

DRIVING TECHNIQUES FOR PHASED-ARRAY SYSTEMS

RELATED APPLICATIONS

[0001] This application claims the benefit of the following eight U.S. Provisional Patent Applications, all of which are incorporated by reference in their entirety:

[0002] 1. Serial No. 62/433,544, filed on December 13, 2016.

[0003] 2. Serial No. 62/436,457, filed on December 20, 2016.

[0004] 3. Serial No. 62/437,791, filed on December 22, 2016.

[0005] 4. Serial No. 62/438,512, filed on December 23, 2016.

[0006] 5. Serial No. 62/594,574, filed on December 5, 2017.

[0007] 6. Serial No. 62/594,687, filed on December 5, 2017.

[0008] 7. Serial No. 62/594,836, filed on December 5, 2017.

[0009] 8. Serial No. 62/594,874, filed on December 5, 2017.

FIELD OF THE DISCLOSURE

[0010] The present disclosure relates generally to an improved technique for driving phased array systems, primarily to produce designed acoustic fields for the creation of mid-air haptics, parametric audio, acoustic levitation and acoustic imaging.

BACKGROUND

[0011] A continuous distribution of sound energy, which we will refer to as an “acoustic field”, can be used for a range of applications including haptic feedback in mid-air.

[0012] In this acoustic field, one or more control points can be defined. These control points can be amplitude modulated with a signal and as a result produce vibrotactile feedback in mid-air. An alternative method to produce feedback is to create control points that are not modulated in amplitude and move them around spatially to create spatio-temporal modulation that can be felt.

[0013] The minimum size and maximum power of the control points created by the device are dependent on the array layout, geometry and relative location of the control point or points to be created. Due to this, it may not be possible to match the behavior of a control point in one location to the behavior in another. Further complications may arise when a control point is moved dynamically through space. The resultant changes in the haptic function of the point experienced by a user are undesirable and often counter-intuitive,

resulting in confusion and loss of experiential immersion. A system that is capable of mitigating the effect of the changes in the air and that is capable of reversing it to provide a consistent haptic experience would be therefore commercially valuable.

[0014] Further, it has been shown that a general non-linear optimization system can be used to produce trap points in the air. But the mechanism used for this purpose has a number of flaws, including but not limited to, not evaluating a potential field appropriate to levitation above phased array systems, using an unnecessarily expensive and time consuming optimization process and being limited only to phase manipulation.

[0015] Addressing these issues is important, as the traps created by existing systems are not as strong as they could be and fail to scale correctly due to these drawbacks. This makes creating systems that exploit acoustic levitation to provide further visuals for furthering multi-modal approaches and alternative uses for haptic systems difficult.

[0016] Further, by defining one or more control points in space, the acoustic field can be controlled. Each point can be assigned a value equating to a desired amplitude at the control point. A physical set of transducers can then be controlled to create an acoustic field exhibiting the desired amplitude at the control points.

[0017] Due to the requirements of existing systems, the natural directional components of acoustic phased arrays have received little attention. For instance, high frequency therapeutic medical systems where the object is to focus energy into a small area is sufficiently approximated by scalar pressure (the change in the number of participating air molecules) as to determine to optimal solution to inject energy, therefore not requiring a vector solution. Since to focus a phased array vector behavior (describing the direction of motion of the air molecules) can be ignored in the majority of cases, accomplishing this efficiently has not yet been a priority.

[0018] However, creating haptic feedback in mid-air involves exchanging vibrations with human skin. This relies completely on the surface of contact with the wavefront and with the momentum carried by the wave, which is a result of engagement with the vector behaviour of the wave. Not only this, since the system is focused such that the angle of incidence differs between transducer elements more and more, the assumed linear relationship between momentum and pressure breaks down. It is not possible to ignore the vector component in this system, and so the scalar pressure is insufficient to consistently find a suitable solution to create a mid-air haptic effect. For this reason, building awareness of this momentum effect into a device is of commercial interest.

[0019] Further, by defining one or more control points in space, the acoustic field can be controlled. Each point can be assigned a value equating to a desired amplitude at the control point. A physical set of transducers can then be controlled to create an acoustic field exhibiting the desired amplitude at the control points.

[0020] It has been shown that the calculation of eigenvectors maps to the solution of some very specific purely quadratic constrained problems in matrix algebra. If a matrix could be formed whose eigenvector solved for the control coefficients of transducers given specific problems in mid-air haptics, the solution could form the basis of a device implementation that exhibits simplifications over existing techniques, and so be of value.

[0021] Further, to use a reduced representation method to describe basis functions, first a shorthand must be developed that enables a simple (low entropy) description of a useful function in terms of details of the acoustic field.

[0022] This description would ideally be something that may be accessible from the point of view of a user.

[0023] Having found such an accessible, low entropy representation, there must also be a way to convert it into its equivalent transducer basis (the transducer basis set that describes the basis of transducer acoustic fields of one of potentially many basis functions in a later solution). This would enable a prospective user to design a variety of different structures within the acoustic field, enabling them to build bespoke systems to maximally manipulate energy flow in the field leading to the creation of a system that may use a number of different effects at once, such as focus points, Bessel beams, bottle beams, bending beams etc. These will extend the concept of control points into control regions which encapsulate a richer method of acoustic field control.

[0024] It was also discovered that there are some problems that must be dealt with in order to use particle velocity as a stand-in for wave direction when considering momentum transfer in haptics.

[0025] Further, performing efficient parallel complex-valued multiplication with the same second operand was a necessity for the large quantities of throughput needed to realize a device with performance metrics in line with those outlined in the requirements documentation. Specifically, to compute in parallel one element of the summation that leads to the computation of a row of a necessary optimization matrix would be sufficient to meet requirements, but only if the subsystem were replicated enough times. For the subsystem to be useful therefore, it must be small and efficient. This pressing need for efficiency yielded a

different viewpoint on an old algorithm that gave rise to a new approach to evaluating complex transcendental functions on top of simple multiplication.

[0026] In the wider world, this resulting algorithm is applicable anywhere an element of processing hardware must interact efficiently or in real-time with complex-valued data whose modulus cannot be ignored. This includes, but is not limited to, communications, including GPS, wireless, radio and satellite transmission and reception systems, audio processing, phased array systems including mid-air haptics, time-of-flight imaging and digital interferometry, computation of scattering effects and polarization, etc. This may translate to a broader scope of opportunity for this invention.

[0027] Volder's algorithm, also known as CORDIC (Volder, J. – The CORDIC computing technique – 1959), can multiply unit complex values, but is incapable of including the modulus in the second operand of the calculation without further computation beyond that of the basic CORDIC algorithm.

[0028] A separate algorithm, invented in 1994 that was implemented by the author some years ago for this reason, is known as BKM after the authors Bajard, Kla and Muller, and further can achieve the goal of full complex-complex multiplication (Bajard, J., Kla, S. and Muller, J. – BKM: A new hardware algorithm for complex elementary functions – 1994). However, this algorithm has an involved implementation and thus does not receive as much attention as it otherwise might. The algorithm can achieve full complex-complex multiplication by reducing the complex value operands to their complex logarithm, adding them and then exponentiating to achieve the result.

[0029] There is a more efficient way to achieve this complex-valued multiplication and also complex-valued division, in both serial and parallel, with BKM and with the new algorithm in some circumstances.

[0030] A 1999 addendum to the BKM algorithm addressed an overly complicated implementation of the complex-valued logarithm, simplifying the implementation somewhat. (Bajard, J. and Imbert, L. – Evaluation of complex elementary functions: a new version of BKM – 1999). This BKM method of 1994 and the 1999 addendum forms the algorithm which in this document will be simply termed as BKM.

[0031] While trying to optimize the algorithm for inclusion into real hardware, the idea of a new technique became apparent that prioritizes ease-of-integration into our, or more generally any, existing system over evaluating the mathematical functions as described. This new technique, when it was developed into a proof of concept code, turned out to have a

number of further emergent properties that make it more promising for general adoption of the hardware implementation than the original algorithm. Additionally, if the exact figures requested of the original technique are required, a further scaling may be applied to retrieve these results, with the algorithm as a whole remaining competitive. In either case, the result is a technique with substantial benefits over the state-of-the-art.

[0032] It is known that a quadratic problem with unit norm may be solved by constructing an matrix whose eigenproblem is equivalent to the quadratic problem. A greater than rank-1 quadratic problem matrix may be generated by summing many rank-1 problem matrices and/or extending the null space of an existing problem matrix. By generating a Krylov subspace of this matrix, an optimized sub-set of the implied basis functions may be generated. By using the Krylov subspace as the basis set to a linear system, a correct or approximated evaluation of the basis set using a number of test functions may be achieved.

[0033] Depending on how many eigenvectors are used and from where they are selected, graceful degradation of a complicated set of many basis functions can be achieved, as the power in each subsequent eigenvector is reduced.

[0034] Further, a complex-valued linear system for defining the result of an acoustic field shape takes n basis functions and recombines these into a resulting field that describes a desired configuration of acoustic pressure at the carrier frequency.

[0035] Up until this point, a control point basis function has only ever been used to actuate control points. Described herein is how to generalize to apply a different set of control test functions to the initial control point or region basis functions.

[0036] This is important because it breaks the link between the initial ‘palette’ of behaviors in the acoustic field and the objective function that the linear system uses. For instance, this may be used to create a basis set of Bessel functions, that may then be used to actuate control points.

BRIEF DESCRIPTION OF THE FIGURES

[0037] The accompanying figures, where like reference numerals refer to identical or functionally similar elements throughout the separate views, together with the detailed description below, are incorporated in and form part of the specification, and serve to further illustrate embodiments of concepts that include the claimed invention, and explain various principles and advantages of those embodiments.

[0038] Figures 1 and 2 are haptic waveforms of arbitrary dynamic range that can be processed through a non-linear remapping function to create a new waveform of controlled dynamic range.

[0039] Figures 3-12 are the effective amplitude modulation waveforms that result from a control points in different configurations with different range control techniques applied.

[0040] Figures 13-16 are projections of acoustic pressure that results from the optimization of Gor'kov-related functions.

[0041] Figure 17 is a graph of the same wave with the same sound pressure level, but scenarios in which the effect of the vector components may give rise to a maximal haptic effect, but also no haptic effect.

[0042] Figure 18 are simulations of measurements of an acoustic field generated by a phased array in dB SPL, dB SIL and dB SVL.

[0043] Figure 19 shows control point plots created by solving the quadratic problems.

[0044] Figure 20 shows cuts through an acoustic field at the edge of the near field region parallel to the flat rectilinear transducer array.

[0045] Figure 21 shows a graph of a Bessel function.

[0046] Figures 22-23 show cross-sections of approximations to high order Bessel and Mathieu beams.

[0047] Figures 24-26 show a bending beam, a bottle beam and a partial section of a bottle beam.

[0048] Figure 27-29 show a beam with increased wavelength, a snaking beam and a helical beam.

[0049] Figure 30-31 shows the effect of removing transducers that contribute waves from the direction that is opposite to that of the desired wave.

[0050] Figure 32-34 show the regions of convergence of the logarithm and exponentiation methods without remapping.

[0051] Figure 35 show a convergence graph of the logarithm and exponentiation methods.

[0052] Skilled artisans will appreciate that elements in the figures are illustrated for simplicity and clarity and have not necessarily been drawn to scale. For example, the dimensions of some of the elements in the figures may be exaggerated relative to other elements to help to improve understanding of embodiments of the present invention.

[0053] The apparatus and method components have been represented where appropriate by conventional symbols in the drawings, showing only those specific details that are pertinent to understanding the embodiments of the present invention so as not to obscure the disclosure

with details that will be readily apparent to those of ordinary skill in the art having the benefit of the description herein.

DETAILED DESCRIPTION

[0054] The foregoing techniques, apparatuses and methods may be used in mid-air haptic systems either in combination or separately.

[0055] 1. DYNAMIC RANGE REDUCTION FOR MID-AIR HAPTIC DEVICES

[0056] I. Psychophysical modelling

[0057] The behavior of the system and the haptic output must be calibrated to the way in which humans perceive touch. For instance, in cases where the control point increases in size, the haptic perception is also increased, although the physical amplitude that is created at the control point does not change. Cross-modal effects, such as from audio, serve to enhance haptic feedback if they are correlated. This leads to discrepancies between physically defined control points, necessitating the use of psychophysical modelling steps in order to achieve equivalent perceptions.

