Mid-Air Haptics: Future Challenges and Opportunities

Orestis Georgiou[1], William Frier[1], Euan Freeman[2], Claudio Pacchierotti[3] and Takayuki Hoshi[4]

[1] Ultraleap Ltd - Bristol, UK
[2] University of Glasgow - Glasgow, UK
[3] CNRS, Univ Rennes, Inria, IRISA - Rennes, France
[4] Pixie Dust - Tokyo, Japan

Abstract: Ultrasound mid-air haptic technology has advanced in many ways over the past decade and has found meaningful application in a plethora of use cases. As the technology matures further and progresses from lab to market, in this chapter we take a step back and discuss three specific directions that we think could result in the greatest impact. Namely, we highlight challenges and opportunities in improving (1) the hardware platforms used, (2) the rendering algorithms employed to create rich haptic sensations, and (3) the resulting user experience and added value the technology can instill to different end-user applications. We hope that this 'wish-list' inspires the mid-air haptics and human computer interaction (HCI) community and others to join our efforts towards a deeper technology understanding, integration, and readiness.

1. Introduction

Since the discovery of ultrasonic mid-air haptic technology in Japan in 2010 [1] and its commercialization by Ultraleap (UK) [2] in 2014, we have seen a rapid and diverse advancement in its development - especially within the last few years. The previous chapters of this book have given a flavor of this development, spanning hardware designs, improved software algorithms for haptic rendering, a focus on enhanced user experiences and immersive applications, and finally a deeper understanding of the physical and acoustic processes involved when inducing contactless touch. Moreover, this book is living proof of a growing and highly interdisciplinary community of researchers who actively publish beyond HCI and haptics journals and conferences.

Despite all this progress, in their recent review, Rakkolainen et al. [3] conclude that ultrasound haptics is still in its infancy, and highlight five directions that can help the technology mature: 1) a greater understanding of perception through more user studies, 2) new haptic rendering methods to allow higher quality haptic output, 3) an exploration of new haptic sensations and textures, 4) improved acoustic solvers, and 5) new applications that leverage the many benefits of mid-air haptic feedback. We agree with their assessment but would also like to close this book by adding further tracks towards technology maturity and expanding on them through our own lens. The aim of this final chapter is therefore to look at the road ahead and highlight key challenges and opportunities associated with the further advancement of this emergent touchless technology. We hope this inspires the mid-air haptics community and others to join us on this journey, as we take the next steps forward towards a deeper technology understanding, integration, and readiness. We group our discussions under the themes of hardware (Section 2), haptic sensations and rendering (Section 3), and user experience and applications (Section 4).

2. Hardware

The majority of mid-air haptic devices currently being used, such as those produced by Ultraleap, are implemented using an array of 256 Murata MA40S4S transducers operating at 40 kHz. These are standard components that are commonly used in applications such as automotive parking and alarm sensors. Other ultrasonic transducers do exist and have been used in multiple mid-air haptic prototypes. However, the MA40S4S is indeed one of the highest performing air-coupled transducers currently on the market. It is a 10 mm barrel transducer that reliably achieves a pressure of ~114 dB SPL (10 Pa RMS) at 30 cm when driven with a 0-20 V square wave. Despite its performance, there are multiple issues with this and similar transducer devices that can inhibit the commercial exploitation of the technology.

Firstly, they are relatively expensive. Air-coupled ultrasound transducers currently cost between $1-3 each (RRP) when ordered in bulk (depending on the brand), with additional driver costs between $0.10-0.30 per channel for each transducer. This results in a per-channel cost of ~$1.2-3.2. For a standard 256-element array, this is a cost of ~$300-800 for the transducers and drivers alone. A reduction in transducer cost would therefore have a significant effect on the commercial opportunities relating to airborne ultrasound technologies in general, lowering the entry cost and making this technology more accessible to new application areas. It should be also emphasized that the current transducers employed for mid-air haptic applications, were originally developed with single-element applications in mind and not the multi-element phased array systems needed here.

