



HAL
open science

PUMAH: Pan-tilt Ultrasound Mid-Air Haptics for larger interaction workspace in virtual reality

Thomas Howard, Maud Marchal, Anatole Lécuyer, Claudio Pacchierotti

► To cite this version:

Thomas Howard, Maud Marchal, Anatole Lécuyer, Claudio Pacchierotti. PUMAH: Pan-tilt Ultrasound Mid-Air Haptics for larger interaction workspace in virtual reality. *IEEE Transactions on Haptics (ToH)*, 2020, 13 (1), pp.38-44. 10.1109/TOH.2019.2963028 . hal-02424247

HAL Id: hal-02424247

<https://inria.hal.science/hal-02424247v1>

Submitted on 26 Dec 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

PUMAH : Pan-tilt Ultrasound Mid-Air Haptics for larger interaction workspace in virtual reality

Thomas Howard, Maud Marchal, Anatole Lécuyer and Claudio Pacchierotti

Abstract—Mid-air haptic interfaces are promising tools for providing tactile feedback in Virtual Reality (VR) applications, as they do not require the user to be tethered to, hold, or wear any system or device. Currently, focused ultrasound phased arrays are the most mature solution for providing mid-air haptic feedback. They modulate the phase of an array of ultrasound emitters so as to generate focused points of oscillating high pressure, eliciting vibrotactile sensations when encountering a user’s skin. While these arrays feature a reasonably large vertical workspace, they are not capable of displaying stimuli far beyond their horizontal limits, severely limiting their workspace in the lateral dimensions. In this paper, we propose an innovative low-cost solution for enlarging the workspace of focused ultrasound arrays. It features 2 degrees of freedom, rotating the array around the pan and tilt axes, thereby significantly increasing the usable workspace and enabling multi-directional feedback. Our hardware tests and human subject study in an ecological VR setting show a 14-fold increase in workspace volume, with focal point repositioning speeds over 0.85m/s while delivering tactile feedback with positional accuracy below 18mm. Finally, we propose a representative use case to exemplify the potential of our system for VR applications.

Index Terms—Mid-Air Haptics, Robotics, Virtual Reality

I. INTRODUCTION

Haptic feedback is an important component of immersive virtual reality (VR) interactions [1]. However, conventional grounded [2] and wearable [3], [4] haptic devices are often still too bulky and obtrusive, adversely affecting the quality of interaction with virtual environments. Mid-air haptic interfaces seek to overcome these limitations by providing cutaneous feedback remotely, without requiring the user to wear or hold any device or tool providing the stimuli. Currently, one of the most mature solutions for providing this type of non-contact feedback is focused ultrasound phased arrays [5], [6]. In these devices, a 2-dimensional grid of ultrasonic transducers is driven such that the phase delay between actuators creates constructive and destructive interference patterns in the sound waves propagated above the device. This results in one or more small localized regions of oscillating high pressure (referred to as focal points), surrounded by regions where the air pressure remains more or less constant. When the skin encounters a focal point, the user feels a localized vibrotactile stimulus. Due to their functioning principle, these arrays are only capable of conveying haptic feedback in the region directly above them, and the stimuli provided are felt the strongest when the skin is parallel to the device plane. Providing haptic feedback in a larger workspace is nonetheless beneficial for a natural and rich interaction in virtual environments. To do so, researchers have proposed different solutions, such as using larger arrays [7], [8], or linking multiple arrays in co-planar [9], [10], non co-planar [11], [12], opposing [13], and surrounding [14], [15] configurations. However, all these systems suffer major drawbacks due to their high cost, complex control, high power requirements, bulkiness, and limited reconfigurability. An alternative approach has been presented by Sand

This project has received funding from the European Union’s Horizon 2020 programme under grant agreement No 801413; project “H-Reality”.