[0058] II. Remapping using a psychohaptic model

[0059] Due to the physical limitations of the playback device, an exact waveform may not be reproducible as specified. The waveform may be constructed in a psychohaptically defined perceptual space. Or it may be comprised of intensities and force measurements on a substantially absolute physical scale, which may be a recording of a haptic texture or an arbitrary constructed waveform potentially expressed alongside further metadata. This waveform represents the behavior of the control point given a ‘perfect’ mid-air haptic device. This is then perceptually evaluated using a psychophysical model that can describe and quantify the perceptual aspects of the recreation of the control point. The control point may be moving while recreating this haptic effect. The movement may be included as part of the perceptual evaluation of the haptic effect. This can then be reinterpreted into the space of output available to the instance of a mid-air haptic system constrained by the user location, relative geometries and available power limits wherein the system is configured to create a substantially equivalent control point.

[0060] The psychohaptic model bestows upon the device the ability to determine how the modulated waveform of the control point may be modified in order to provoke substantially the same haptic response as a perfect device. The process of achieving this remapping may also include a psychoacoustic modelling step in order to prevent the resulting haptic waveform from becoming noisy or to select a more audibly acceptable modulation which is

close by in terms of a perceptual haptic space, but has less noise. It is known that many waveforms may be made more haptically powerful by increasing the amount of audible noise. In some situations, this may be noisy enough to be untenable, while in others this may be desirable. A psychoacoustic model may then be used in tandem to provide an audio to haptic trade-off setting that biases the output to make it either quieter or stronger.

[0061] III. Psychohaptic units

[0062] If the waveform is defined in a psychohaptic space, wherein amplitudes are intended to feel equivalent, then a unit must be designed that takes into account the psychohaptic effect of the produced control points in mid-air. The most notable of these is the integration of tactile sensations across the hand. The focal region involved for each control point changes size and shape again as a function of the properties of the physical device, and this must be counteracted in the construction of a psychohaptic unit. Such a psychohaptic unit will change what physical control point amplitude it will evaluate to depending of the size and shape of the focal region, counteracting the effects of the integration of touch sensation over a large area. In this way, this approach allows the same number of psychohaptic units to encode an equivalent sensation, preserving haptic equivalence across devices and physical instantiations.

[0063] IV. Dynamic waveform remapping

[0064] In order to achieve the best mapping into the target device, the complete waveform and its trajectory may be considered. However, due to interactivity and thus low latency being a requirement, only part or even individual samples of the waveform may be available. Although the whole waveform may be remapped into the capabilities of the target device, to do this requires that the waveform must be known beforehand and so can be considered to be a pre-processing step on known waveforms. An alternative is to remap the physical properties of each waveform sample in amplitude, because remapping amplitude samples can be achieved in real-time with interactive waveforms, making it more suitable for the provision of haptic feedback in real-time.

[0065] A variety of functions may be considered to be good candidates to provide a remapping for control point amplitude samples. These include log-linear curves as well as sigmoidal functions, which may convincingly reproduce perceptual sensory input on devices with low dynamic range in other modalities. While these curves have generally not required the specification of upper limits in the mapping function, the dynamic range may be more effectively used given that there is a finite quantity of power available as demonstrated in Figures 1 and 2. Figure 1 shows a haptic waveform of arbitrary dynamic range 100 that can

be processed through a non-linear remapping function 110 to create a new waveform of controlled dynamic range 120. This can be used to take a haptic waveform of arbitrary dynamic range and convert it into a smaller range. As the waveform can be infinite in amplitude, this does not fill the entire dynamic range of the output, resulting in unused output range. In Figure 2, similar to Figure 1, the non-linear function 210 remaps the amplitude waveform 200 into the amplitudes available to the target device 220. Awareness of the extent of the wave enables the remapping function to be tailored exactly to the output range specific to the device. This allows the full range of the device output to be exploited.

[0066] For illustration, a minimalist linear rescaling is used in the following examples to describe how the mappings prevent undesirable behavior in mid-air haptic devices.

[0067] Shown in Figure 3 is a graph 300 where the first step towards achieving the remapping of the waveform into a reproducible perception is to feed data into the solver function that describes the waveform and its limits 300. In Figure 3 the solid line is A'_1 and the dotted line is $A'_{1,range}$. This is the input to the simple waveform remapping function. The y-axis here is in units internal to the device.

[0068] This is achieved by defining the maximum amplitude that needs to be reproduced to emit the waveform without clipping – this is denoted at each sample time to be $A'_{1,range}$. This is then fed into the solver as input alongside the input waveform amplitude, A'_1 .

[0069] Shown in Figure 4 is a graph 400 that shows the result of attempting to reproduce this waveform without applying any changes to the input data. The result is that the waveform becomes clipped due to a lack of available power to recreate the necessary amplitude at the control point. In Figure 4, the solid line is $A'_{1,simulated}$, which is the result of $amplitude(A'_1, \chi_1, T)$ and the result of attempting to output directly the user-specified amplitudes in Figure 1 without regard for the dynamic range of the device.

[0070] Defining the physical configuration of active transducers in the system as $\mathbf{T}$ and a desired control point amplitude A'_1 at position χ_1 , auxiliary capabilities of the solver function for the mid-air haptic system can be used to define the function amplitude $(A'_1, \chi_1, \mathbf{T})$, which yields either the desired carrier amplitude at the control point if the given device is physically capable of providing this at the given time or alternatively the closest amplitude that can be reproduced at the control point position.

[0071] Evaluating $A_{1,limit} := amplitude(A'_{1,range}, \chi_1, \mathbf{T})$ then yields the maximum of the range to which the input amplitude can be mapped at this point in time. This does not need to be precise, as the control point does not move much from sample to sample, the value of

$A_{1,limit}$ from the previous iteration of the control point may be used. This value may also be filtered in order to remove rounding errors and other sources of noise. In this case, the remapping function takes a single sample of the waveform and remaps this from the interval $[0, A'_{1,range})$ onto the interval $[0, A_{1,limit})$. Shown in Figure 5 is a graph 500 that produces a direct map from $A'_1 \in [0, A'_{1,range})$ onto $A_1 \in [0, A_{1,limit})$, going through the remapping function 500. By remapping the amplitude from A'_1 into the range obtained by $A_{1,limit} := \text{amplitude}(A'_{1,range}, \chi_1, T)$ to the value A_1 , the clipping of the waveform is prevented and $A'_{1,simulated} = A_1$.

[0072] V. Remapping for multiple simultaneous control points

[0073] To use the system for multiple control points at once the ranges must be evaluated the same time so that the solver can provide a guarantee that the waveforms can be represented correctly, as in the two sine wave case example in Figure 6. Shown is a graph 600 where two control points are created with different amplitudes A'_1 and A'_2 at different points in time. One is moving towards and the other moving across with respect to the transducer array. They are set up to oscillate in amplitude with two sine waves of different frequencies.

[0074] Unlike a single control point wherein violating the power limitations of the device results in underpowered, clipped waveforms, with multiple control points going over the limitations of the device promotes crosstalk between the points, muddling their haptic effects. Shown in Figure 7 is graph 700 showing an example of two moving control points. Figure 7 shows an attempt to recreate the two control points using the transducer array that leads to complex and unpredictable behavior as different resonances between the two points are created and replaced, but ultimately fail to reach the desired level.

[0075] Shown in Figure 8 is a graph 800 that uses the auxiliary solver evaluation to determine the power limitations present at both points then can be written as

$\{A_{1,limit}, A_{2,limit}\} := \text{amplitude}(A'_{1,range}, A'_{2,range}, \chi_1, \mathbf{T})$, leading to the shown waveform.

Here, one control point is moving toward the device with the potential to become stronger and the other is moving across the face of the device and away with the potential to become weaker. As can be seen from Figure 8, due to the dynamic range limiting generated by the solver solution, their waveforms may then be represented as intended and without crosstalk. Thus Figure 8 shows two points that have their amplitude envelopes reduced to a level that is guaranteed by the solver to be recreated faithfully without error, or with a controlled known level of error. They are the same power because this is the configuration that minimizes the squared amplitude error with the idealized amplitude levels.

[0076] VI. Power redistribution across multiple simultaneous control points

[0077] The amplitude ranges may be set to different values and even be set dynamically through time in order to manage power levels. In this way, manipulating the ranges can determine the proportion of the array output that is concentrated at a given control point. This is demonstrated for a simple maneuvering of power from one control point to another in Figures 9 and 10, but also a more complicated out-of-phase relationship that makes behavioral dynamics between control points such as that demonstrated in Figure 11 and 12 possible.

[0078] Shown in Figure 9 is a graph 900 with input waveforms and ranges for the two input control points. These have different waveform ranges, which also double as hints to allow each point to take more or less of the available power from the device. Shown in Figure 10 is a graph 1000 having the resulting output from the inputs in Figure 9. In this configuration the power produced by the array is moved from one control point onto the other while they are moving away from and towards the array respectively. Due to this, more power can be used than if it were expected that the control points would be of equal amplitude, as was expected for Figures 6, 7 and 8.

[0079] Shown in Figure 11 is a graph 1100 with alternating functions wherein one control point or control point group ramps up where the others die away in amplitude may be modelled effectively by matching the range to the sinusoid. The result of this is that each waveform is reproduced according to the power requirement of each side of the modulated system, while guaranteeing each side of the system has access to full array power. The result is time slicing without the potential for arbitrary output clipping. Shown in Figure 12 is a graph 1200 having the resulting output from the inputs in Figure 11, but without the grouping described. This results in a reduction in output dynamic range because the power budget has been shared unnecessarily between all control points.

[0080] 2. EFFICIENT CREATION OF TRAP POINTS FOR ACOUSTIC LEVITATION**[0081]** I. Optimization objectives for levitating spherical particles

[0082] The optimization problem required for levitating particles centers around the use of the Gor'kov energy potential. This is a scalar function that evaluates the energy of each possible position of a particle in an acoustic field. Finding minima or low energy states at spatial positions corresponds to finding potentially stable trap points. However, there are acoustic effects that the Gor'kov potential does not take into account, such as acoustic streaming wherein the acoustic waves induce bulk movement in the fluid. For this reason, it is preferable to also ensure that the trap does not occur in a region of fluctuating pressure,

which is often determined to be a viable minimum in the energy potential, as these fluctuations are destructive due to the other factors that are not included in the definition of the potential.

[0083] As the direction and magnitude of outside forces such as gravity and acoustic streaming effects cannot be derived from the Gor'kov potential alone, it is often preferable to calculate the Laplacian of the Gor'kov scalar potential. This gives a further degree of control over which axes to prioritize over others in order to generate a more stable arrangement of inwardly pushing forces.