The second transducer issue is their size. Most 40 kHz transducers are large both in terms of area and volume, each occupying a 10 mm diameter barrel cylinder and 10 mm in height. A reduction in transducer area and volume would significantly improve their integration options inside or on the surface of products, e.g., an automotive dashboard. As noted in previous chapters, a mid-air haptic device needs to be above a certain minimum area (aperture) to enable focusing acoustic pressure to a point within a reasonable mid-air interaction range, e.g., 10-50 cm. Modular solutions have been proposed to that effect allowing for some product design flexibility.

Another issue with ultrasound transducer technologies used in mid-air haptic applications is their carrier frequency of 40 kHz. This is above the range of human hearing but is not above the hearing range of some animals, including cats and dogs. Higher carrier frequency devices have been presented in a few prototypes [4], however, it is not well understood if higher frequencies have any significant positive or negative effects on the haptic sensation itself. For example, higher frequency transducers would produce a smaller, sharper, focal point. That could be good (e.g., higher resolution rendering), or bad (e.g., weaker haptics), and should be studied in depth before new hardware is proposed.

The last two transducer-related issues are power consumption and heat dissipation. The Murata devices for example consume ~250 mW when driven at maximum strength with a 0-20V 50% duty cycle square wave. This results in a standard 256-element array having a power consumption in excess of 64 W in the maximum output scenario. This is not an insignificant amount of power and may cause issues with some consumer applications, e.g., limiting the potential for portable devices with integrated power units. In consequence, the heat generated by this power is also significant. Specifically, most 40 kHz transducers are quite inefficient as less than 10% of the consumed power is converted to ultrasound. The remainder of the power consumed is wasted as heat, either in the transducer or the driver circuit. Therefore, a key challenge is managing the heat dissipation through cooling units, software control, or new materials.

Advances in transducer technology would certainly mitigate many of the above problems. Until then, however, a good approach is to simply reduce the number of transducers to the minimum needed by the application at hand. The 256 element array employed by Ultraleap development kits produced is overkill for many applications. Meaningful mid-air haptic sensations and interactions can be generated with as little as 64 transducers, therefore, slashing costs, power, and heat simultaneously.

The size of the array also directly affects the renderable workspace, which spans the region directly above the device. For example, the Ultraleap Stratos Explore device (16x16 transducers array) features a workspace of 0.055 m3, shaped like an ellipsoid spanning from 40 mm to 700 mm above the array, with a maximum lateral radius of 320 mm [5]. Providing haptic feedback in a bigger workspace is of course possible by using larger or multiple ultrasound arrays, but - depending on the application at hand - other approaches might be equally effective. A notable example is that of Sand et al. [6], who mounted an ultrasound array onto a head-mounted display so that the renderable workspace moves together with the user. More recently, Howard et al. [5] presented a 2-degrees-of-freedom robot for rotating a 16x16 ultrasound array around the pan and tilt axes, achieving a 14-fold increase in workspace volume with respect to using the array in its standard static position. Innovative approaches such as these ones can significantly reduce the cost of the haptic system with respect to the renderable workspace.

Other hardware-related opportunities that we foresee include the enhancement of mid-air haptic devices with additional sensors and networking capabilities. Already, the technology is directly coupled with independent hand-tracking sensors, however, there are others that could lead to improved operation of the haptics device, such as microphones, thermometers, IR, and humidity sensors, to name a few. Finally, the endowment of mid-air haptic devices with WiFi or other wireless communication modules could enable over-the-air synchronization and distributed operation of multiple devices, potentially opening up a number of new commercial opportunities and applications.

Finally, most mid-air haptic devices look like engineering prototypes, at best enclosed inside a black box with a perforated grid cover to hide their inner workings whilst still allowing ultrasound propagation. To date, little attention has been given towards product design and product integration requirements. Mid-air haptic devices are for the most part not intended to be stand-alone products. Rather, they are likely to offer the most value to the user experience when integrated inside an automotive dashboard, under a digital signage monitor, or as a VR table-top accessory, etc. As noted in previous chapters, the location and orientation of mid-air haptic devices have a direct impact on the quality of the haptic sensation. Therefore, any hardware product design effort must also consider system performance as well as ergonomics, UX, aesthetics and limit any interactions with nearby sensors, e.g., how the acoustic fields could influence the performance of nearby microphones and earpieces, or how any electromagnetic fields on the driver circuit boards can interfere with nearby electrical and electronic equipment. Hayward et al. [7] have for example have proposed enclosing the haptic device in a Faraday cage to enable its use with sensitive neurological monitoring devices.