T. Howard and C. Pacchierotti are with Univ Rennes, CNRS, Inria, IRISA – Rennes, France. {thomas.howard, claudio.pacchierotti}@irisa.fr

A. Lécuyer is with Univ Rennes, Inria, CNRS, IRISA – Rennes, France. anatole.lecuyer@inria.fr

M. Marchal is with Univ Rennes, INSA, IRISA, Inria, CNRS – Rennes, France. maud.marchal@irisa.fr




Fig. 1. Assembled (left) and exploded (right) view of the PUMAH system design. The array (1) is mounted on an aluminium holding plate (2) using 3D-printed ABS clips at the corners (3). The plate rotates around the $Pivot_x$ -axis (tilt) within an aluminium tubing and ABS frame (4), which itself rotates around the device vertical axis $Base_z$ (pan). The axes are driven by HiTec HS645-MG (5) and HS625-MG (6) servomotors. They are mounted on bearings held within aluminium chassis (7), relieving the motor shafts of any radial loads. The complete system is mounted on a 3D-printed ABS foot (8), which can be screwed to a supporting structure or mounted on a tripod using M6 screws.

et al. [16], who mounted a focused ultrasound array onto a head-mounted display (HMD), so that the workspace moves with the user. However, this approach is still problematic due to the limits imposed on the array size and added mass on the head. Furthermore, the device is only capable of providing haptic feedback right in front of the HMD, and only to the skin facing the array. To address this important limitation, we propose PUMAH, a novel 2-degrees-of-freedom (2-DoF) pan-tilt mount for a focused ultrasound array, the Ultrahaptics Stratos Explore [17]. While our device is built around this model of array, it can easily be adapted to any other focused ultrasound array of similar dimensions. The proposed system has a very low added cost (under EUR 150), and it is able to significantly increase the rendering workspace as well as the haptic rendering quality by actively reorienting itself toward the interaction region. Rotating the device also offers new possibilities for delivering sensations from multiple different directions. This paper presents the PUMAH device (Sec. II) for the first time, along with an evaluation of its dynamics (Sec. III) and workspace via a user study (Sec. IV). Finally, we show the viability of our device in an immersive VR scenario (Sec. V).

II. MECHANICAL DESIGN AND COMPONENTS

A. Mechanical Design

Our device functions as a servo-driven pan and tilt mount for a focused ultrasound array (see Fig. 1). The pan and tilt servo motors are HiTec HS645-MG and HS625-MG [18], respectively. They were chosen for their low cost, high holding torques of 0.94Nm and 0.66Nm and high rotation speeds of 300°/s and 400°/s, respectively. All cables running to the array are secured to the back of the aluminium holding plate so as to prevent disconnecting and damage during device operation, and the servomotor controller is attached to the aluminium and ABS rotating structure (parts 4 in Fig. 1). We define the central reference position P_{ref} as 0° on the pan axis and

0° on the tilt axis. In this configuration, the array face is located in a horizontal plane facing upward, with the array y-axis y_{arr} pointing away from the user. We refer to the PUMAH base z-axis as the pan axis and to the pivot x-axis as the tilt axis. The intersection of these axes defines the device center of rotation, which we will refer to as “device pivot” in the remainder of the paper. From P_{ref} , the tilt axis allows a $+85^\circ$ upward and a -130° downward tilting motion. The pan axis allows a $+90^\circ$ (rightward) and -90° (leftward) rotation of the array.

B. Electronics and Driver

The servomotors are driven using a Pololu Mini Maestro-24 USB servo controller board. The board is powered by a 5V, 2.4A DC power supply and receives position commands generated within a VR application running on a computer connected via USB. To limit unnecessary bandwidth usage and computations on the micro-controller, updates to the target positions for both axes are only sent if they differ from the previous target position, at a rate up to the frame rate of the VR application (approx. 90Hz).

C. Control

Pan and tilt error angles are computed based on the angular error between the array normal and the vector from the pivot to a target defined on an application-specific basis. They are then used as commands for the servomotors on each axis. Targets can range from static positions (e.g. see Sec. IV on workspace experimental evaluation) to dynamic tracking of the user’s hand (e.g. see the presented use case in Sec. V). The internal motor servoing mechanism was found to be sufficient to obtain fast and accurate responses to commands (see Sec. III), hence the system is driven in an open loop configuration. To ensure accuracy despite non-linearity in the relationship between pulse width modulation (PWM) and servo shaft angle, we perform a calibration step where the pulse width is recorded at fixed intervals in shaft angle. The servos are then considered as having a linear response (in shaft angle) between these recorded positions, such that the PWM for a desired angle can be obtained from linear interpolation between two known values.