[0084] II. Evaluating the Gor'kov potential

[0085] The Gor'kov potential field, an energy field that enables spherical particles to float due to the existence of energy minima, is described as (in Bruus 2012):

$$U = \frac{4\pi}{3} a^3 \left[\frac{1}{2} (\kappa_0 - \kappa_p) \tilde{p}^2 - \frac{3}{4} \frac{2(\rho_p - \rho_0)\rho_0}{2\rho_p + \rho_0} \|\tilde{\mathbf{u}}\|^2 \right],$$

$$\kappa_0 = \frac{1}{\rho_0 c_0^2},$$

$$\kappa_p = \frac{1}{\rho_p c_p^2},$$

where c_0 is the speed of sound through the acoustic medium, c_p is the speed of sound through the spherical particle, ρ_0 is the density of the acoustic medium, ρ_p is the density of the spherical particle, the left hand term involving time-averaged pressure $\tilde{p}$ is a representation of potential energy and the right hand term involving the particle velocity vector $\tilde{\mathbf{u}}$ is conversely related to the potential energy. As the Gor'kov potential makes the assumption of existing in the active or far-field, these pressure and velocity terms can be exchanged for time-harmonic terms that are in phase. This suggests that the time-harmonic Gor'kov potential field can be rewritten as:

$$U = \frac{4\pi}{3} a^3 \left[\frac{1}{2} (\kappa_0 - \kappa_p) \|\tilde{p} e^{-i\omega t}\|^2 - \frac{3}{4} \frac{2(\rho_p - \rho_0)\rho_0}{2\rho_p + \rho_0} \|\tilde{\mathbf{u}} e^{-i\omega t}\|^2 \right],$$

[0086] Extending this to many sources yields:

$$U = \frac{4\pi}{3} a^3 \left[\frac{1}{2} (\kappa_0 - \kappa_p) \left\| \sum_{q=1}^n \tilde{p}_q e^{-i(\omega t + \varphi_q)} \right\|^2 - \frac{3}{4} \frac{2(\rho_p - \rho_0)\rho_0}{2\rho_p + \rho_0} \left\| \sum_{q=1}^n \tilde{u}_q e^{-i(\omega t + \varphi_q)} \right\|^2 \right],$$

[0087] The problem with this is that the cross-terms effectively mean that this becomes:

$$U = \frac{4\pi}{3} a^3 \left[\frac{1}{2} (\kappa_0 - \kappa_p) \sum_{q=1}^n \tilde{p}_q e^{-i(\omega t + \varphi_q)} \cdot \overline{\sum_{q=1}^n \tilde{p}_q e^{-i(\omega t + \varphi_q)}} - \frac{3}{4} \frac{2(\rho_p - \rho_0)\rho_0}{2\rho_p + \rho_0} \sum_{q=1}^n \tilde{u}_q e^{-i(\omega t + \varphi_q)} \cdot \overline{\sum_{q=1}^n \tilde{u}_q e^{-i(\omega t + \varphi_q)}} \right]$$

resulting in terms that are quadratic and grow in the number of transducers squared.

[0088] III. Optimizing the potential through an eigenproblem

[0089] The potential definition, neglecting time harmonic effects and considering the pressure and particle velocities input into each transducer, can be written as:

$$x^H = [\alpha_1 e^{-i\varphi_1} \quad \dots \quad \alpha_q e^{-i\varphi_q} \quad \dots \quad \alpha_n e^{-i\varphi_n}],$$

$$\begin{aligned}
M_{q,r} = & \sum_{j=1}^m -\frac{4\pi}{3} a^3 \left(\frac{1}{2} (\kappa_0 - \kappa_p) \tilde{p}_{j,q} \tilde{p}_{j,r} \right. \\
& - \frac{3}{4} \frac{2(\rho_p - \rho_0)\rho_0}{2\rho_p + \rho_0} (\tilde{u}_{j,q,x} \tilde{u}_{j,r,x} + \tilde{u}_{j,q,y} \tilde{u}_{j,r,y} \\
& \left. + \tilde{u}_{j,q,z} \tilde{u}_{j,r,z}) \right) \tilde{A}_{jq} \tilde{A}_{jr} e^{+i\varphi_{j0q}} e^{-i\varphi_{j0r}}, \\
U = & \mathbf{x}^H \mathbf{M} \mathbf{x},
\end{aligned}$$

where j is the index of the trap desired, and the extra $\tilde{A}_{jq} \tilde{A}_{jr} e^{+i\varphi_{j0q}} e^{-i\varphi_{j0r}}$ the phase and amplitude changes when moving from transducers q and r to trap point j . Using the constraint $\mathbf{x}^H \mathbf{x} = 1$, to ensure that the transducers are powered to a non-zero level, the optimal complex transducer activation coefficients to achieve a potential well are given by the eigenvector of $\mathbf{M}$ with the maximum eigenvalue.

[0090] This choice of $\mathbf{x}$ maximises the objective function $\mathbf{x}^H \mathbf{M} \mathbf{x}$ which because of the negative sign in the definition of $M_{q,r}$, yields a minimum in the Gor'kov potential field. This is because $\mathbf{x}^H \mathbf{x} = 1$ is the definition of an eigenvector, as well as $\mathbf{x}^H \mathbf{M} \mathbf{x} = \mathbf{x}^H \lambda \mathbf{x}$, given that by definition when λ , an eigenvalue of the matrix, is maximized while $\mathbf{x}^H \mathbf{x} = 1$, this must yield a maximum, if it exists. The standard statement of the problem then becomes:

$$\begin{aligned}
& \text{maximize} && \mathbf{x}^H \mathbf{M} \mathbf{x} \\
& \text{subject to} && \mathbf{x}^H \mathbf{x} = 1, \\
& \text{and} && \mathbf{x} \in \mathbb{C}^n.
\end{aligned}$$

[0091] However, as previously discussed in the literature, minimizing the energy of the Gor'kov field generates a focus and not actually a trap point. To generate a trap point, the pressure fluctuations at the trap are minimized. This is non-trivial, as the definitions of the pressure fluctuations are complex-valued and represent a fluctuating quantity that cannot be meaningfully minimized but brought instead to zero.

[0092] IV. Constructing a null space inside the eigenproblem to generate pressure zeroes

[0093] Zero pressure conditions can be constructed as a series of constraints that evaluate to a zero pressure at given points in the acoustic field. The pressure can only evaluate to a non-zero value at a given location if there exists part of the solution vector that is parallel to the

known optimal pressure focusing solution in the high dimensional phase space. For a trap point j there is a vector $\mathbf{c}_j'^H$ which is:

$$\mathbf{c}_j'^H = [\tilde{A}_{j1}e^{-i\varphi_{j01}} \quad \dots \quad \tilde{A}_{jq}e^{-i\varphi_{j0q}} \quad \dots \quad \tilde{A}_{jn}e^{-i\varphi_{j0n}}]$$

[0094] Therefore, the constraints for zero pressure at the trap point j may be described as the constraining equation:

$$\mathbf{c}_j'^H \mathbf{x} = 0.$$

[0095] As it does not matter what the amplitudes are for this and the constraints need to be unit length, so therefore:

$$\mathbf{c}_j^H = \frac{\mathbf{c}_j'^H}{\sqrt{\mathbf{c}_j'^H \mathbf{c}_j'}}.$$

[0096] The standard statement of the problem is therefore:

$$\begin{aligned} &\text{maximize} && \mathbf{x}^H \mathbf{M} \mathbf{x} \\ &\text{subject to} && \mathbf{x}^H \mathbf{x} = 1, \\ &&& \mathbf{c}_j^H \mathbf{x} = 0, \quad j \in 1, \dots, m, \\ &\text{and} && \mathbf{c}_j^H, \mathbf{x} \in \mathbb{C}^n. \end{aligned}$$

[0097] Taking an eigensystem result from Golub (1973) and restating it for complex-values yields the following argument:

[0098] Given the above system with $m = 1$, a characteristic equation can be written for the system as:

$$\varphi(\mathbf{x}, \lambda, \mu) = \mathbf{x}^H \mathbf{M} \mathbf{x} - \lambda(\mathbf{x}^H \mathbf{x} - 1) + 2\mu \mathbf{x}^H \mathbf{c},$$

where λ and μ are Lagrange multipliers. This can be differentiated by the conjugate of the $\mathbf{x}$ vector, yielding the system:

$$\mathbf{M} \mathbf{x} - \lambda \mathbf{x} + \mu \mathbf{c} = 0,$$

[0099] Multiplying through by $\mathbf{c}^H$ yields:

$$\mathbf{c}^H \mathbf{M} \mathbf{x} - \mathbf{c}^H \lambda \mathbf{x} + \mathbf{c}^H \mu \mathbf{c} = 0,$$

$$\mathbf{c}^H \mathbf{M} \mathbf{x} - \mathbf{c}^H \mathbf{x} \lambda + \mathbf{c}^H \mathbf{c} \mu = 0,$$

[00100] Assuming $\mathbf{c}$ is unit length yields $\mathbf{c}^H \mathbf{c} = 1$, the earlier problem definition yields $\mathbf{c}^H \mathbf{x} = 0$, and substituting both produces:

$$\mu = -\mathbf{c}^H \mathbf{M} \mathbf{x}.$$

[00101] Substituting this back into the earlier derivative yields:

$$\mathbf{M} \mathbf{x} - \mathbf{c}^H \mathbf{M} \mathbf{x} \mathbf{c} = \lambda \mathbf{x},$$

[00102] Which is:

$$(\mathbf{I} - \mathbf{c} \mathbf{c}^H) \mathbf{M} \mathbf{x} = \lambda \mathbf{x},$$

[00103] As this is the statement of the eigensystem, yet $(\mathbf{I} - \mathbf{c} \mathbf{c}^H) \mathbf{M}$ may not be symmetric making eigenvector calculation difficult. However, $\mathbf{I} - \mathbf{c} \mathbf{c}^H$ can be shown to be idempotent $((\mathbf{I} - \mathbf{c} \mathbf{c}^H)^2 = \mathbf{I} - \mathbf{c} \mathbf{c}^H)$ and eigenvalues are unchanged by ordering, yielding:

$$(\mathbf{I} - \mathbf{c} \mathbf{c}^H) \mathbf{M} (\mathbf{I} - \mathbf{c} \mathbf{c}^H) \mathbf{x} = \lambda \mathbf{x},$$

[00104] Extending this to more constraint vectors, there is no problem with further multiplication of the left hand side:

$$(\mathbf{I} - \mathbf{c}_2 \mathbf{c}_2^H) (\mathbf{I} - \mathbf{c}_1 \mathbf{c}_1^H) \mathbf{M} (\mathbf{I} - \mathbf{c}_1 \mathbf{c}_1^H) (\mathbf{I} - \mathbf{c}_2 \mathbf{c}_2^H) \mathbf{x} = \lambda \mathbf{x},$$

Which finally after many multiplications yields:

$$(\mathbf{I} - \mathbf{C} \mathbf{C}^\dagger) \mathbf{M} (\mathbf{I} - \mathbf{C} \mathbf{C}^\dagger) \mathbf{x} = \lambda \mathbf{x},$$

$$\mathbf{C} = [\mathbf{c}_1 \quad \cdots \quad \mathbf{c}_j \quad \cdots \quad \mathbf{c}_m],$$

where $\mathbf{C}^\dagger$ denotes the Moore-Penrose pseudoinverse.

[00105] In summary, the solution to the initial problem can be found by determining the eigenvector with greatest eigenvalue of $(\mathbf{I} - \mathbf{C} \mathbf{C}^\dagger) \mathbf{M} (\mathbf{I} - \mathbf{C} \mathbf{C}^\dagger)$. The multiplies by $\mathbf{I} - \mathbf{C} \mathbf{C}^\dagger$ can be viewed as adding rank-1 changes that add basis vectors to the null space of the solution matrix.

[00106] Finally, it is useful to add support for calculating the Laplacian of the Gor'kov potential in a simple way. The Gor'kov potential solution found herein yields a trap point, but there is no control over from which directions the forces are applied from, and there may even be directions from which there is no force. To achieve this, there is a need to optimize in terms of the most diverging point in the gradient of the potential field and then reweight the gradients of this field. This is then, generally speaking, the Laplacian of the Gor'kov field.