3. Haptic Sensations and Rendering Algorithms

In analogy to 2D and 3D graphical rendering, mid-air haptic rendering relies on spatial and temporal modulation techniques that change properties of the acoustic focus, so as to create vibrations on the skin that imbue tactile properties and characteristics. Initially, the acoustic focus, which forms the basis for mid-air haptics, was amplitude modulated (AM) to create a localized vibrotactile sensation, while later the focus was moved around in space to create small lateral modulations (LM) [8] or to trace out larger tactile shapes using so-called Spatio-temporal modulation (STM) [9]. Techniques that use acoustic holography [10], multiple focal points [11], or a blend of AM and STM [12] have also emerged and appear to be more suitable at delivering different haptic sensations in different settings. However, none of these modulation techniques adapt or take into consideration the heterogeneity of the human skin, the density, and types of mechanoreceptors being targeted, nor any effects of wave interference on the skin surface [13]. Also, no mid-air haptic demonstrators and prototypes that we are aware of have ever been tailored to a particular user group demographic or person, despite the great differences and preferences displayed by end-users. Mid-air haptics, unlike many other haptic technologies, has the customization potential to be able to address many of these limitations. Beyond biophysical heterogeneity across users, there is also output heterogeneity across the interaction zone. For instance, Raza et al. [14] proposed an algorithm that tuned parameters such as intensity and AM frequency so that the haptic sensation was consistent within the interaction volume. To that end, we call for more research to investigate the perception of ultrasound haptic sensations and explore ways these can be improved through more sophisticated rendering methods.

Psychophysics quantitatively investigates the relationship between physical stimuli and the sensations and perceptions they produce. In the case of mid-air haptics, this presents a great challenge since the possible stimuli cover a very large parameter space. Stimuli can vary in size, shape, intensity, target location, temporal and spatial waveform just to name a few dimensions. Further, mapping all these stimuli parameters onto the perceptual space they relate to is a grand and taxing challenge, especially since the latter combines both functional and non-functional characteristics. What we mean by that is that different haptic stimuli can lead to low-level sensations such as perceptual and two-point discrimination thresholds, mid-level haptic properties such as roughness and curvature, and high-level haptic experiences such as valence and sense of agency. The latter is especially important ever since the discovery that mid-air haptics can be used to target non-glabrous (hairy) skin with the possibility of inducing affective haptics in social touch applications [15]. Here, it is also worth mentioning how Frier et al. [16] showed how perceptual results can be used to optimize haptic performance by using fewer resources than the capabilities of state-of-the-art hardware. Indeed, sometimes, in haptics "less is more" [17], meaning that we can achieve higher performance or improved sensations by providing less rich feedback through simpler devices. While counter-intuitive, such results are well-known in the research community and they are at the core of the great popularity of many cutaneous haptics solutions, such as wearable haptic interfaces [18]. Indeed, cutaneous haptics has already been successfully employed in many high-impact scenarios, such as medical robotics, industrial remote manipulation, and micro-robotic assembly. However, for this approach to be successful, it is necessary to know which are the most important stimuli and sensations to deliver, so as to focus the limited actuation capabilities of this technology where it counts most; hence the importance of studying the perceptual aspects of this technology and the needs of the application at hand. An important forward step to that end is the availability of easy-to-use software tools that can facilitate psychophysical studies, enabling researchers who are not experts in the technology to design and output consistent mid-air haptic stimuli. While there exist some attempts in this direction [19], experimental platforms and frameworks for running perceptual studies on ultrasound mid-air haptics are still rare.