D. Integration into VR Environments

Integrating the device into a VR environment to provide mid-air haptic feedback requires two main steps: an initial calibration and a definition of the target orientation (e.g. through user hand tracking). For calibration, the pan-tilt mount is first driven to P_{ref} (0° tilt, 0° pan). The transformation between device coordinates and virtual world coordinates is then obtained by measuring three calibration points on the device. From these points, the virtual model of the device (and the entire virtual scene) are rotated and translated such that they coincide with positions of their real-world counterparts. Once the virtual and real devices are aligned, a target in the workspace needs to be defined and its position computed (e.g. in the case of a purely virtual target) or acquired through tracking (e.g. in the case of the hand). The vector v_{target} between the device pivot and target is computed at every frame. The tilt and pan angular errors can be computed by projecting this vector onto the array y-z-plane and device base x-z-plane, respectively. To verify the correct execution of the tracking and angle computations, the virtual model of the device is rotated by $tilt_{error}$ around $pivot_x$, then by pan_{error} around $base_z$ (see Fig. 1). The virtual model of the device thus acts as an ideal representation of our system’s target behavior. The computed tilt and pan error angles are also translated into target PWM values for the servo-motors based on the method discussed in Section II-C and transmitted to the Maestro controller via USB.

III. CONTROL EVALUATION

We performed a set of experiments to evaluate the performances and limits of the PUMAH’s design and control scheme.

A. Step Response

The step response of the device was evaluated for each axis individually by driving it to a $+45^\circ$ target angle, then commanding a -90° rotation and recording the actual rotation angle using a Vive tracker rigidly attached to the array and aluminium holder plate (see Fig. 2 left). The position was left to settle, after which the device was driven back to its original position with a $+90^\circ$ rotation, allowing another measurement of the step response to be performed. This process was repeated a total of 20 times for each axis, then another 40 times with both axes driven simultaneously in identical (both axes rotating towards angles of the same sign) and opposing (both axes rotating towards angles of opposing signs) configurations respectively.




Fig. 2. Left: Experimental set-up for the evaluation of the device’s performance. A Vive tracker (1) is rigidly attached to the holding plate and array. Right: To compute the device’s actual pan and tilt angles, a calibration step is performed to compute the pivot’s (P) position from the Vive tracker data. The horizontal and vertical planes as well as the pan and tilt axes are computed based on Vive tracker acquisitions respectively at fixed tilt and pan angles. For each point x of a subsequent acquisition, we compute the horizontal and vertical projections x_h and x_v . The angle between $tilt_0$ and $\overrightarrow{0_P x_h}$ yields the current pan and the angle between \overrightarrow{pan} and $\overrightarrow{P x_v}$ yields the current tilt angle.

The system’s behavior was found to be very repeatable, hence Fig. 3 shows 4 representative plots of the step responses obtained. The pan axis response shows a typical under-damped second order system response, with a mean rise time (0% to 100%) of 0.49s (SD = 0.0073s) and a mean settling time of 0.73s (SD = 0.0088s). The axis exhibits large overshoot (mean 14.48°, i.e. 16%, SD = 0.63°) but maintains minimal final error (mean 0.58°, i.e. 0.6%, SD 0.3°). For a focal point generated at a distance of 45cm above the array (i.e. at the limit of the usable workspace, see Sec. IV), this final error would translate to a positioning inaccuracy of approx. 4.5mm, about half the size of a focal point. The tilt axis response also shows the profile of an under-damped second order system, with both rise time (mean 0.48s, SD = 0.097s) and settling time (mean 0.63s, SD = 0.14s) similar to those of the pan axis, but with a higher variability. This axis also overshoots (mean 8.88°, i.e. 9.8%, SD = 1.7°) and shows larger final error than the pan axis (mean 1.73°, i.e. 1.9%, SD = 0.26), which can be attributed to uneven load distribution around the axis as well as residual calibration errors. For a focal point generated at a distance of 45cm above the array, this final error would translate to a positioning inaccuracy of approx. 13mm. When simultaneously driving the pan and tilt axes, we observe similar response profiles but with larger final error on the pan axis (mean 2.32° for pan ; mean 1.75° for tilt). The final errors also show larger variability (SD 1.39° for pan ; SD 1.17° for tilt), as well as some cross-contamination between axes (90° steps on the tilt axis introduce a mean error of 1.08° on




Fig. 3. Example plots of step responses obtained for the pan only (top left), tilt only (bottom left), and combined (right) trials. The dotted blue lines show target positions while the full red lines show the actual positions.