[00107] V. Modifications to generate the Laplacian of the Gor'kov field

[00108] The modifications to minimize the Laplacian of the Gor'kov field are only necessarily applicable to the $\mathbf{M}$ matrix. If the effects of wavefront curvature from each transducer are discarded and regard them as plane waves only, then the singularities that complicate the formulation are removed and particle velocities become constants. The resulting plane waves from each transducer can be rewritten in the form:

$$\phi_q(\chi - \mathbf{O}) = \mathbf{A}e^{-i(2\pi(\hat{n}_x x + \hat{n}_y y + \hat{n}_z z)/\lambda + \varphi)}.$$

[00109] The matrix element can then, in the Gor'kov without Laplacian case, be represented as the equivalent system:

$$\mathbf{M}_{q,r} = \sum_{j=1}^m \mathbf{A}_{j,q,r} e^{-i(2\pi(\hat{n}_{jq,x}x + \hat{n}_{jq,y}y + \hat{n}_{jq,z}z)/\lambda + \varphi)} e^{+i(2\pi(\hat{n}_{jr,x}x + \hat{n}_{jr,y}y + \hat{n}_{jr,z}z)/\lambda + \varphi)}.$$

where all the constants in $\mathbf{M}_{q,r}$ are merged into $\mathbf{A}_{j,q,r}$. Taking the Laplacian of this formulation then yields the modified matrix formulation:

$$\mathbf{M}'_{q,r} = \sum_{j=1}^m \mathbf{A}_{j,q,r} \frac{4\pi^2}{\lambda^2} \|\hat{n}_{jq} - \hat{n}_{jr}\|^2 e^{-i(\frac{2\pi}{\lambda}\hat{n}_{jq}\cdot\chi_j + \varphi_{jq})} e^{+i(\frac{2\pi}{\lambda}\hat{n}_{jr}\cdot\chi_j + \varphi_{jr})}.$$

[00110] Using this value in place of the corresponding components of the $\mathbf{M}$ matrix yields a solution in terms of optimising for the Laplacian of the Gor'kov potential. As the term $\|\hat{n}_{jq} - \hat{n}_{jr}\|^2$ can be rewritten as:

$$\|\hat{n}_{jq} - \hat{n}_{jr}\|^2 = (\hat{n}_{jq,x} - \hat{n}_{jr,x})^2 + (\hat{n}_{jq,y} - \hat{n}_{jr,y})^2 + (\hat{n}_{jq,z} - \hat{n}_{jr,z})^2$$

[00111] It can also be split into contributions from x , y and z directions:

$$\mathbf{w}_x(\hat{n}_{jq,x} - \hat{n}_{jr,x})^2 + \mathbf{w}_y(\hat{n}_{jq,y} - \hat{n}_{jr,y})^2 + \mathbf{w}_z(\hat{n}_{jq,z} - \hat{n}_{jr,z})^2$$

[00112] In this way, weighting can be applied to generate constraint axes and bias the forces generated from each direction.

[00113] VI. Implementation in a device

[00114] The system in this document represents a device that can be used to compute one or more levitation traps based on a simple eigensystem solver. This can potentially be implemented in a device firmware in order to create an embodiment capable of using acoustic levitation to constrain, trap and transport items.

[00115] This is illustrated in Figures 13-16.

[00116] Figure 13 shows a Laplacian of Gor'kov with equal direction weighting. The left side 1310 is the x-y plane parallel to transducer array. The right side 1320 is the x-z plane perpendicular to and slicing through the flat transducer array.

[00117] Figure 14 shows Laplacian of Gor'kov with heavily weighted z-direction. The left side 1410 is the x-y plane parallel to transducer array. The right side 1420 is the x-z plane perpendicular to and slicing through the flat transducer array.

[00118] Figure 15 shows a Gor'kov maximization with zero pressure condition. The left side 1510 is the x-y plane parallel to transducer array, which describes a helical trap as is shown by the non-uniform shade around a ring of uniform intensity. The right side 1520 is the x-z plane perpendicular to and slicing through the flat transducer array underneath.

[00119] In Figure 16 further eigenmodes 1610, 1620, 1630 correspond to trap solutions that also solve the core problem. These are less effective as they do not maximize the objective function. This is because the value of the objective function, the modified Gor'kov potential, is less for these modes that are orthogonal to the dominant eigenvector. The value of the objective function achieved is reflected in the eigenvalues associated with each eigenvector. Eigenvectors associated with smaller eigenvalues represent these alternative less effective solutions.

[00120] 3. PHASED ARRAY DEVICE FOCUSING WITH VECTOR EFFECTS WITH APPLICATIONS TO MID-AIR HAPTICS

[00121] I. Active acoustic intensity

[00122] Active acoustic intensity I_a can be defined using the pressure p and particle velocity vector $\mathbf{u}$ of the acoustic medium. For an acoustic wave, this can be described as:

$$I_a = \frac{1}{2} \Re(\mathbf{p} \cdot \bar{\mathbf{u}}),$$

where the active intensity is a purely real vector quantity. In the near field of a wave source, the reactive field dominates instead $I_r = \frac{1}{2} \Im(\mathbf{p} \cdot \bar{\mathbf{u}})$. However, in the far field, so beyond one or two wavelengths, $I_r \rightarrow 0$. For this reason, I_a will only be considered and assumed to be guaranteed real.

[00123] Considering a monochromatic plane wave source, where:

$$\begin{aligned} \mathbf{p} &= \tilde{\mathbf{p}} e^{-i(\omega t + \varphi)}, \\ \mathbf{u} &= \tilde{\mathbf{u}} e^{-i(\omega t + \varphi)}, \end{aligned}$$

results in:

$$I_a = \frac{1}{2} (\tilde{p} e^{-i(\omega t + \varphi)} \cdot \overline{\tilde{u} e^{-i(\omega t + \varphi)}}),$$

[00124] where $\tilde{p}$ is the pressure amplitude, and $\tilde{u}$ is the particle velocity amplitude in each direction.

[00125] The sound intensity and thus energy flux transmitted by the wave along a given axis is of interest so the dot product to measure the scalar contribution is:

$$I_a \cdot \hat{n} = \frac{1}{2} (\tilde{p} e^{-i(\omega t + \varphi)} \cdot \overline{\tilde{u} e^{-i(\omega t + \varphi)}} \cdot \hat{n})$$

[00126] Due to the form of the acoustic wave equations:

$$\begin{aligned} \frac{\partial^2 p}{\partial t^2} - c_0^2 \nabla^2 p &= 0, \\ \frac{\partial^2 u}{\partial t^2} - c_0^2 \nabla^2 u &= 0, \end{aligned}$$

the superposition principle must hold for both pressure and particle velocity vectors.

Assuming n acoustic sources and approximating each as a plane wave at the point of interest, then yields the equation:

$$I_a \cdot \hat{n} = \frac{1}{2} \left(\sum_{q=1}^n \tilde{p}_q e^{-i(\omega t + \varphi_q)} \cdot \overline{\sum_{q=1}^n \tilde{u}_q e^{-i(\omega t + \varphi_q)}} \cdot \hat{n} \right),$$

which for efficiency it is desirable to solve for the appropriate variable using a linear system.

This means it cannot be solved for the acoustic intensity in a given direction, as it is made up of two independent variables. Existing techniques solve for the complex-valued scalar pressures $\tilde{p}_q e^{-i\varphi_q}$ and then assume that the direction of the each of the individual waves do not matter or are roughly parallel. Instead, solve for the particle velocity vector $\tilde{u}_q e^{-i\varphi_q}$ because this directly influences the momentum (and the kinetic energy), which for the application of haptic feedback in mid-air is important.

[00127] This leads to two potential techniques to extend current solution mechanisms to harness the complex-valued particle velocity vector $\tilde{u} e^{-i\varphi}$ to create solutions that respect the directivity of the waves and their resulting energy transfer.

[00128] II. Solution method for complex-valued pressure

[00129] The solution method for complex-valued pressure involves using an acoustic model for pressure. This model can be queried, given a transducer location alongside a control point location, and the resultant phasor at the point in the field determined. This is a complex pressure, $\alpha_{j0q} = \tilde{p}_{j0q} e^{-i\varphi_{j0q}}$, where j corresponds to the index of the sample point in the acoustic field and q the index numbering of the actuated transducer. This enables the creation of an underdetermined linear system:

$$\begin{bmatrix} \alpha_{101} & \dots & \alpha_{10q} & \dots & \alpha_{10n} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \alpha_{j01} & \dots & \alpha_{j0q} & \dots & \alpha_{j0n} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \alpha_{m01} & \dots & \alpha_{m0q} & \dots & \alpha_{m0n} \end{bmatrix} \begin{bmatrix} x_1 \\ \vdots \\ x_q \\ \vdots \\ x_n \end{bmatrix} = \begin{bmatrix} b_1 \\ \vdots \\ b_j \\ \vdots \\ b_m \end{bmatrix},$$

where $\mathbf{b}$ are the intended phasors in the air at locations $1, \dots, j, \dots, m$ and $\mathbf{x}$ are the transducer activation coefficients from the transducers $1, \dots, q, \dots, n$, the initial phase and amplitude required to create the behaviour in the air that is specified by $\mathbf{b}$.

[00130] However, as $\mathbf{b}$ is complex-valued and contains both phase and amplitude specification, the phase offset can be freely modified. This can be leveraged to find a phase value that opens up the largest possible subspace of amplitudes. This is achieved by finding the principal eigenvector $\mathbf{y}$ of the system, $L\mathbf{y} = \lambda\mathbf{y}$:

$$L = \begin{bmatrix} A_{c_1} & \dots & A_{c_j} k_j (\overline{\alpha_{j0}} \cdot \alpha_{10}) & \dots & A_{c_m} k_m (\overline{\alpha_{m0}} \cdot \alpha_{10}) \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ A_{c_1} k_1 (\overline{\alpha_{10}} \cdot \alpha_{j0}) & \dots & A_{c_j} & \dots & A_{c_m} k_m (\overline{\alpha_{m0}} \cdot \alpha_{j0}) \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ A_{c_1} k_1 (\overline{\alpha_{10}} \cdot \alpha_{m0}) & \dots & A_{c_j} k_j (\overline{\alpha_{j0}} \cdot \alpha_{m0}) & \dots & A_{c_m} \end{bmatrix},$$

[00131] in which A represents amplitudes desired at each point in the acoustic field and k a per column normalization applied that cancels the dot product to one in each diagonal of the matrix, for each control point in the acoustic field $1, \dots, j, \dots, m$. The normalized components of the eigenvector $\mathbf{y}$ of this matrix when multiplied with the initial amplitudes A will then correspond to a phasor vector $\mathbf{b}$ that promotes reinforcement amongst each of the points of pressure in the field.

[00132] Least norm solutions of underdetermined problems can be solved using the form:

$$\mathbf{x}_{leastnorm} = \mathbf{A}^H (\mathbf{A} \mathbf{A}^H)^{-1} \mathbf{b}.$$

[00133] Critical to this solution is the inverse of the matrix $\mathbf{C} = \mathbf{A} \mathbf{A}^H$, which is the first step of solving the system via Cholesky decomposition, and is written:

$$C = \begin{bmatrix} \overline{\alpha_{10}} \cdot \alpha_{10} & \dots & \overline{\alpha_{j0}} \cdot \alpha_{10} & \dots & \overline{\alpha_{m0}} \cdot \alpha_{10} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \overline{\alpha_{10}} \cdot \alpha_{j0} & \dots & \overline{\alpha_{j0}} \cdot \alpha_{j0} & \dots & \overline{\alpha_{m0}} \cdot \alpha_{j0} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \overline{\alpha_{10}} \cdot \alpha_{m0} & \dots & \overline{\alpha_{j0}} \cdot \alpha_{m0} & \dots & \overline{\alpha_{m0}} \cdot \alpha_{m0} \end{bmatrix},$$

meaning it can help derive the eigensystem matrix L as well. However, also working with the system in the form:

$$\mathbf{w}_{focus} = C^{-1} \mathbf{b},$$

yields a very small linear system and subsequent solution size, which can be solved with any technique (QR decomposition, Cholesky factorization, etc.) and kept in this diminutive state until emitted. Neither this C matrix nor this $\mathbf{w}_{focus}$ vector has an index in transducers, until computing:

$$\mathbf{x}_{leastnorm} = A^H \mathbf{w}_{focus}.$$

[00134] It is often worth multiplying through by A^H at the last possible moment to convert back into $\mathbf{x}_{leastnorm}$ transducer activations that multiply the wave emitted from each transducer, as the transducers can be numerous.

