate that UWB is an ideal match to this kind of applications. This enables opportunities for novel music and art formats exploiting the interplay of musical/sonic content with tactile haptics [42]. Moreover, our findings can also be exploited in support of individuals with hearing impairments, who may leverage the tactile sensory channel to mitigate the deficits at the auditory level [43], [44]. We hope that the present study inspires further research efforts by the academic and industrial research community seizing the many opportunities unlocked by the use of UWB in artistic settings.

References

[1] B. Remache-Vinueza et al., "Audio-tactile rendering: A review on technology and methods to convey musical information through the sense of touch," *Sensors*, vol. 21, no. 19, 2021, Art. no. 6575.

[2] G. W. Young, D. Murphy, and J. Weeter, "Haptics in music: The effects of vibrotactile stimulus in low frequency auditory difference detection tasks," *IEEE Trans. Haptics*, vol. 10, no. 1, pp. 135–139, Jan.–Mar. 2017.

[3] S. Papetti and C. Saitis, Eds., *Musical Haptics* (ser. Springer Series on Touch and Haptic Systems). Cham, Switzerland: Springer, 2018.

[4] L. Turchet, C. Fischione, G. Essl, D. Keller, and M. Barthet, "Internet of musical things: Vision and challenges," *IEEE Access*, vol. 6, pp. 61994–62017, 2018.

[5] A. Haynes, J. Lawry, C. Kent, and J. Rossiter, "FeelMusic: Enriching our emotive experience of music through audio-tactile mappings," *Multimodal Technol. Interact.*, vol. 5, no. 6, p. 29, 2021.

[6] L. Turchet, T. West, and M. Wanderley, "Touching the audience: Musical haptic wearables for augmented and participatory live music performances," *Pers. Ubiquitous Comput.*, vol. 25, no. 4, pp. 749–769, 2021.

[7] Y. Yamazaki, H. Mitake, and S. Hasegawa, "Implementation of tension-based compact necklace-type haptic device achieving widespread transmission of low-frequency vibrations," *IEEE Trans. Haptics*, vol. 15, no. 3, pp. 535–546, Jul.–Sep. 2022.

[8] "What product developers & public locations need to know about auracast broadcast audio," Accessed: Dec. 06, 2024. [Online]. Available: https://www.bluetooth.com/bluetooth-resources/le-audio-webinar/

[9] S. Hayward, K. van Lopik, C. Hinde, and A. West, "A survey of indoor location technologies, techniques and applications in industry," *Internet Things*, vol. 20, 2022, Art. no. 100608.

[10] A. Alarifi et al., "Ultra wideband indoor positioning technologies: Analysis and recent advances," *Sensors*, vol. 16, no. 5, p. 707, 2016.

[11] SPARK Microsystems, "SPARK SR1020 datasheet, rev. 1.3," Accessed: Jun. 18, 2024. [Online]. Available: https://www.sparkmicro.com/wp-content/uploads/2024/03/datasheet_SR1020_customer.pdf