the pan axis, while 90° steps on the pan axis introduce a mean error of 1.93° on the tilt axis). Closer analysis of the final errors shows a cyclical profile consistent with the effects of play within the system. Design changes aiming at reducing play could probably improve our device's positional accuracy. Rise times (means: 0.56s for pan, 0.52s for tilt) and settling times (mean 0.74s for pan, 0.6s for tilt) remain practically unchanged, as does overshoot (mean 10.65° for pan, 6.5° for tilt). For a focal point at the limits of the usable workspace, the final error would translate to a positioning inaccuracy of approx. 18mm, and one could expect a delay of 0.74s in obtaining a tactile sensation at a desired target.

B. Sinusoidal Response

To evaluate the frequency response of the PUMAH, it was initially set to its central position P_{ref} . Then, the tilt and pan axes were individually driven for a duration of 30s with commands in the form of $target = A \sin(2\pi ft)$ [°], with amplitudes A set to 45° and 90° and frequency f varying between 0.1Hz and 4Hz over 13 steps (respectively steps of 0.1Hz below 1Hz, and 1Hz steps above). Actual array positions were recorded identically to the above step response protocol.

Fig. 4 shows the obtained sinusoidal responses in the form of bode plots. In the worst case scenario ($A=90^\circ$), driving the device at rates above 0.6Hz leads to severely degraded performance, while in the mean scenario ($A=45^\circ$), performance only starts degrading above 1Hz, with a cutoff frequency around 1.5Hz. Phase shift remains more or less constant below 0.6Hz regardless of the scenario. Above this value, the tilt axis tends to get out of phase more rapidly than the pan axis. In practical terms, this means that the device is capable of consistently tracking movements up to 0.85m/s (for a target at the edge of the usable workspace 0.45m from the array).

C. Device Accuracy, Repeatability and Drift

The results from the step response evaluation shown in Fig. 3 allow a first estimate of the device accuracy and repeatability. Fig. 5 shows that the accuracy (measured as mean absolute angular error) is similar for both axes, at 1.42° in the mean for the pan axis and 2.07° in the mean for the tilt axis. We observe a slightly more repeatable behavior on the pan axis (1.52° interquartile range for pan against 2.03° for tilt), which is probably due to a more even load distribution when compared to the tilt axis. Changing the device holding plate to better distribute loads around the tilt axis may improve repeatability. Given the current relatively low spatial resolution of the Stratos Explore's tactile feedback (the perceived diameter of a focal point is around 1cm), these positioning errors are not expected to have a major




Fig. 4. Bode plots obtained for the sinusoidal responses for the pan axis (blue) and tilt axis (green). The dotted lines show the frequency response in the worst-case scenario (amplitude $A = 90^\circ$) while the full lines show the frequency response for a mean scenario (amplitude $A = 45^\circ$).




Fig. 5. Device angular errors per axis. The red line shows the mean (device accuracy) while the horizontal extent of the boxes shows axis repeatability.

impact on the quality of the delivered tactile feedback. However, if significant improvements to the device accuracy were required, closed-loop angular position control of axis angles would likely need to be implemented, which would impact device cost and complexity.

IV. RENDERING WORKSPACE

To estimate the workspace within which sensations can be rendered by the PUMAH, we first performed a user study to experimentally determine the Stratos Explore's workspace. We then performed an identical study to determine the usable workspace of the PUMAH, thus quantifying improvements in workspace volume and haptic feedback quality. User studies were chosen over e.g. microphone measurements as a more direct way of evaluating user experience of the display's output, given that the relationship between generated acoustic radiation pressure and perceived stimulus intensity is still an open research question.

A. Experimental Procedure and Design

A total of 15 subjects (11 male, 4 female, 13 right-handed, 2 left-handed, ages 22-36 (mean: 26y, SD: 3.5y)) performed the experiment after providing written informed consent.

a) *Experimental Setup*: Subjects faced a small table onto which the device was mounted (see Fig. 6 left). They wore a HTC Vive Pro HMD, tracked using a pair of Lighthouse base stations, displaying a simple virtual environment containing a model of the table without showing the device. A Vive Tracker was attached with a pair of straps in such a way as to keep the palm unobstructed, tracking their dominant hand. They held a Vive controller in their non dominant hand, which they used to respond to questions during the experiment.