[00135] III. Solving for one or more anisotropic control points

[00136] Creating the control points in the acoustic field such that they each have a normal vector or direction specified alongside them is one way to incorporate the momentum requirement into a phased array solver. Giving each control point a normal vector effectively means solving for the x, y and z particle velocities and thus components of the wave momentum at the point of interest, rather than solving for pressure at each control point. To achieve this, we expand the linear system equation from before to include three components for each particle velocity in the place of one for pressure defined previously. These three components, then specified as, $\alpha_{x,j0q}$, $\alpha_{y,j0q}$ and $\alpha_{z,j0q}$, are defined as:

$$\alpha_{x,j0q} = \|\tilde{u}_{j0q}\| e^{-i\varphi_{j0q}} \cdot \hat{\Delta}_{x,j0q} \cdot \hat{n}_{x,j},$$

$$\alpha_{y,j0q} = \|\tilde{u}_{j0q}\| e^{-i\varphi_{j0q}} \cdot \hat{\Delta}_{y,j0q} \cdot \hat{n}_{y,j},$$

$$\alpha_{z,j0q} = \|\tilde{u}_{j0q}\| e^{-i\varphi_{j0q}} \cdot \hat{\Delta}_{z,j0q} \cdot \hat{n}_{z,j},$$

in each case, $\hat{\Delta}_{j0q}$ is the vector denoting the direction of the wave propagation from the transducer q at the control point j , and $\hat{n}_j$ is the normal vector direction of the control point j , across whose position the momentum is measured. The core linear system to solve then becomes:

$$\begin{bmatrix} \alpha_{x,101} & \cdots & \alpha_{x,10q} & \cdots & \alpha_{x,10n} \\ \alpha_{y,101} & \cdots & \alpha_{y,10q} & \cdots & \alpha_{y,10n} \\ \alpha_{z,101} & \cdots & \alpha_{z,10q} & \cdots & \alpha_{z,10n} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \alpha_{x,j01} & \cdots & \alpha_{x,j0q} & \cdots & \alpha_{x,j0n} \\ \alpha_{y,j01} & \cdots & \alpha_{y,j0q} & \cdots & \alpha_{y,j0n} \\ \alpha_{z,j01} & \cdots & \alpha_{z,j0q} & \cdots & \alpha_{z,j0n} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \alpha_{x,m01} & \cdots & \alpha_{x,m0q} & \cdots & \alpha_{x,m0n} \\ \alpha_{y,m01} & \cdots & \alpha_{y,m0q} & \cdots & \alpha_{y,m0n} \\ \alpha_{z,m01} & \cdots & \alpha_{z,m0q} & \cdots & \alpha_{z,m0n} \end{bmatrix} \begin{bmatrix} x_1 \\ \vdots \\ x_q \\ \vdots \\ x_n \end{bmatrix} = \begin{bmatrix} b_{x,1} \\ b_{y,1} \\ b_{z,1} \\ \vdots \\ b_{x,j} \\ b_{y,j} \\ b_{z,j} \\ \vdots \\ b_{x,m} \\ b_{y,m} \\ b_{z,m} \end{bmatrix}$$

which specifies the three components of velocity for particles of the acoustic medium for each control point. However, when using the reduced representation approach to the solution of the system, that is the specification of the $\mathbf{C}$ matrix to reduce the amount of effort required, the number of entries of this matrix increases by a factor of nine, making the problem significantly more intensive to compute. For each element of the matrix $\mathbf{C}$ previously,

$$\overline{\alpha_{k0}} \cdot \alpha_{j0},$$

becomes the submatrix:

$$\begin{bmatrix} \overline{\alpha_{x,k0}} \cdot \alpha_{x,j0} & \overline{\alpha_{y,k0}} \cdot \alpha_{x,j0} & \overline{\alpha_{z,k0}} \cdot \alpha_{x,j0} \\ \overline{\alpha_{x,k0}} \cdot \alpha_{y,j0} & \overline{\alpha_{y,k0}} \cdot \alpha_{y,j0} & \overline{\alpha_{z,k0}} \cdot \alpha_{y,j0} \\ \overline{\alpha_{x,k0}} \cdot \alpha_{z,j0} & \overline{\alpha_{y,k0}} \cdot \alpha_{z,j0} & \overline{\alpha_{z,k0}} \cdot \alpha_{z,j0} \end{bmatrix}.$$

[00137] This remains well behaved when using the existing approach in the eigensystem matrix $\mathbf{L}$, which can still be constructed using the same approach, although the user input amplitudes become $A_{x,j} = \tilde{u}_{x,j}$, $A_{y,j} = \tilde{u}_{y,j}$ and $A_{z,j} = \tilde{u}_{z,j}$ or $\mathbf{A}_j = \|\tilde{\mathbf{u}}_j\| \hat{\mathbf{n}}_j$. The cross terms of the eigensystem matrix can be viewed as performing a measurement on one axis at the position of one control point when considering the impact of adding the wave on another axis

on a further control point. While the $\mathbf{w}_{focus}$ vector now contains three times as many elements, the multiplication by $\mathbf{A}^H$ can take the system back to the transducer activation coefficients required to reproduce the desired momentum vectors in the air above the device. [00138] This solution is particularly interesting for the application of mid-air haptic feedback because it opens the possibility to create waves that approach the hand from different angles. This may enable the acoustic waves to be able to induce shear waves in the skin on the boundary, which opens up an entirely new set of sensations that may be created using mid-air haptic techniques, at the expense of the extra required computation.

[00139] IV. Solving for the particle speed at one or more control points in the acoustic medium

[00140] Another technique to achieve the conversion from a solver that targets pressure values to one that is capable of targeting momentum is to create a linear system solver that works entirely with the acoustic particle speed, which (similarly to the original pressure approach) is also complex-valued due to the monochromatic harmonic time. Further, it is scalar. This can be achieved by calculating the direction of greatest particle velocity at a control point and thus the resultant wave direction $\tilde{\mathbf{n}}_j$, then multiplying through to make each matrix entry a single complex-valued scalar. As the dot product affects all the sampling directions by a scale factor equally across x , y and z , this can be applied as a post-process to the initial matrix construction. However, the direction of the resulting wave must be known before the solver can function and therefore cannot be influenced by the solver. When applied to the construction of mid-air haptic devices, this approach will be most suitable for small devices where mid-air but relatively close range feedback is best and is an extension to current solution mechanisms.

[00141] To achieve this, the acoustic model for each transducer is recast as a problem to determine particle speed in the acoustic medium. The alpha vector thus is redefined to yield a complex-valued particle speed as measured along the direction of greatest motion, $\alpha_{j0q} = \tilde{\mathbf{u}}_{j0q} e^{-i\varphi_{j0q}} \cdot \hat{\mathbf{n}}_j$, as the system that needs solving for a given control point is of the form $(\sum_{q=1}^n \tilde{\mathbf{u}}_{j0q} e^{-i(\omega t + \varphi_{j0q})}) \cdot \hat{\mathbf{n}}_j$. However, unfortunately $\hat{\mathbf{n}}_j$ cannot be known *a priori*, and it is required to construct the linear system matrix:

$$\begin{bmatrix} \alpha_{101} & \dots & \alpha_{10q} & \dots & \alpha_{10n} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \alpha_{j01} & \dots & \alpha_{j0q} & \dots & \alpha_{j0n} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \alpha_{m01} & \dots & \alpha_{m0q} & \dots & \alpha_{m0n} \end{bmatrix} \begin{bmatrix} x_1 \\ \vdots \\ x_q \\ \vdots \\ x_n \end{bmatrix} = \begin{bmatrix} b_1 \\ \vdots \\ b_j \\ \vdots \\ b_m \end{bmatrix}.$$

[00142] Therefore, the acoustic model can be modified to produce the values:

$$\tilde{u}'_{x,j0q} = \|\tilde{u}_{j0q}\| \cdot \hat{\Delta}_{x,j0q},$$

$$\tilde{u}'_{y,j0q} = \|\tilde{u}_{j0q}\| \cdot \hat{\Delta}_{y,j0q},$$

$$\tilde{u}'_{z,j0q} = \|\tilde{u}_{j0q}\| \cdot \hat{\Delta}_{z,j0q},$$

where

$$\alpha_{j0q} = \begin{bmatrix} \tilde{u}'_{x,j0q} \\ \tilde{u}'_{y,j0q} \\ \tilde{u}'_{z,j0q} \end{bmatrix} e^{-i\varphi_{j0q}} \cdot \hat{n}_j.$$

[00143] These values must then be used to construct the dot products of the $\mathbf{C}$ matrix:

$$\mathbf{C} = \begin{bmatrix} \overline{\alpha_{10}} \cdot \alpha_{10} & \dots & \overline{\alpha_{j0}} \cdot \alpha_{10} & \dots & \overline{\alpha_{m0}} \cdot \alpha_{10} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \overline{\alpha_{10}} \cdot \alpha_{j0} & \dots & \overline{\alpha_{j0}} \cdot \alpha_{j0} & \dots & \overline{\alpha_{m0}} \cdot \alpha_{j0} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \overline{\alpha_{10}} \cdot \alpha_{m0} & \dots & \overline{\alpha_{j0}} \cdot \alpha_{m0} & \dots & \overline{\alpha_{m0}} \cdot \alpha_{m0} \end{bmatrix},$$

[00144] But there are two potential routes to achieve this. The dot product required is defined by

$$\overline{\alpha_{k0}} \cdot \alpha_{j0} = \sum_{q=1}^n \left(\begin{bmatrix} \tilde{u}'_{x,k0q} \\ \tilde{u}'_{y,k0q} \\ \tilde{u}'_{z,k0q} \end{bmatrix} e^{i\varphi_{k0q}} \cdot \hat{n}_k \right) \left(\begin{bmatrix} \tilde{u}'_{x,j0q} \\ \tilde{u}'_{y,j0q} \\ \tilde{u}'_{z,j0q} \end{bmatrix} e^{-i\varphi_{j0q}} \cdot \hat{n}_j \right).$$

[00145] This can be written as:

$$\overline{\alpha_{k0}} \cdot \alpha_{j0} = \left(\left(\sum_{q=1}^n \begin{bmatrix} \tilde{u}'_{x,k0q} \\ \tilde{u}'_{y,k0q} \\ \tilde{u}'_{z,k0q} \end{bmatrix} e^{i\varphi_{k0q}} \otimes \begin{bmatrix} \tilde{u}'_{x,j0q} \\ \tilde{u}'_{y,j0q} \\ \tilde{u}'_{z,j0q} \end{bmatrix} e^{-i\varphi_{j0q}} \right) \cdot \hat{n}_k \right) \cdot \hat{n}_j$$

where $\otimes$ is the dyadic product. This requires nine separate summations to be maintained per matrix element (three for the diagonals) in order to compute the normal vectors afterwards.

This is beneficial in that it results in an initial matrix configuration for $\mathbf{C}$ that has many similarities to that of the anisotropic control point approach, so therefore may allow for common implementation infrastructure that supports both techniques. However, in many implementations, the alternative of having a two stage modelling system instead may be more desirable than holding these nine separate summation variables. In such a system, the first stage would run a simplified model that computes only velocity amplitude and wave direction, yielding the three velocity directions per control point summations $\sum_{q=1}^n \tilde{u}'_{x,j0q}$, $\sum_{q=1}^n \tilde{u}'_{y,j0q}$ and $\sum_{q=1}^n \tilde{u}'_{z,j0q}$. As this is separated from the phase calculation this may be a simplified model. These summations each give the component of the final strongest amplitude, and so can be used to calculate the direction over which to conduct the summations in the second stage.

[00146] In summary:

$$\mathbf{u}_j = \begin{bmatrix} \sum_{q=1}^n \tilde{u}'_{x,j0q} \\ \sum_{q=1}^n \tilde{u}'_{y,j0q} \\ \sum_{q=1}^n \tilde{u}'_{z,j0q} \end{bmatrix}, \hat{n}_j = \frac{\mathbf{u}_j}{\|\mathbf{u}_j\|}.$$

where in the second stage, the model is able to compute:

$$\alpha_{j0q} = \begin{bmatrix} \tilde{u}'_{x,j0q} \\ \tilde{u}'_{y,j0q} \\ \tilde{u}'_{z,j0q} \end{bmatrix} e^{-i\varphi_{j0q}} \cdot \hat{n}_j,$$

inside the transducer model before commencing the other stages of the process. The remaining stages proceed in much the same way as before, only that the directions over which the velocities act are different.

[00147] This approach allows the application of mid-air haptic feedback to more directly address the momentum transfer to the hand, rather than the pressure. This leads to improvements in repeatability of sensations and therefore the measurement of perceptual behavior, which is critical to the expansion of mid-air haptics as a commercialized technology.

[00148] V. Attached Figures

[00149] Figure 17 shows the same wave with the same sound pressure level (SPL) yields different effects at a surface dependent on the angle of incidence. On the left side 1710 a wave (left) travelling parallel to a surface (right), impacts with some energy E , reflects and leaves with slightly less energy as a result of momentum exchange. On the right side 1720 a wave that travels perpendicularly to the same surface does not reflect and exchange energy, therefore leaving with the same quantity of energy. As a result, it is impossible for a haptic effect to be imparted.