This section has so far reflected on the challenges of using perceptual knowledge to inform the design and presentation of haptic sensations. Closely related to this is the challenge of implementing such designs, so that they can be evaluated and deployed in real usage scenarios. As an emerging technology, there are only a limited number of tools to support the design and implementation of haptic sensations. The Ultraleap Sensation Editor and Controls Suite are a few examples from which designers can select and customize a limited number of templates. Better software support is needed to allow haptic designers to be more expressive and to enable them to explore novel and bespoke haptic designs. Such “no-code” and “easy-to-use” software tools will help the technology become more accessible to a wider audience of designers, practitioners, makers, and academics in other disciplines.

Another aspect that has seen rapid developments in recent years is that of ultrasound field computation and display. Early mid-air haptic prototypes were capable of individually controlling the phase and amplitude of hundreds of transducers, to focus waves at different 3D positions in space. However, these early devices were limited to just one focus point, had low resolution, and limited refresh rates (i.e., the rate at which transducer values could be updated). Since then, we have seen a rapid increase in phase and amplitude resolution thus improving the granularity of the resulting target field and the number of focus points that can be generated simultaneously. There has also been a rapid increase in computational solving speed thus improving the refresh rate at which focal point properties can be changed, in turn enabling the rendering of more advanced haptic sensations such as multi-point STM [20] and PRO-STM [21]. Higher refresh rates allow for smoother transitions between fields thus influencing the amount of unwanted audible sounds, a byproduct of abruptly changing acoustic fields [22]. All this progress has been a result of the adaptation of optical holography techniques [10], efficient eigenproblem solving operations [23], the use of Gerchberg-Saxton (GS) type of iterative phase retrieval algorithms [20], and their efficient implementation on GPUs or FPGAs. In the near future, we expect to see algorithmic advancements and extensions that are closely coupled to advancements in transducer technology as discussed in the previous section. The opportunity here is for improved acoustic holography and solver techniques that can produce sharper focus states while also reducing grating lobe and acoustic streaming phenomena, and also unlocking new spatial frequencies, and thus potentially enhancing the resulting haptic sensations.

Two major limitations of all these solvers is that they are not time-domain accurate, and are not aware of any scattering effects caused by users, i.e., they assume a free field where sound waves are not reflected or scattered off any objects or obstacles. The first problem becomes important only when a focal point moves extremely quickly and the distance between the focus and the ultrasound sources is large, resulting in acoustic interference and aliasing effects to manifest due to heterogeneous speed of sound delays between sound field frames; a kind of acoustic field inertia. Therefore, solvers should not assume that changes to the acoustic field happen instantaneously, and should therefore take into consideration the history of pressure fields preceding new update frames (memory), and if possible optimize for future ones (forecast). The second problem becomes important when complex objects enter the interaction region of the ultrasound display. While reflections of ultrasonic waves off of flat surfaces have been leveraged before to create tactile holograms and illusions [24], partial obstruction (e.g., by a car gear-stick) or indeed the ultrasonic scattering off the user's hands have not been considered in-depth, with the expectation of the work by Inoue et al. [25]. Often, we assume that the user's hand being targeted by a focal point is flat and homogeneous and thus the impinging ultrasound experiences a single specular reflection. However, many mid-air interactions described in this book include gestures like pinching, grasping, and pointing. Here, ultrasound may bounce around the semi-clenched palm inducing all sorts of unexpected tactile sensations. Including such calculations into the acoustic solver requires a closer coupling with the hand-tracking system and a modified Huygens model of acoustic propagation.

Another limitation with existing solvers is that there is no feedback control with regards to the state of the ultrasound transducer board or the environment in which it is operating in. For example, transducers may experience a shift in their resonance frequency of up to a few kHz or a complete phase inversion due to small changes in temperature or humidity levels. Including onboard sensor feedback and solver flexibility can account for such dynamic changes and thus improve the haptic performance of the technology significantly.

Finally, and accentuated by the 2020-2022 global chip shortage crisis that has sent semiconductor component costs rocketing, we can identify the need for modular and scalable hardware platforms, as well as simplified and much more efficient solvers that can be implemented on smaller and cheaper processing units or on the cloud. In summary, there are many gains to be made in exploring more sophisticated acoustic solvers and more efficient hardware implementations, towards the goal of creating higher quality mid-air haptic experiences.