Fig. 6. Setup for the experiment in VR. Right: VR views during the experiment. Left: Subjects stand in front of the table (5) onto which our device (4) is mounted. They wear a HMD (1) displaying a virtual environment containing a model of the table (8), target cones (7) or a question (9). Their dominant hand is tracked (2), allowing the VR view to show a hand avatar with a green disk co-located with their actual palm (6) which represents the interaction region they are to use to swipe through the target cones (7). Subjects use the buttons on a Vive controller (3) held in their non-dominant hand to answer questions (9) after each trial.

b) *Experimental Procedure:* In the virtual environment, the user could at all times see the table, a hand avatar superimposed on their actual hand as well as a model of the Vive controller (see Fig. 6 right). Before the experiment began, a calibration step was performed to ensure proper alignment of the real and virtual hand locations. Subjects then spent a short time familiarizing themselves with a reference stimulus considered strong (stimulus presented 20cm above the array origin at 155dB SPL). Subjects began by touching a virtual blue sphere with their hand to launch a trial. After that, the sphere disappeared and a virtual cone appeared at a randomly selected position in the virtual environment. The targets cones always pointed towards the device pivot. Subjects were instructed to slowly swipe the palm of their hand through the cone, starting at the base and moving towards the tip (see Fig. 6). During the time the cone is displayed, the array attempted to render a small circle (diameter 2cm) centred on the mid-point of the cone vertical axis and in a plane parallel to the array. The circle was rendered using spatiotemporal modulation [19] at a drawing frequency of 100Hz and maximum intensity. Once subjects swiped through the cone, it disappeared and was replaced with a question asking the subjects how strongly they felt tactile feedback at the target, if at all. Responses were provided via a 3-point scale with the options “Not at all”, “Weak” and “Strong”. To limit the risk of collisions between the user’s hand and the PUMAH, bounding boxes around the PUMAHs actual position were displayed whenever the user’s hand got too close to it. The experiment was divided into two conditions with 3 repetitions each. In the first condition (static), the array was kept fixed at P_{ref} , facing upward. This condition represents the current uses of the Ultrahaptics array in the literature. In the second condition (moving), our device rotated so as to minimize the angular error between the array vertical axis and the vector pointing from the pivot to the currently displayed target. This condition represents a standard use-case for the PUMAH. We randomized the order of conditions between subjects and the target presentation order within each trial. A set of 45 points distributed throughout the space above the array plane was sampled in the static condition, while 40 additional points were sampled all around the device pivot in the moving condition (see Fig. 7). The point locations were chosen based on a pilot study determining the limits of the array workspace via a method of adjustments. During this pilot study, subjects were tasked with moving a focal point along the array z-axis as well as along the




Fig. 7. Target locations in the experiment, relative to the device location. The 45 targets for the static condition are shown in blue. The additional 40 points shown in green were only sampled in the moving condition as they lay far outside the expected array workspace limits based on our pilot study.

x- and y-axes at different heights so as to determine the locations at which the focal point became imperceptible.

B. Results and Discussion




Fig. 8. (A): Experimentally-determined rendering workspace limits for the static and (B): for the moving conditions. The blue rectangles show the array plane at P_{ref} . The blue dots show the locations sampled in the experiment, while the red disks have a diameter proportional to the mean of participant’s median responses for each point. The red convex hull shows the region within which points were detected as “Strong” more than 50% of the time, while the yellow convex hull shows the region within which points were at least detected as “Weak” more than 50% of the time. The plot origins are located at the device pivot and the axes are aligned with the array axes at P_{ref} .

The median of each subject’s 3 responses for a given point in a given condition was considered representative of the subjective intensity of the tactile stimulus at that location. Fig. 8A shows the obtained workspace limits for the static array while Fig. 8B shows the workspace limits for the PUMAH. For the static array, the outermost points of the sampled volume consistently yielded reports of no stimuli felt, except for the uppermost central point along the vertical axis, which was reported as weak. This allows