[12] DecaWave Ltd., "DW1000 data sheet, version 2.19," 2017.
[13] L. Turchet, "Smart musical instruments: Vision, design principles, and future directions," *IEEE Access*, vol. 7, pp. 8944–8963, 2019.
[14] F. Hachem, D. Vecchia, M. L. Damiani, and G. P. Picco, "Fine-grained stop-move detection in UWB-based trajectories," in *Proc. 2022 IEEE Int. Conf. Pervasive Comput. Commun*, 2022, pp. 111–118.
[15] T. Istomin et al., "Janus: Dual-radio accurate and energy-efficient proximity detection," in *Proc. ACM Interactive Mobile Wearable Ubiquitous Technol.*, 2021, vol. 5, no. 4, pp. 1–33.
[16] H. Zhang et al., "Cost-effective wearable indoor localization and motion analysis via the integration of UWB and IMU," *Sensors*, vol. 20, no. 2, p. 344, 2020.
[17] A. Li, E. Bodanese, S. Poslad, T. Hou, K. Wu, and F. Luo, "A trajectory-based gesture recognition in smart homes based on the ultrawideband communication system," *IEEE Internet Things J.*, vol. 9, no. 22, pp. 22861–22873, Nov. 2022.
[18] *Standard for Low-Rate Wireless Networks*, IEEE Standard 802.15.4-2020 (Revision of IEEE Standard 802.15.4-2015), 2020.
[19] T. Kern et al., "Actuator design," in *Engineering Haptic Devices*. Berlin, Germany: Springer, 2022, pp. 309–429.
[20] S. Nanayakkara, E. Taylor, L. Wyse, and S. Ong, "An enhanced musical experience for the deaf: Design and evaluation of a music display and a haptic chair," in *Proc. SIGCHI Conf. Hum. Factor Comput. Syst*, 2009, pp. 337–346.
[21] L. Hayes, "Skin music (2012): An audio-haptic composition for ears and body," in *Proc. SIGCHI Conf. Creativity Cogn.*, 2015, pp. 359–360.
[22] E. Gunther and S. O'Modhrain, "Cutaneous grooves: Composing for the sense of touch," *J. New Music Res.*, vol. 32, no. 4, pp. 369–381, 2003.
[23] I. Hwang, H. Lee, and S. Choi, "Real-time dual-band haptic music player for mobile devices," *IEEE Trans. Haptics*, vol. 6, no. 3, pp. 340–351, Jul.–Sep. 2013.
[24] S. Bhagwati et al., "Musicking the body electric: The "body: Suit: Score" as a polyvalent score interface for situational scores," in *Proc. Int. Conf. Technol. Music Notation Representation*, 2016.
[25] J. A. McDowell and D. J. Furlong, "Haptic-listening and the classical guitar," in *Proc. Conf. New Interfaces Musical Expression*, 2018, pp. 293–298.
[26] J. Armitage and K. Ng, "Augmented opera performance," in *Proc. Inf. Technol. Performing Arts, Media Access, Entertainment*, Springer, 2013, pp. 276–287.
[27] J. Armitage and K. Ng,, "Configuring a haptic interface for music performance," in *Proc. Conf. Electron. Vis. Arts. Brit. Comput. Soc.*, 2015, pp. 41–45.
[28] S. Merchel and M. E. Altinsoy, "The influence of vibrations on musical experience," *J. Audio Eng. Soc.*, vol. 62, no. 4, pp. 220–234, 2014.
[29] I. J. Hirsh and C. E. Sherrick Jr., "Perceived order in different sense modalities," *J. Exp. Psychol.*, vol. 62, no. 5, p. 423–432, 1961.
[30] M. Altinsoy, "Perceptual aspects of auditory-tactile asynchrony," in *Proc. 10th Int. Cong. Sound Vib.*, 2003, pp. 3831–3838.
[31] M. E. Altinsoy, "Auditory tactile interaction in virtual environments," Ph.D. dissertation, Technische Universität Dresden, Dresden, Germany, 2006.
[32] M. Daub and M. E. Altinsoy, "Audiotactile simultaneity perception of musical-produced whole-body vibrations," in *Proc. Joint Congr. CFA/DAGA*, 2004, vol. 4, pp. 111–112.
[33] S. Kim, W. L. Martens, and K. Walker, "Perception of simultaneity and detection of asynchrony between audio and structural vibration in multimodal music reproduction," in *Proc. Audio Eng. Soc. Conv.*, 2006.
[34] W. L. Martens and W. Woszczyk, "Perceived synchrony in a bimodal display: Optimal intermodal delay for coordinated auditory and haptic reproduction," in *Proc. Int. Conf. Auditory Display*, 2004.
[35] J. Navarra, S. Soto-Faraco, and C. Spence, "Adaptation to audiotactile asynchrony," *Neurosci. Lett.*, vol. 413, no. 1, pp. 72–76, 2007.
[36] J. G. W. Wildenbeest, D. A. Abbink, C. J. M. Heemskerk, F. C. T. van der Helm, and H. Boessenkool, "The impact of haptic feedback quality on the performance of teleoperated assembly tasks," *IEEE Trans. Haptics*, vol. 6, no. 2, pp. 242–252, Apr.–Jun. 2013.
[37] J. Qin et al., "Effect of packet loss on collaborative haptic interactions in networked virtual environments: An experimental study," *Presence*, vol. 22, no. 1, pp. 36–53, Feb. 2013.
[38] K. Antonakoglou, X. Xu, E. Steinbach, T. Mahmoodi, and M. Dohler, "Toward haptic communications over the 5G tactile internet," *IEEE Commun. Surveys Tuts.*, vol. 20, no. 4, pp. 3034–3059, Fourthquarter 2018.
[39] A. McPherson and V. Zappi, "An environment for submillisecond-latency audio and sensor processing on BeagleBone black," in *Proc. Audio Eng. Soc. Conv.*, 2015.
[40] H. Yao and V. Hayward, "Design and analysis of a recoil-type vibrotactile transducer," *J. Acoust. Soc. America*, vol. 128, no. 2, pp. 619–627, 2010.
[41] L. Vignati, S. Zambon, and L. Turchet, "A comparison of real-time Linux-based architectures for embedded musical applications," *J. Audio Eng. Soc.*, vol. 70, no. 1–2, pp. 83–93, 2022.
[42] S. C. Aker et al., "Effect of audio-tactile congruence on vibrotactile music enhancement," *J. Acoust. Soc. America*, vol. 152, no. 6, pp. 3396–3409, 2022.
[43] F. Sorgini, R. Caliò, M. C. Carrozza, and C. M. Oddo, "Haptic-assistive technologies for audition and vision sensory disabilities," *Disabil. Rehabil.: Assistive Technol.*, vol. 13, no. 4, pp. 394–421, 2018.
[44] M. D. Fletcher, "Can haptic stimulation enhance music perception in hearing-impaired listeners?," *Front. Neurosci.*, vol. 15, 2021, Art. no. 723877.