4. User Experience and Applications

A holistic approach to UX is required to advance the applicability and design of mid-air haptics and how they can be tailored to particular applications. While we know that the underlying technology can deliver value by increasing interface usability, improving gesture learning and recall, reducing cognitive load, enhancing a sense of agency, reducing visual distraction, supporting error recovery, etc., we do not know which of these to prioritize and optimize for each target application (automotive, AR/VR gaming, training and simulation, public displays, etc.). Therefore, we call for more studies that develop non-singular prototype systems, starting from the ground up, and that do not just demonstrate a specific capability or function but rather deliver value and improved UX throughout. Frameworks to do so have been presented by Kim et al. [26] for generic haptic experience design, however, expanding and further specializing those towards mid-air haptics comes with its own challenges due to the large haptic design space associated with gesture input and rendering techniques.

Another obstacle we foresee here is the lack of appropriate publication venues being targeted by mid-air haptic researchers. Namely, most published papers on mid-air haptics have to date been addressing the haptics, acoustics, and HCI communities, which to a first approximation reward novelty and methodology. We, therefore, call for more cross-disciplinary collaborations that produce case-study results that present and discuss how different solutions were thought up, designed, developed, and evaluated. Two examples to that end include the works by O'Conaill [27] and Young et al. [28], where the authors give more emphasis on the project delivery process, requirement considerations, and the resulting added UX value, rather than just highlighting a singular contribution or incremental improvement. Such explorations would benefit considerably towards the quest for a "killer application" for mid-air haptic technology.

Another largely unexplored aspect of mid-air haptics is how this technology interacts with other sensory and technological modalities. In terms of the multi-sensory aspect, very few studies have investigated the interplay between mid-air touch and visuals, sounds, and smells. We know, for example, that visuo-haptic feedback can improve the precision of mid-air grasping operations of virtual objects [29], while the right audio-haptic composition can improve the experience of a holographic light-switch button [30]. More such studies are required if we are to imbue unobtrusive touch sensations to the emergent Metaverse paradigm but also to more near-term applications such as automotive human-machine interfaces and digital signage displays.

In terms of technological interactions and integrations, very few studies have coupled mid-air haptics to other novel sensors or data streams. For example, most realizations of the technology take hand and finger positional data as their main input and output a pre-calibrated mid-air tactile sensation in return. The pre-calibrated tactile sensation typically represents some properties of the widget or object being interacted with, e.g., its location, size, shape, state, function, or texture. Instead, Romanus et al. [31] have included to the input stream the user's heart rate as measured by a wearable sensor; the heart rate modulates the presented tactile sensation in real-time. In the age of the Internet of Things (IoT), where smart sensors are prolific and produce rich data streams about our lives, homes, and daily objects we interact with, we see an opportunity to enrich the mid-air haptic effect, both in space, time and form as to represent additional properties such as the weather, the flow of incoming traffic, the urgency of an incoming email.

Multi-haptic explorations have also been very limited. With the exception of works by Ochai et al. [32] and Fan et al. [33] where femtosecond-laser light fields and cable-driven force-feedback, respectively, were combined with ultrasonic acoustic fields to produce novel haptic sensations, very few other examples of multi-haptic interfaces exist. We would therefore like to invite the exploration of mixed haptic interfaces where ultrasound mid-air haptics are utilized together with other contact (wearables, tangibles, grounded, surfaces, etc.) or non-contact (infrared, laser, air jets, electrostatic, etc.) haptic stimuli. These novel prototypes could potentially induce new touch sensations and enable new collaboration synergies between academic groups or commercial exploitation opportunities between industry players.

Considering that touch is a fundamentally complex sense, we believe that in order to advance the technology from CRT, to SD, to HD, to 4K (in analogy to visual displays) we need to first deepen our understanding of the whole mid-air haptic pipeline, from hardware to perception to human experience. That is, how do the acoustic vibrations emitted by a device interact with human perceptual systems and, in turn, become meaningful tactile sensations that encode information, or cause an affective or emotive response? Obrist et al. [34] for example demonstrated a non-arbitrary mapping between emotions and different haptic descriptions (e.g., patterns and frequencies) pointing towards a massive gap in our ability to model and design sensations in a predictable and controlled manner. To make progress in this grand challenge, we advocate for the need to build links between acoustic models of ultrasound vibrations, finite element models of skin vibrations and mechanoreceptor firings, neurocognitive models of low-level tactile thresholds and high-level experiences, and finally connect these to a user-facing interface or application. Implicit in all this is the need for accurate, reliable, and open datasets that can be used to build and validate such models. Particularly, where there is data, there is also the enticing prospect of leveraging artificial intelligence (AI) methods to create predictive and generative haptic design tools and enhanced applications. In this, we highlight the potential for valuable contributions from different disciplines and the benefits of a more multi-disciplinary approach to research on this topic.