[00150] In Figure 18, shown in the top two rows 1810, 1820 are simulations of measurements of an acoustic field generated by a phased array in dB SPL, dB SIL and dB SVL respectively of two different focusses created by the time of flight approach, where light gray is +10 dB, white is 0 dB and dark gray is -10 dB. The bottom row 1830 (dB SPL – dB SIL), (dB SPL – dB SVL) and (dB SIL – dB SVL) respectively with the same color scheme reveals the differences between each measurement method. The SVL field measurement simulation, as it is based on the velocity of the particles of the medium, can allow waves to cancel further due to additional vector component cancellation. As a result, dB SVL is always less than dB SIL which is less than dB SPL, although they are similar distances apart in the logarithmic dB space due to the multiplication of their respective components and effective geometric mean of the SIL measurement.

[00151] 4. EIGENSYSTEM SOLUTIONS FOR QUADRATIC PROBLEMS IN MID-AIR HAPTIC SYSTEMS

[00152] I. Rayleigh-Ritz theorem for quadratic optimization

[00153] Suppose a matrix $\mathbf{M}$ is a square positive semi-definite Hermitian $d \times d$ matrix. Since the matrix is square and Hermitian, each eigenvalue must be real. Since it is positive semi-definite these eigenvalues must be positive as $\mathbf{x}^H \mathbf{M} \mathbf{x} \geq 0$.

[00154] The eigenvalues may then be organized as:

$$\lambda_{min} \leq \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_d \leq \lambda_{max}.$$

[00155] The λ_{min} and λ_{max} of the matrix can then be shown to be solutions to the optimization problems:

$$\lambda_{min} = \min_{x \neq 0} \frac{x^H M x}{x^H x} = \min_{|x|^2=1},$$

and

$$\lambda_{max} = \max_{x \neq 0} \frac{x^H M x}{x^H x} = \max_{|x|^2=1}.$$

[00156] This is the Rayleigh-Ritz theorem. It follows that the x vector from each of these can be found as the corresponding eigenvector to the eigenvalue found.

[00157] II. Quadratic optimization for phased array systems

[00158] When considering phased array systems, it is useful to define as x^H as:

$$x^H = [\alpha_1 e^{-i\varphi_1} \quad \dots \quad \alpha_q e^{-i\varphi_q} \quad \dots \quad \alpha_n e^{-i\varphi_n}],$$

with n the number of monochromatic sound sources, so that this vector solution can be used to drive the array of source transducers. Then define the matrix M as an $n \times n$ Hermitian matrix, such that when x is a unit norm ($|x|^2 = 1$) eigenvector:

$$\begin{aligned} x^H M x &= \\ \alpha_1 e^{-i\varphi_1} M_{1,1} \alpha_1 e^{+i\varphi_1} &+ \\ \alpha_1 e^{-i\varphi_1} M_{1,2} \alpha_2 e^{+i\varphi_2} &+ \\ \dots & \\ \alpha_n e^{-i\varphi_n} M_{n,n-1} \alpha_{n-1} e^{+i\varphi_{n-1}} &+ \\ \alpha_n e^{-i\varphi_n} M_{n,n} \alpha_n e^{+i\varphi_n} &= \\ \lambda_{max}. \end{aligned}$$

[00159] This then represents the form of the optimization. All that is left to do is to fill in the matrix $\mathbf{M}$ with the particular pure quadratic (as the matrix $\mathbf{M}$ is required to be symmetric or Hermitian) objective function that is required.

[00160] There are five potentially useful purely quadratic forms for the objective function that produce Hermitian matrices $\mathbf{M}$ and so can be maximized using existing eigenvector solvers to obtain a solved field that could be used to drive mid-air ultrasonic haptics:

[00161] 1. Maximum pressure optimization (total SPL): a. Maximizes $p\bar{p}$, proportional to the square of the pressure; b. This is the objective function used generally for driving phased array systems.

[00162] 2. Maximum particle velocity optimization (total SVL): a. Maximizes $\mathbf{u} \cdot \bar{\mathbf{u}}$ proportional to the square of the particle velocity; b. This takes into account the directionality of the waves, as this is necessary to obtain consistent haptic perception.

[00163] 3. Maximum particle velocity optimization, but along a given direction (usable SVL): a. Maximizes $(\mathbf{u} \cdot \hat{\mathbf{n}})(\overline{\mathbf{u} \cdot \hat{\mathbf{n}}})$ proportional to the square of particle velocity along a given direction; b. This takes into account both the directionality of the wave and the impact normal with the (skin) surface $\hat{\mathbf{n}}$.

[00164] 4. Maximum acoustic intensity (energy flux density) optimization (total SIL): a.

Maximizes $\sqrt{\frac{1}{4}(\mathbf{p}\bar{\mathbf{u}} \cdot \mathbf{u}\bar{\mathbf{p}})}$ proportional to acoustic intensity; b. This takes into account the directionality of the waves, as this is necessary to obtain consistent haptic perception.

[00165] 5. Maximum acoustic intensity (energy flux density) optimization along a given direction (usable SIL): a. Maximizes $\frac{1}{4}(\mathbf{p}\bar{\mathbf{u}} \cdot \hat{\mathbf{n}} + \bar{\mathbf{p}}\mathbf{u} \cdot \hat{\mathbf{n}})$ proportional to acoustic intensity along a given direction; b. This takes into account both the directionality of the wave and the impact normal with the (skin) surface $\hat{\mathbf{n}}$.

[00166] III. Optimize to maximize pressure

[00167] The quadratic form optimizer capable of maximizing pressure at each control point would be modeling the quantity:

$$\left(\sum_{q=1}^n \tilde{p}_{jq} \alpha_q e^{-i(\omega t + \varphi_{j0q} + \varphi_q)} \right) \left(\sum_{r=1}^n \tilde{p}_{jr} \alpha_r e^{+i(\omega t + \varphi_{j0r} + \varphi_r)} \right),$$

at each control point j , yielding the summation of the objective function

$$\sum_{j=1}^m w_j \left(\sum_{q=1}^n \tilde{p}_{jq} \alpha_q e^{-i(\varphi_{jq} + \varphi_q)} \right) \left(\sum_{r=1}^n \tilde{p}_{jr} \alpha_r e^{+i(\varphi_{jr} + \varphi_r)} \right),$$

which makes each element of the $n \times n$ matrix whose dominant eigenvector is to be solved for:

$$M_{q,r} = \sum_{j=1}^m w_j \tilde{p}_{jq} \tilde{p}_{jr} e^{-i\varphi_{jq}} e^{+i\varphi_{jr}},$$

where $\tilde{p}_{jq}$ is the positive and real amplitude of the wave pressure and $e^{-i\varphi_{jq}}$ the complex phase offset, on the wave travelling from transducer q at control point j from the initial emission. The weight w_j is used to reweight the quadratic $p\bar{p}$ term to provide relative amplitude control among the control points. The global amplitude scaling is provided by the unit norm condition on the input power, $\mathbf{x}^H \mathbf{x} = 1$. This, combined with the quadratic reweighting factor and a control loop can be used to address individual pressure amplitudes precisely.

[00168] IV. Optimize to maximize particle velocity

[00169] The quadratic form optimizer capable of maximizing particle velocity would be modelling the quantity:

$$\begin{aligned} & \left(\sum_{q=1}^n \tilde{u}_{xjq} \alpha_q e^{-i(\omega t + \varphi_{jq} + \varphi_q)} \right) \left(\sum_{r=1}^n \tilde{u}_{xjr} \alpha_r e^{+i(\omega t + \varphi_{jr} + \varphi_r)} \right) + \\ & \left(\sum_{q=1}^n \tilde{u}_{yjq} \alpha_q e^{-i(\omega t + \varphi_{jq} + \varphi_q)} \right) \left(\sum_{r=1}^n \tilde{u}_{yjr} \alpha_r e^{+i(\omega t + \varphi_{jr} + \varphi_r)} \right) + \\ & \left(\sum_{q=1}^n \tilde{u}_{zjq} \alpha_q e^{-i(\omega t + \varphi_{jq} + \varphi_q)} \right) \left(\sum_{r=1}^n \tilde{u}_{zjr} \alpha_r e^{+i(\omega t + \varphi_{jr} + \varphi_r)} \right) \end{aligned}$$

for each control point j . This would yield the summation of the objective function as:

$$\sum_{j=1}^m \left(\left(\sum_{q=1}^n \tilde{u}_{xjq} \alpha_q e^{-i(\varphi_{j0q} + \varphi_q)} \right) \left(\sum_{r=1}^n \tilde{u}_{xjr} \alpha_r e^{+i(\varphi_{j0r} + \varphi_r)} \right) + \right. \\ \left. \left(\sum_{q=1}^n \tilde{u}_{yjq} \alpha_q e^{-i(\varphi_{j0q} + \varphi_q)} \right) \left(\sum_{r=1}^n \tilde{u}_{yjr} \alpha_r e^{+i(\varphi_{j0r} + \varphi_r)} \right) + \right. \\ \left. \left(\sum_{q=1}^n \tilde{u}_{zjq} \alpha_q e^{-i(\varphi_{j0q} + \varphi_q)} \right) \left(\sum_{r=1}^n \tilde{u}_{zjr} \alpha_r e^{+i(\varphi_{j0r} + \varphi_r)} \right) \right),$$

which makes each element of the $n \times n$ matrix whose dominant eigenvector is to be solved for:

$$M_{q,r} = \sum_{j=1}^m w_j (\tilde{u}_{xjq} \tilde{u}_{xjr} + \tilde{u}_{yjq} \tilde{u}_{yjr} + \tilde{u}_{zjq} \tilde{u}_{zjr}) e^{-i\varphi_{j0q}} e^{+i\varphi_{j0r}},$$

where u_{xjq} is the x direction part of the induced particle velocity amplitude (which is real and positive) and $e^{-i\varphi_{j0q}}$ the complex phase offset, on the wave travelling from transducer q at control point j from the initial emission. The weight w_j is used to reweight the quadratic $\mathbf{u} \cdot \bar{\mathbf{u}}$ term to provide relative particle velocity amplitude control among the control points. The global amplitude scaling is provided by the unit norm condition on the input power, $\mathbf{x}^H \mathbf{x} = 1$. This, combined with the quadratic reweighting factor and a control loop can be used to address individual velocity amplitudes precisely.

[00170] V. Optimize to maximize particle velocity in a set of given directions

[00171] The quadratic form optimizer capable of maximizing particle velocity in a set of given directions would be modeling the quantity:

$$\left(\sum_{q=1}^n \tilde{u}_{xjq} \hat{n}_{xj} \alpha_q e^{-i(\omega t + \varphi_{j0q} + \varphi_q)} \right) \left(\sum_{r=1}^n \tilde{u}_{xjr} \hat{n}_{xj} \alpha_r e^{+i(\omega t + \varphi_{j0r} + \varphi_r)} \right) + \\ \left(\sum_{q=1}^n \tilde{u}_{yjq} \hat{n}_{yj} \alpha_q e^{-i(\omega t + \varphi_{j0q} + \varphi_q)} \right) \left(\sum_{r=1}^n \tilde{u}_{yjr} \hat{n}_{yj} \alpha_r e^{+i(\omega t + \varphi_{j0r} + \varphi_r)} \right) + \\ \left(\sum_{q=1}^n \tilde{u}_{zjq} \hat{n}_{zj} \alpha_q e^{-i(\omega t + \varphi_{j0q} + \varphi_q)} \right) \left(\sum_{r=1}^n \tilde{u}_{zjr} \hat{n}_{zj} \alpha_r e^{+i(\omega t + \varphi_{j0r} + \varphi_r)} \right)$$

for each control point j . This would yield the summation of the objective function as:

$$\begin{aligned} & \sum_{j=1}^m \left(\left(\sum_{q=1}^n \tilde{u}_{xjq} \hat{n}_{xj} \alpha_q e^{-i(\varphi_{j0q} + \varphi_q)} \right) \left(\sum_{r=1}^n \tilde{u}_{xjr} \hat{n}_{xj} \alpha_r e^{+i(\varphi_{j0r} + \varphi_r)} \right) + \right. \\ & \left(\sum_{q=1}^n \tilde{u}_{yjq} \hat{n}_{yj} \alpha_q e^{-i(\varphi_{j0q} + \varphi_q)} \right) \left(\sum_{r=1}^n \tilde{u}_{yjr} \hat{n}_{yj} \alpha_r e^{+i(\varphi_{j0r} + \varphi_r)} \right) + \\ & \left. \left(\sum_{q=1}^n \tilde{u}_{zjq} \hat{n}_{zj} \alpha_q e^{-i(\varphi_{j0q} + \varphi_q)} \right) \left(\sum_{r=1}^n \tilde{u}_{zjr} \hat{n}_{zj} \alpha_r e^{+i(\varphi_{j0r} + \varphi_r)} \right) \right), \end{aligned}$$

which makes each element of the $n \times n$ matrix whose dominant eigenvector is to be solved for:

$$M_{q,r} = \sum_{j=1}^m w_j (\tilde{u}_{xjq} \tilde{u}_{xjr} \hat{n}_{xj}^2 + \tilde{u}_{yjq} \tilde{u}_{yjr} \hat{n}_{yj}^2 + \tilde{u}_{zjq} \tilde{u}_{zjr} \hat{n}_{zj}^2) e^{-i\varphi_{j0q}} e^{+i\varphi_{j0r}},$$

where $\tilde{u}_{xjq}$ is the x direction part of the induced particle velocity amplitude (which is real and positive) and $e^{-i\varphi_{j0q}}$ the complex phase offset, on the wave travelling from transducer q at control point j from the initial emission. The weight w_j is used to reweight the quadratic $(\mathbf{u} \cdot \hat{\mathbf{n}})(\overline{\mathbf{u} \cdot \hat{\mathbf{n}}})$ term to provide relative amplitude control among the control points. The global amplitude scaling is provided by the unit norm condition on the input power, $\mathbf{x}^H \mathbf{x} = 1$. This, combined with the quadratic reweighting factor and a control loop can be used to address individual amplitudes precisely.

[00172] VI. Active Acoustic Intensity

[00173] Active acoustic intensity $\mathbf{I}_a$ can be defined using the pressure p and particle velocity vector $\mathbf{u}$ of the acoustic medium. For an acoustic wave, this can be described as:

$$\mathbf{I}_a = \frac{1}{2} \Re(\mathbf{p} \cdot \overline{\mathbf{u}}),$$

where the active intensity is a purely real vector quantity. In the near field of a wave source, the reactive field dominates instead $\mathbf{I}_r = \frac{1}{2} \Im(\mathbf{p} \cdot \overline{\mathbf{u}})$. However, in the far field, so beyond one or two wavelengths, $\mathbf{I}_r \rightarrow \mathbf{0}$. For this reason, only $\mathbf{I}_a$ is considered and assumed to be guaranteed real. Considering a monochromatic plane wave source, where:

$$\begin{aligned} \mathbf{p} &= \tilde{p} e^{-i(\omega t + \varphi)}, \\ \mathbf{u} &= \tilde{u} e^{-i(\omega t + \varphi)}, \end{aligned}$$

we have:

$$I_a = \frac{1}{2} (\tilde{p} e^{-i(\omega t + \varphi)} \cdot \overline{\tilde{u} e^{-i(\omega t + \varphi)}}),$$

where $\tilde{p}$ is the pressure amplitude, and $\tilde{u}$ is the particle velocity amplitude in each direction.

[00174] The sound intensity and thus energy flux transmitted by the wave along a given axis is also of interest so the dot product to measure the scalar contribution is formed:

$$I_a \cdot \hat{n} = \frac{1}{2} (\tilde{p} e^{-i(\omega t + \varphi)} \cdot \overline{\tilde{u} e^{-i(\omega t + \varphi)}} \cdot \hat{n})$$

[00175] Due to the form of the acoustic wave equations:

$$\frac{\partial^2 p}{\partial t^2} - c_0^2 \nabla^2 p = 0,$$

$$\frac{\partial^2 u}{\partial t^2} - c_0^2 \nabla^2 u = 0,$$

the superposition principle must hold for both pressure and particle velocity vectors. Assuming n acoustic sources and approximating each as a plane wave at the point of interest, then yields the equations:

$$I_a = \frac{1}{2} \left(\sum_{q=1}^n \tilde{p}_q e^{-i(\omega t + \varphi_q)} \cdot \overline{\sum_{r=1}^n \tilde{u}_r e^{-i(\omega t + \varphi_r)}} \right).$$

$$I_a \cdot \hat{n} = \frac{1}{2} \left(\sum_{q=1}^n \tilde{p}_q e^{-i(\omega t + \varphi_q)} \cdot \overline{\sum_{r=1}^n \tilde{u}_r e^{-i(\omega t + \varphi_r)}} \cdot \hat{n} \right).$$

[00176] It is known that in a regime in question $\Im(p \cdot \bar{u}) \rightarrow 0$, and p and $\bar{u}$ are in phase, then it is known that the placement of the conjugation does not matter and it is possible to reverse order of the conjugations to conjugate pressure instead of particle velocity (this can be verified by expanding the above equations and using commutativity to reverse the order of the complex exponentials in each term). From inference, this can be written in a way that generates a Hermitian matrix $\mathbf{M}$:

$$I_a = \frac{1}{4} \left(\sum_{q=1}^n \tilde{p}_q e^{-i(\omega t + \varphi_q)} \cdot \overline{\sum_{r=1}^n \tilde{u}_r e^{-i(\omega t + \varphi_r)}} + \right.$$

$$\sum_{q=1}^n \tilde{u}_q e^{-i(\omega t + \varphi_q)} \cdot \overline{\sum_{r=1}^n \tilde{p}_r e^{-i(\omega t + \varphi_r)}} \Bigg),$$

$$I_a \cdot \hat{n} = \frac{1}{4} \left(\sum_{q=1}^n \tilde{p}_q e^{-i(\omega t + \varphi_q)} \cdot \overline{\sum_{r=1}^n \tilde{u}_r e^{-i(\omega t + \varphi_r)}} \cdot \hat{n} + \right.$$

$$\left. \sum_{q=1}^n \tilde{u}_q e^{-i(\omega t + \varphi_q)} \cdot \hat{n} \cdot \overline{\sum_{r=1}^n \tilde{p}_r e^{-i(\omega t + \varphi_r)}} \right).$$

[00177] VII. Optimize to maximize acoustic intensity

[00178] The first quantity that generates a scalar value is:

$$\frac{1}{4} \sum_{q=1}^n \tilde{p}_{jq} \alpha_q e^{-i(\omega t + \varphi_{j0q} + \varphi_q)} \cdot \sum_{r=1}^n [\tilde{u}_{xjr}, \tilde{u}_{yjr}, \tilde{u}_{zjr}] \alpha_r e^{+i(\omega t + \varphi_{j0r} + \varphi_r)} \cdot$$

$$\sum_{s=1}^n [\tilde{u}_{xjs}, \tilde{u}_{yjs}, \tilde{u}_{zjs}] \alpha_s e^{-i(\omega t + \varphi_{j0s} + \varphi_s)} \cdot \sum_{v=1}^n \tilde{p}_{jv} \alpha_v e^{+i(\omega t + \varphi_{j0v} + \varphi_v)}$$

but then assuming it is in the far field, this may also be written as:

$$\frac{1}{4} \sum_{q=1}^n \tilde{p}_{jq} \alpha_q e^{-i(\omega t + \varphi_{j0q} + \varphi_q)} \cdot \sum_{r=1}^n \tilde{p}_{jr} \cdot [\hat{d}_{xjr}, \hat{d}_{yjr}, \hat{d}_{zjr}] \alpha_r e^{+i(\omega t + \varphi_{j0r} + \varphi_r)} \cdot$$

$$\sum_{s=1}^n \tilde{p}_{js} \cdot [\hat{d}_{xjs}, \hat{d}_{yjs}, \hat{d}_{zjs}] \alpha_s e^{-i(\omega t + \varphi_{j0s} + \varphi_s)} \cdot \sum_{v=1}^n \tilde{p}_{jv} \alpha_v e^{+i(\omega t + \varphi_{j0v} + \varphi_v)}$$

this is a fourth power of the pressure summation, but due to the structure of the complex exponentials two of these are the same and so can be removed through a square root.

[00179] The quadratic form optimizer capable of maximizing acoustic intensity would therefore be modelling the quantity:

$$\frac{1}{2} \sum_{q=1}^n \sum_{r=1}^n \tilde{p}_{jq} \alpha_q e^{-i(\omega t + \varphi_{j0q} + \varphi_q)} \tilde{p}_{jr} \alpha_r e^{+i(\omega t + \varphi_{j0r} + \varphi_r)} \sqrt{[\hat{d}_{xjq}, \hat{d}_{yjq}, \hat{d}_{zjq}] \cdot [\hat{d}_{xjr}, \hat{d}_{yjr}, \hat{d}_{zjr}]}$$

for each control point j . This would yield the summation of the objective function as:

$$\frac{1}{2} \sum_{j=1}^m \left(\left(\sum_{q=1}^n \tilde{p}_{jq} \alpha_q e^{-i(\varphi_{jq} + \varphi_q)} \right) \left(\sum_{r=1}^n \tilde{p}_{jr} \alpha_r e^{+i(\varphi_{jr} + \varphi_r)} \right) \right. \\ \left. \left(\sum_{q=1}^n \sum_{r=1}^n \sqrt{[\hat{d}_{xjq}, \hat{d}_{yjq}, \hat{d}_{zjq}] \cdot [\hat{d}_{xjr}, \hat{d}_{yjr}, \hat{d}_{zjr}]} \right) \right),$$

which makes each element of the $n \times n$ matrix whose dominant eigenvector is to be solved for:

$$M_{q,r} = \frac{1}{2} \sum_{j=1}^m w_j \left(\tilde{p}_{jq} \tilde{p}_{jr} \sqrt{[\hat{d}_{xjq}, \hat{d}_{yjq}, \hat{d}_{zjq}] \cdot [\hat{d}_{xjr}, \hat{d}_{yjr}, \hat{d}_{zjr}]} \right) e^{-i\varphi_{jq}} e^{+i\varphi_{jr}}$$

where $\hat{d}_{xjq}$ is the x direction part of the wavefront normal vector, $\tilde{p}_{jq}$ is the positive and real amplitude of the wave pressure and $e^{-i\varphi_{jq}}$ the complex phase offset, on the wave travelling from transducer q at control point j from the initial emission. The weight w_j is used to reweight the quadratic term which really remains the square root of the function involving a fourth power of pressure, $\sqrt{\frac{1}{4}(\mathbf{p}\bar{\mathbf{u}} \cdot \mathbf{u}\bar{\mathbf{p}})}$ to provide relative particle velocity amplitude control among the control points. The global amplitude scaling is provided by the unit norm condition on the input power, $\mathbf{x}^H \mathbf{x} = 1$. This, combined with the quadratic reweighting factor and a control loop can be used to address individual velocity amplitudes precisely.