Luca Turchet (Senior Member, IEEE) received the master's degree in computer science from the University of Verona Verona, Italy, in 2006, the degrees in classical guitar and composition from the Music Conservatory of Verona, in 2007 and 2009, respectively, the degree in electronic music from the Royal College of Music of Stockholm, Stockholm, Sweden, in 2015, and the Ph.D. degree in media technology from Aalborg University Copenhagen, Copenhagen, Denmark, in 2013. He is currently an Associate Professor with the Department of Information Engineering and Computer Science of the University of Trento, Trento, Italy.

Christian Sassi received the bachelor's degree in computer, communications and electronic engineering from the University of Trento, Trento, Italy, in 2023, with a final project based on ultra wideband (UWB) for the development of a haptic application, and the master's degree with the Department of Information Engineering and Computer Science, University of Trento. His research interests include UWB, music technology, and cybersecurity.

Davide Vecchia received the M.Sc. degree in computer science and the Ph.D. degree in information and communication technology from University of Trento, Trento, Italy. He is currently a Postdoctoral Researcher with the Department of Information Engineering and Computer Science, University of Trento. His research interests include ultra wideband communication and localization in large-scale networks. He was the recipient of Best Paper awards at IPSN (2023) and IPIN (2019).

Gian Pietro Picco (Senior Member, IEEE) is currently a Professor with the University of Trento, Trento, Italy. His research interests include software engineering, middleware, and distributed systems, and focuses on low-power wireless networking, localization for Internet of Things, and cyber-physical systems. He was the recipient of several scientific awards. He is also with organizing committees and editorial boards of flagship venues in the fields above. He is the founding Editor-in-Chief of the ACM Transactions on Internet of Things.