Finally, while the scientific community of mid-air haptics is active, international, interdisciplinary, and growing, the industrial community is rather localized and represented by mainly Ultraleap, who is the majority intellectual property (IP) holder, and their direct partnerships with Original Equipment Manufacturers (OEMs) like Hosiden or Bosch who focus on integrating and assembling application-oriented products for a specific market, e.g., automotive human-machine interfaces (HMIs). To that end, we hope to see more companies, both small startups, and large OEMs, reach out to academic expert groups (directly or via the haptics and HCI communities) with the aim to engage with this enabling technology and explore its commercial use in a larger variety of end-user applications, for example, in accessibility, automotive, sterile interfaces in medical and public spaces, AR/VR training, robotics, gaming, the arts and immersive experiences, wellness, and e-commerce or product showcasing. We hope that this book can help facilitate such explorations.

5. Conclusions and Perspectives

Digital haptic interfaces have been around for several decades, using diverse form factors and actuation mechanisms to enrich the tactile experience of user interfaces across many application areas. However, except for some notable exceptions (e.g., vibrations on smartphones and game controllers), haptic technologies have yet to reach the diffusion we believe they deserve. The sense of touch is indeed a fundamental part of human experience, enabling and defining how we interact with the world around us - it is impossible to imagine a true immersive interaction without appealing to our most visceral sense. The timeliness for integrating haptics into digital or remote experiences has gained accelerated momentum due to the COVID-19 pandemic, which has forced individuals, businesses, and governments to shift activities from the real to the digital at a rapid pace. However, most of these digital activities are restricted to the visual and auditory modalities, severely limiting the immersiveness and richness of the targeted interactive experience. Nonetheless, the way we teach, learn, socialize, and play involves all our senses, well beyond what we receive during a typical video conference call. What became apparent through this real-to-digital substitution exercise is that audio-visual media were triumphant, while other senses were absent and truly missed. In this respect, ultrasound mid-air haptic technology can play a revolutionary role. Being able to provide compelling and rich haptic sensations without having to touch or wear anything can take haptic technology to the next level, finally able to reach the popularity of audio and visual feedback modalities.

In this chapter we reflected on our own experiences and the many viewpoints presented in this book, to identify a number of promising challenges and opportunities that we think should be addressed to help mature mid-air haptic technology so that its potential can be realized, e.g., hardware and software improvements, new haptic rendering methods, a deeper understanding of UX, and more application-relevant prototype explorations. However, these directions of inquiry are mostly intrinsic to the technology and the scientific and industrial communities involved. Extrinsic factors such as the ethics and standardization efforts of mid-air touch should also be addressed as they hold great potential in accelerating its trajectory. For example, given that touch is a very intimate and expressive sense which is crucial to our development and wellbeing [35], mid-air haptics researchers and practitioners should anticipate, reflect on, engage with and act upon possible negative societal or environmental impacts of the development of these technologies (e.g., by following a reflective framework for responsible innovation like AREA). Being proactive about engaging with such issues can help with the successful rollout of new technologies, by increasing peoples' willingness to accept and adopt them. As an example of why this is important, consider the issues encountered by early attempts to commercialize augmented reality glasses. Many people perceived this technology as being 'not for them' and were concerned about potential privacy violations. Such issues largely stemmed from poor communication and engagement with the communities where these devices were deployed. Stakeholders in haptics are encouraged to engage with ethical and social concerns around technologies like this. For example, Jewitt et al. [36] have presented a manifesto of 10 statements that are aimed to help haptic designers and developers involved in developing digital touch technologies. This could be further specialized to touchless and mid-air haptic solutions to provide a framework for the responsible development and deployment of mid-air haptic interfaces.