us to confidently consider the convex hull of the points reported at least as “weak” over 50% of the time (yellow in Fig. 8A) as a good approximation of the perceptible workspace. This yields a volume of 0.055m^3 shaped like an asymmetrical ellipsoid (oblate at the bottom, prolate at the top) spanning from $z = 40\text{mm}$ to $z = 700\text{mm}$ above the device with a lateral radius of approx. 320mm at its widest. Within this volume lies a smaller similar-shaped volume of 0.0016m^3 within which targets were identified as “strong” more than 50% of the time (red in Fig. 8A). This “strong” region spans $z = 40\text{mm}$ to $z = 430\text{mm}$ with a radius of approx. 140mm at its widest point ($z = 70\text{mm}$). While these threshold envelopes are interesting from a psychophysical perspective, in terms of rendering, the volume within which tactile stimuli are systematically detected by a significant proportion of users could also be a relevant measure. We therefore computed the envelope of our sampled points which were reported at least as “weak” for *all* of our subjects. This yields an intermediary volume of 0.0254m^3 , referred to as the “usable volume” (shown in green in Fig. 9 left). For the moving array, we obtained a “strong” volume of approx. 0.122m^3 , spanning an approximately half-ellipsoidal region above the horizontal plane bisecting the device at the level of the pivot, spanning 0.9m back to front, 0.45m side to side and 0.45m upward. This amounts to a 76-fold increase in volume (i.e. approx. 4-fold increase along the linear dimension) for the “strong” region. While this increase is already substantial, we expected to see this volume extend somewhat symmetrically around the x-y-plane. However, as can be seen in Fig. 9 (right), this is not the case. The fact that targets beneath the x-y-plane were predominantly reported as “weak” could stem from our sampling of the region being too sparse. Alternatively, it could also reflect perceptual effects where focal points projected from above onto the palm create different sensations from identical focal points projected from below onto a palm. Further investigation would be required to determine each factor’s contribution to this observation. Also, contrary to our results for the static array, almost all targets including those on the outer edges of the 0.95m^3 sampled region are detected at least as “weak” more than 50% of the time. While there is a possibility that the PUMAH’s actual workspace extends beyond the sampled 0.73m radius spherical region, the moving condition required more movement from the subjects when reaching for the targets, with rapid hand movements easily generating airflow which might be mistaken for a weak stimulus. As for our evaluation of the static array, we calculate the volume within which points are detected by all users 100% of the time as 0.368m^3 , yielding a 14-fold increase in usable workspace volume (in green in Fig. 9).

Given the observed increase in size of the region where tactile stimuli were reported as “strong” (red in Figs. 8 and 9), it appears that beyond extending the array’s usable workspace, the PUMAH allows for improved tactile feedback quality within the workspace of the static array by optimizing the position of the array relative to a user’s hand. We therefore analysed the points within the static device’s workspace (points within the yellow envelope in Fig. 8A) to quantify these improvements. Fig. 10 shows the results of this analysis.

All points within the static device’s workspace were detected identically or better in the moving condition compared to the static condition. Points along $\overrightarrow{Base_z}$ were detected identically in the moving condition (blue in Fig. 10). The points furthest from the array vertical axis and close to the board saw the largest increase in stimulus quality, going from undetected to detected as strong stimuli (red in Fig. 10). The remainder of the points saw varying increases in stimulus quality in the moving condition, confirming that our device both increases the usable workspace and optimizes tactile stimulus delivery.





Fig. 9. Cross-sections of the three volumes computed based on user responses for the static (left) and moving (right) conditions. The yellow region shows the points reported at least as weak more than 50% of the time (detection threshold), the green region shows the points reported at least as weak 100% of the time (usable volume for rendering) and the red region shows the points reported as strong more than 50% of the time. The plot origins are located at the device pivot and the axes are aligned with the array axes at P_{ref} .

V. USE CASE

To showcase the capabilities of our device and its viability for VR applications, we developed a VR demo allowing users to interact with and explore virtual 3D shapes of various sizes throughout the workspace. The virtual environment contains a holographic projector displaying a circular menu of 4 shapes around it (see Fig. 11 top). In the menu, the shapes are shown as miniatures (each approx. 10cm across) and slowly rotate about their centre of mass. Users hold a Vive controller in their left hand, allowing to rotate the menu in 90° steps as well as to blow up the object directly in front of them. Doing this hides the menu and centers the currently selected object on the projector, tripling its size (see Fig. 11 bottom). The user’s hand position is tracked using a Vive tracker and the device axis rotation angles are adjusted so as to minimize angular error between the array normal vector and the vector between the pivot and user’s palm. Tactile feedback is provided when the user’s hand collides with the virtual objects, giving the compelling feel of a touchable hologram (see video in supplemental material). Collisions between the user and PUMAH are avoided using bounding boxes displayed in the same way as described in Section IV-A.