Finally, like with any new technology, standards can ensure that mid-air haptic devices and their induced sensations are consistent in terms of quality and operation and that they adhere to any relevant health and safety considerations. Currently, different prototypes produce different acoustic fields, through a different Application Programming Interface (API), using different electronic components, thus inducing a different haptic sensation. All this uncertainty makes it difficult to compare, reproduce, and reuse, therefore hindering the advancement of the technology, its interoperability, and compatibility with other systems and platforms. To avoid a low-growth and divided future, the mid-air haptics community needs to align on terminology, and recommended practices across all levels of the haptic stack.

Ultrasound haptic technology has advanced considerably over the past decade, accompanied by the emergence of new interaction paradigms, new hardware innovations, and new knowledge about our sense of touch and the role it plays in interactive experiences. These new insights have helped establish new areas of haptics and HCI research and have led to the growth of a diverse and multi-disciplinary research community. This book is a product of that community. Its chapters reflect on the formative years of mid-air haptic technology and the many achievements that have led to the current state-of-the-art. Some offer a retrospective review of key research insights and aim to inform designers, practitioners, and researchers so that they can make the best use of it. Others look to the future to inspire and inform the crucial next steps in advancing the technology. We hope that you share our enthusiasm for mid-air haptics and our excitement for where it goes next, and we hope that this book will be a useful companion as you too may decide to contribute to its bright future.

Acknowledgments

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreements No 801413 (H-Reality), and 101017746 (Touchless).

References

[1] Hoshi, Takayuki, et al. "Noncontact tactile display based on radiation pressure of airborne ultrasound." IEEE Transactions on Haptics 3.3 (2010): 155-165.

[2] https://www.ultraleap.com/

[3] Rakkolainen, Ismo, et al. "A survey of mid-air ultrasound haptics and its applications." IEEE Transactions on Haptics 14.1 (2020): 2-19.

[4] Ito, Mitsuru, et al. "High spatial resolution midair tactile display using 70 kHz ultrasound." International Conference on Human Haptic Sensing and Touch Enabled Computer Applications. Springer, Cham, 2016.

[5] Howard, Thomas, et al. "PUMAH: Pan-tilt Ultrasound Mid-Air Haptics for larger interaction workspace in virtual reality." *IEEE transactions on haptics* 13.1 (2019): 38-44.

[6] Sand, Antti, et al. "Head-mounted display with mid-air tactile feedback." *Proceedings of the 21st ACM Symposium on Virtual Reality Software and Technology*. 2015.

[7] Hayward, Nick, et al. "A novel ultrasonic haptic device induces touch sensations with potential applications in neuroscience research." *2020 IEEE International Ultrasonics Symposium (IUS)*. IEEE, 2020.

[8] Takahashi, Ryoko, Keisuke Hasegawa, and Hiroyuki Shinoda. "Lateral modulation of midair ultrasound focus for intensified vibrotactile stimuli." International Conference on Human Haptic Sensing and Touch Enabled Computer Applications. Springer, Cham, 2018.

[9] Frier, William, et al. "Using spatiotemporal modulation to draw tactile patterns in mid-air." International Conference on Human Haptic Sensing and Touch Enabled Computer Applications. Springer, Cham, 2018.

[10] Morales, Rafael, et al. "Generating Airborne Ultrasonic Amplitude Patterns Using an Open Hardware Phased Array." Applied Sciences 11.7 (2021): 2981.

[11] Carter, Tom, et al. "UltraHaptics: multi-point mid-air haptic feedback for touch surfaces." Proceedings of the 26th annual ACM symposium on User interface software and technology. 2013.

[12] Hajas, Daniel, et al. "Mid-air haptic rendering of 2D geometric shapes with a dynamic tactile pointer." IEEE transactions on haptics 13.4 (2020): 806-817.

[13] Frier, William, et al. "Simulating airborne ultrasound vibrations in human skin for haptic applications." IEEE Access (2022).