VI. CONCLUSION AND PERSPECTIVES

We presented a novel pan-tilt-mount for a focused ultrasound mid-air haptic interface, making it capable of providing tactile feedback across a significantly larger workspace and from multiple directions, with a very low added material and computational cost. Our device is shown to functionally integrate into VR environments. It is capable of tracking user or target movements up to 0.85m/s in the worst case configuration, while delivering tactile feedback with positional accuracy below 18mm . Through user studies, we validate a 14-fold increase in the usable rendering workspace as well as the ability for providing tactile feedback in all directions except directly below the device. Furthermore, comparative analyses of the subjectively perceived strength of focal points at different locations in the workspace show that the device provides the possibility of optimizing tactile feedback depending on the orientation of the user’s hand. With adequate modelling or tracking of device positional errors, dynamic repositioning of the focal point during movement could possibly compensate for the PUMAH’s inaccuracy, and will be a focus of future work. While current results will be used to improve future designs, our main focus will now target new possibilities for



X	Y	Z	Mean Median Response (STATIC)		Mean Median Response (MOVING)		p-Value (WMW-test)
			Mean	Median	Mean	Median	
207	0	92	0,07	NOT AT ALL	1,8	STRONG	5,11E-07
-207	0	92	0,13	NOT AT ALL	1,8	STRONG	9,53E-07
0	207	92	0	NOT AT ALL	1,87	STRONG	1,86E-07
0	-207	92	0	NOT AT ALL	1,93	STRONG	1,29E-07
-527	0	258	0	NOT AT ALL	1	WEAK	7,54E-07
0	-123	99	0,93	WEAK	2	STRONG	1,29E-07
0	-107	469	1,2	WEAK	1,73	STRONG	4,31E-03
0	-212	467	1	WEAK	1,6	STRONG	2,08E-03
523	0	460	0,2	NOT AT ALL	0,73	WEAK	4,31E-03
0	-523	460	0,27	NOT AT ALL	0,8	WEAK	4,31E-03
166	0	271	1,4	WEAK	1,93	STRONG	6,43E-03
0	166	271	1,3	WEAK	1,87	STRONG	9,06E-03
323	0	264	1	WEAK	1,67	STRONG	2,53E-03
0	-323	264	0,87	WEAK	1,8	STRONG	1,49E-05
-323	0	264	1,07	WEAK	1,67	STRONG	7,98E-03
527	0	258	0,07	NOT AT ALL	1	WEAK	4,30E-06
0	527	258	0,3	NOT AT ALL	1,2	WEAK	1,18E-04
0	-527	258	0,13	NOT AT ALL	1	WEAK	1,99E-05
0	123	99	1,3	WEAK	1,87	STRONG	3,65E-03
123	0	99	1,13	WEAK	1,73	STRONG	3,74E-03
-123	0	99	1,07	WEAK	1,93	STRONG	2,04E-05
-212	0	467	1	WEAK	1,4	WEAK	2,54E-02
0	523	460	0,27	NOT AT ALL	0,67	WEAK	3,28E-02
-166	0	271	1,7	STRONG	1,93	STRONG	1,54E-02
0	-166	271	1,47	WEAK	1,93	STRONG	1,54E-02
0	323	264	1,07	WEAK	1,53	STRONG	3,74E-02

Fig. 10. Point-by-point comparison of subjective reports of tactile feedback intensity between stimuli delivered in the static and moving conditions, within the static array’s workspace. Blue: no change in reported intensity between conditions. White: a change was observed but of an insufficient magnitude to affect point attribution to one of the three volumes previously computed. Yellow: between conditions, points changed either from not detected to detected, from detected to within the usable workspace or from within the usable workspace to detected as strong. Green: between conditions, points either went from not detected to within the usable workspace or from detected to detected as strong. Red: between conditions, points went from not detected to detected as strong. Points for which the difference in responses between conditions was found to be statistically significant (Mann-Whitney-Wilcoxon tests) are highlighted using a black cross. The table shows the difference in responses as well as computed p-values for each of these points, color-coded to match the plot. All coordinates are provided in [mm] relative to the pivot.

interaction techniques and rendering. Actively reorienting focused ultrasound arrays, array sub-units or individual transducers opens up a range of mid-air rendering possibilities, in particular regarding interaction with and discrimination of virtual 3D haptic shapes. Furthermore, the motion of the device itself may be used to produce interesting haptic motion effects. Our device was shown to improve a single array’s workspace, however it still targets only applications for one or two hands close together. Since VR applications will call for bimanual or possibly multi-user interactions, we plan mount two arrays mounted back to back on a single PUMAH and work with multiple PUMAH devices in future work. Finally, the PUMAH could also serve to investigate mid-air vibrotactile stimulus perception depending on the user’s posture and position relative to the stimuli.

REFERENCES

[1] J. Kreimeier *et al.*, “Evaluation of different types of haptic feedback influencing the task-based presence and performance in virtual reality,” in *Proc. 12th ACM Int. Conf. on Pervasive Technologies Related to Assistive Environments*, 2019, pp. 289–298.




Fig. 11. VR scene showing interaction with the hologram menu (top), with haptic feedback for multiple objects in the workspace, and with an enlarged hologram (bottom), showcasing haptic interaction from multiple angles.

[2] V. Hayward *et al.*, “Haptic interfaces and devices,” *Sensor Review*, vol. 24, no. 1, pp. 16–29, 2004.

[3] C. Pacchierotti *et al.*, “Wearable haptic systems for the fingertip and the hand: Taxonomy, review, and perspectives,” *IEEE Trans. Haptics*, vol. 10, no. 4, pp. 580–600, 2017.

[4] J. Perret *et al.*, “Touching virtual reality: A review of haptic gloves,” in *Proc. of Int. Conf. on New Actuators*, 2018, pp. 1–5.

[5] T. Iwamoto *et al.*, “Airborne ultrasound tactile display,” in *Proc. of ACM SIGGRAPH New Tech Demos*, 2008.

[6] T. Carter *et al.*, “Ultrahaptics: multi-point mid-air haptic feedback for touch surfaces,” in *Proc. ACM Symp. User interface software & technology*, 2013, pp. 505–514.

[7] T. Hoshi *et al.*, “Non-contact tactile sensation synthesized by ultrasound transducers,” in *Proc. IEEE World Haptics Conf.*, 2009, pp. 256–260.

[8] K. Hasegawa *et al.*, “Aerial display of vibrotactile sensation with high spatial-temporal resolution using large-aperture airborne ultrasound phased array,” in *Proc. IEEE World Haptics Conf.*, 2013, pp. 31–36.

[9] G. Korres *et al.*, “Haptogram: Ultrasonic point-cloud tactile stimulation,” *IEEE Access*, vol. 4, pp. 7758–7769, 2016.

[10] K. Hasegawa *et al.*, “Aerial vibrotactile display based on multiunit ultrasound phased array,” *IEEE Trans. Haptics*, vol. 11, no. 3, pp. 367–377, July 2018.

[11] T. Hoshi *et al.*, “Adding tactile reaction to hologram,” in *Proc. IEEE Int. Symp. Robot & Human Interactive Communication*, 2009, pp. 7–11.

[12] T. Hoshi, “Development of aerial-input and aerial-tactile-feedback system,” in *Proc. IEEE World Haptics Conf.*, 2011, pp. 569–573.

[13] T. Hoshi *et al.*, “Three-dimensional noncontact manipulation by opposite ultrasonic phased arrays,” *Japanese J. Applied Physics*, vol. 53, 2014.

[14] S. Inoue *et al.*, “Horn: The hapt-optic reconstruction,” in *Proc. ACM SIGGRAPH Emerging Tech.*, 2014, pp. 11:1–11:1.

[15] —, “Active touch perception produced by airborne ultrasonic haptic hologram,” in *Proc. IEEE World Haptics Conf.*, 2015, pp. 362–367.

[16] A. Sand *et al.*, “Head-mounted display with mid-air tactile feedback,” in *Proc. ACM Symp. Virtual Reality Software & Technology*, 2015, pp. 51–58.

[17] [Online]. Available: www.ultrahaptics.com/products-programs/

[18] [Online]. Available: <https://hitecrcd.com/products/servos/sport-servos>

[19] B. Kappus *et al.*, “Spatiotemporal modulation for mid-air haptic feedback from an ultrasonic phased array,” *J. Acoustical Soc. of America*, vol. 143, pp. 1836–1836, 2018.