[14] Raza, Ahsan, et al. "Perceptually correct haptic rendering in mid-air using ultrasound phased array." IEEE Transactions on Industrial Electronics 67.1 (2019): 736-745.

[15] Pittera, Dario, et al. "" I can feel it coming in the hairs tonight": Characterising Mid-Air Haptics on the Hairy Parts of the Skin." IEEE Transactions on Haptics (2021).

[16] Frier, William, et al. "Sampling strategy for ultrasonic mid-air haptics." Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems. 2019.

[17] Berger, Christopher C., et al. "The uncanny valley of haptics." Science Robotics 3.17 (2018): Art-No.

[18] Pacchierotti, Claudio, et al. "Wearable haptic systems for the fingertip and the hand: taxonomy, review, and perspectives." IEEE transactions on haptics 10.4 (2017): 580-600.

[19] Mulot, Lendy, et al. "DOLPHIN: A Framework for the Design and Perceptual Evaluation of Ultrasound Mid-Air Haptic Stimuli." ACM Symposium on Applied Perception 2021. 2021.

[20] Plasencia, Diego Martinez, et al. "GS-PAT: high-speed multi-point sound-fields for phased arrays of transducers." *ACM Transactions on Graphics (TOG)* 39.4 (2020): 138-1.

[21] Barreiro, Hector, Stephen Sinclair, and Miguel A. Otaduy. "Path Routing Optimization for STM Ultrasound Rendering." *IEEE transactions on haptics* 13.1 (2020): 45-51.

[22] Suzuki, Shun, et al. "Reducing amplitude fluctuation by gradual phase shift in midair ultrasound haptics." IEEE transactions on haptics 13.1 (2020): 87-93.

[23] Long, Benjamin, et al. "Rendering volumetric haptic shapes in mid-air using ultrasound." *ACM Transactions on Graphics (TOG)* 33.6 (2014): 1-10.

[24] Monnai, Yasuaki, et al. "HaptoMime: mid-air haptic interaction with a floating virtual screen." *Proceedings of the 27th annual ACM symposium on User interface software and technology*. 2014.

[25] Inoue, Seki, Yasutoshi Makino, and Hiroyuki Shinoda. "Mid-air ultrasonic pressure control on skin by adaptive focusing." *International Conference on Human Haptic Sensing and Touch Enabled Computer Applications*. Springer, Cham, 2016.

[26] Kim, Erin, and Oliver Schneider. "Defining Haptic Experience: Foundations for Understanding, Communicating, and Evaluating HX." *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 2020.

[27] O'Conaill, Brid, et al. "Improving immersive experiences for visitors with sensory impairments to the aquarium of the pacific." *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*. 2020.

[28] Young, Gareth, et al. "Designing mid-air haptic gesture controlled user interfaces for cars." *Proceedings of the ACM on Human-Computer Interaction* 4. EICS (2020): 1-23.

[29] Frutos-Pascual, Maite, et al. "Evaluation of ultrasound haptics as a supplementary feedback cue for grasping in virtual environments." *2019 International Conference on Multimodal Interaction*. 2019.

[30] Ozkul, Cisem, David Geerts, and Isa Rutten. "Combining Auditory and Mid-Air Haptic Feedback for a Light Switch Button." *Proceedings of the 2020 International Conference on Multimodal Interaction*. 2020.

[31] Romanus, Ted, et al. "Mid-air haptic bio-holograms in mixed reality." *2019 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*. IEEE, 2019.

[32] Ochiai, Yoichi, et al. "Cross-field aerial haptics: Rendering haptic feedback in air with light and acoustic fields." *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 2016.

[33] Fan, Liqiang, Aiguo Song, and Haochen Zhang. "Haptic Interface Device Using Cable Tension Based on Ultrasonic Phased Array." *IEEE Access* 8 (2020): 162880-162891.

[34] Obrist, Marianna, et al. "Emotions mediated through mid-air haptics." *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. 2015.

[35] Field, Tiffany. "*Touch*". 2nd Ed. Cambridge, Massachusetts, US: The MIT Press. 2014.

[36] Jewitt, Carey, et al. "Manifesto for Digital Social Touch in Crisis." *Frontiers in Computer Science* (2021): 97.