[00180] VIII. Optimize to maximize acoustic intensity in a set of given directions

[00181] The quadratic form optimizer capable of maximizing acoustic intensity along a given direction would be modelling the quantity:

$$\frac{1}{4} \sum_{q=1}^n \tilde{p}_{jq} \alpha_q e^{-i(\omega t + \varphi_{jq} + \varphi_q)} \cdot \sum_{r=1}^n (\tilde{u}_{xjr} \hat{n}_{xj} + \tilde{u}_{yjr} \hat{n}_{yj} + \tilde{u}_{zjr} \hat{n}_{zj}) \alpha_r e^{+i(\omega t + \varphi_{jr} + \varphi_r)} + \\ \frac{1}{4} \sum_{q=1}^n (\tilde{u}_{xjq} \hat{n}_{xj} + \tilde{u}_{yjq} \hat{n}_{yj} + \tilde{u}_{zjq} \hat{n}_{zj}) \alpha_q e^{-i(\omega t + \varphi_{jq} + \varphi_q)} \cdot \sum_{r=1}^n \tilde{p}_{jr} \alpha_r e^{+i(\omega t + \varphi_{jr} + \varphi_r)}$$

but then assuming it is in the far field, this may also be written as:

$$\frac{1}{4} \sum_{q=1}^n \tilde{p}_{jq} \alpha_q e^{-i(\omega t + \varphi_{jq} + \varphi_q)} \cdot \sum_{r=1}^n \tilde{p}_{jr} [\hat{d}_{xjr}, \hat{d}_{yjr}, \hat{d}_{zjr}] \cdot [\hat{n}_{xj}, \hat{n}_{yj}, \hat{n}_{zj}] \alpha_r e^{+i(\omega t + \varphi_{jr} + \varphi_r)} +$$

$$\frac{1}{4} \sum_{q=1}^n \tilde{p}_{jq} [\hat{d}_{xjq}, \hat{d}_{yjq}, \hat{d}_{zjq}] \cdot [\hat{n}_{xj}, \hat{n}_{yj}, \hat{n}_{zj}] \alpha_q e^{-i(\omega t + \varphi_{jq} + \varphi_q)} \cdot \sum_{r=1}^n \tilde{p}_{jr} \alpha_r e^{+i(\omega t + \varphi_{jr} + \varphi_r)}$$

for each control point j . This would yield the summation of the objective function as:

$$\begin{aligned} & \frac{1}{4} \sum_{j=1}^m \left(\left(\sum_{q=1}^n \tilde{p}_{jq} \alpha_q e^{-i(\varphi_{jq} + \varphi_q)} \right) \right. \\ & \left(\sum_{r=1}^n \tilde{p}_{jr} [\hat{d}_{xjr}, \hat{d}_{yjr}, \hat{d}_{zjr}] \cdot [\hat{n}_{xj}, \hat{n}_{yj}, \hat{n}_{zj}] \alpha_r e^{+i(\varphi_{jr} + \varphi_r)} \right) + \\ & \left(\sum_{q=1}^n \tilde{p}_{jq} [\hat{d}_{xjq}, \hat{d}_{yjq}, \hat{d}_{zjq}] \cdot [\hat{n}_{xj}, \hat{n}_{yj}, \hat{n}_{zj}] \alpha_q e^{-i(\varphi_{jq} + \varphi_q)} \right) \\ & \left. \left(\sum_{r=1}^n \tilde{p}_{jr} \alpha_r e^{+i(\varphi_{jr} + \varphi_r)} \right) \right), \end{aligned}$$

which makes each element of the $n \times n$ matrix whose dominant eigenvector is to be solved for:

$$\begin{aligned} M_{q,r} = & \frac{1}{4} \sum_{j=1}^m w_j (\tilde{p}_{jq} \tilde{p}_{jr} (\hat{d}_{xjq} \hat{n}_{xj} + \hat{d}_{yjq} \hat{n}_{yj} + \hat{d}_{zjq} \hat{n}_{zj} + \\ & \hat{d}_{xjr} \hat{n}_{xj} + \hat{d}_{yjr} \hat{n}_{yj} + \hat{d}_{zjr} \hat{n}_{zj})) e^{-i\varphi_{jq}} e^{+i\varphi_{jr}}, \end{aligned}$$

where $\tilde{u}_{xjq}$ is the x direction part of the induced particle velocity amplitude (which is real and positive), $\tilde{p}_{jq}$ is the positive and real amplitude of the wave pressure and $e^{-i\varphi_{jq}}$ the complex phase offset, on the wave travelling from transducer q at control point j from the initial emission. The weight w_j is used to reweight the quadratic $\frac{1}{4}(\overline{p\mathbf{u}} \cdot \hat{\mathbf{n}} + \bar{p}\mathbf{u} \cdot \hat{\mathbf{n}})$ term to provide relative particle velocity amplitude control among the control points. The global amplitude scaling is provided by the unit norm condition on the input power, $\mathbf{x}^H \mathbf{x} = 1$. This, combined with the quadratic reweighting factor and a control loop can be used to address individual velocity amplitudes precisely.

[00182] IX. Finding the weighting factors

[00183] The Rayleigh-Ritz theorem shows that any of the above matrices can be maximized by finding the dominant eigenvector and allowing the power to increase up to the physical limitations of the transducers. This can be achieved through the power iteration. By

constructing the matrix $\mathbf{M}$ while individually defining each of the sum of control points separately, when multiplying through by the vector $\mathbf{x}$ the evaluation of the objective function can be performed as an additional side effect. This can be used to find appropriate values for the weights w_j such that they allow the maximization to find amplitudes that when scaled are appropriate to the solution of the system.

[00184] As the power iteration describes:

$$\mathbf{x}_a = \frac{\mathbf{M}\mathbf{x}_{a-1}}{\|\mathbf{M}\mathbf{x}_{a-1}\|},$$

where:

$$\mathbf{x} = \lim_{a \rightarrow \infty} \mathbf{x}_a.$$

[00185] We can use the power iteration to determine the objective function at each control point.

[00186] Defining:

$$M_{q,r} = \sum_{j=1}^m M_{q,r,j},$$

and:

$$\mathbf{M} = \sum_{j=1}^m \mathbf{M}_j.$$

[00187] On each iteration it can instead compute:

$$\mathbf{x}_{a,j} = \mathbf{M}_j \mathbf{x}_{a-1},$$

so, the objective function at each control point can be evaluated as:

$$o_j(\mathbf{x}_{a-1}) = \mathbf{x}_{a-1}^H \mathbf{x}_{a,j}.$$

[00188] Finally, as before $\mathbf{x}_a$ can be computed:

$$\mathbf{x}_a = \frac{\sum_{j=1}^m \mathbf{x}_{a,j}}{\|\sum_{j=1}^m \mathbf{x}_{a,j}\|},$$

where the final eigenvector solution is as above. However, calculating it in this way allows $o_j(\mathbf{x}_{a-1})$ to yield an intermediate result for the control point j . To fix the level at which this is representative of the desired control point output, the average objective function value is taken in order to allow the set of values $o_j(\mathbf{x}_{a-1})$ to be directly compared to the objective values given by the user. In this case, the weighting factors can be updated for the next iteration using any derivative-free method meant for the optimization of expensive functions, such as the Nelder-Mead algorithm for example.

[00189] After some number of iterations of the power method and coupled derivative-free optimization, the intermediate result $\mathbf{x}_a$ should be used to drive the transducers. This will therefore allow the system to track the desired values of the objective functions through time, while allowing for some error. The advantage of this system is that it is a very simple way to drive a transducer array. The largest disadvantage is that many calculations are required, but the calculations are effectively only complex multiplications which with sufficient hardware are straightforwardly implemented.

[00190] X. Attached Figure

[00191] Shown in Figure 19 are: Top left 1910: Control point created by solving the quadratic particle velocity problem with a non-perpendicular vector, shown as a cut through the x - y plane. Top right 1920: The same control point displayed as a cut through the x - z plane. Bottom left 1930: For comparison, a control point created by solving the quadratic pressure problem, shown as a cut through the x - y plane. Bottom right 1940: A cut through the solution to the pressure problem across the x - z plane.

[00192] 5. CONTROL REGIONS FOR PHASED ARRAYS SPECIFIED THROUGH QUADRATIC PROBLEMS

[00193] This disclosure describes a method to automatically generate basis functions using user supplied design parameters. These are input to an efficient optimization over a quadratic function in the space of a given transducer basis set in an online manner. This gives the user a design space, accessible through an online optimization that may generate a variety of different effects. Using a linear operator, termed a ‘quadratic control region’ as it is a control region for a quadratic problem, a matrix may be constructed by evaluating the dyadic product over the transducer effects. The dominant eigenvector of this matrix may be chosen as the best solution from a set of orthonormal solutions which are identically the eigenvectors of this matrix. The optimization produces as output the transducer activation coefficients – the raw complex values needed to drive the individual transducer elements to recreate the acoustic field or phenomena with the desired properties.

[00194] These raw complex values may then be used to further comprise the basis functions of a later linear system whose solution enables the user to produce all the desired effects simultaneously whilst reweighting the output power available to each.

[00195] I. Linear Control Region

[00196] A linear control region, which is defined as an operator that performs an integral with an linear acoustic quantity, such as pressure, with a complex valued weighting function, $w(x, y, z)$:

$$f_p(x, y, z) = \iiint_{\Omega} \overline{w(x, y, z)} \cdot p(x, y, z) dx dy dz$$

may be used to evaluate an acoustic field. It should be noted that the complex conjugate is used to maximize function evaluation when $w(x, y, z)$ is paired with similar phase distributions in the acoustic quantity. Any approximation to this integral (such as sampling) can be used or indeed this may be used to approximate any linear operator over the field of acoustic pressure. A further integral may be constructed to optimize a control region using a vector particle velocity, given it is also linear, wherein:

$$f_v(x, y, z) = \iiint_{\Omega} \overline{\mathbf{w}(x, y, z)} \cdot \mathbf{u}(x, y, z) dx dy dz,$$

with $\mathbf{u}$ the particle velocity of the medium, and $\mathbf{w}$ a three-dimensional vector field in which each dimension is complex-valued, which serves as a weighting function with the same purpose as above. In this way, different patterns of acoustic quantities may be evaluated in a linear fashion, maximizing a particular phase of the result when phases align with the ‘template’ created by the scalar or vector function in w .

[00197] It also should be noted that $w(x, y, z)$ may be largely zero (or comprised of zero vectors) to create a single evaluation of multiple seemingly disjoint regions. Any integral approximation (such as sampling) can be used as this formulation can extend to any linear operator over the field of a linear acoustic quantity.

[00198] A phased array may be described by a series of complex valued coefficients that controls each transducer. Driving an array of acoustic transducers at a fixed frequency with such an input $\mathbf{x}$ vector to maximise an acoustic quantity encoded in a matrix $\mathbf{M}$ may be expressed as a linear algebra problem, so that the intention is to:

$$\begin{aligned} &\text{maximize} && \mathbf{x}^H \mathbf{M} \mathbf{x} \\ &\text{subject to} && \mathbf{x}^H \mathbf{x} = \mathbf{1}, \\ &\text{and} && \mathbf{x} \in \mathbb{C}^n. \end{aligned}$$

[00199] This may be solved by taking the dominant eigenvector of the matrix $\mathbf{M}$, as the statement of the problem is also the definition of an eigenvector, here $\mathbf{x}$. Building the matrix $\mathbf{M}$ may be achieved by first considering the linear acoustic quantities emitted by each transducer as a basis function in the global space of the acoustic quantity in question. This is a way to achieve the maximization of pressure and other acoustic quantities at a single point, described by an absolute position in x , y and z . However, using the described control region integral, the $q(x, y, z)$ that governs the acoustic quantities emitted into the field of each individual transducing element may be exchanged for $f(x, y, z)$, yielding the column vector:

$$\mathbf{f} = \begin{bmatrix} f_1(x, y, z) \\ \vdots \\ f_n(x, y, z) \end{bmatrix},$$

[00200] By invoking linearity, the complex activation values driving each of the transducers may be multiplied leading to the vector:

$$\mathbf{m} = \begin{bmatrix} x_1 \cdot f_1(x, y, z) \\ \vdots \\ x_n \cdot f_n(x, y, z) \end{bmatrix},$$

where n is the number of transducers. Given that this $\mathbf{m}$ is the result of a linear operator operating on a linear field, it may be Hermitian transposed and multiplied by itself to generate a quadratic form. Finally,

$$\mathbf{m}^H \mathbf{m} = \mathbf{x}^H \mathbf{M} \mathbf{x} = \mathbf{x}^H \mathbf{f} \mathbf{f}^H \mathbf{x}.$$

[00201] By modifying the weighting function $w(x, y, z)$ and solving for the eigenvectors of the matrix $\mathbf{M}$, a variety of different beams, surfaces and other acoustic phenomenon may be created and maximized, producing optimally actuated regions in the acoustic field.

[00202] The null space generation technique discussed previously (“Eigensystem solutions for quadratic problems in mid-air haptic systems”) may also be used to extend this method to describe acoustic levitation, as well as allowing for null steering and the formation of mid-air texture by introducing null points into the field at positions to describe structure. It should be

noted that this process increases the rank of the matrix, so simplification is not possible should this be used.

[00203] Equally, multiple optimizations may be carried out by adding together ***M*** matrices for a number of control regions. This may also include matrices that represent control points. This generates a matrix that has increased rank, so some simplification methods are not possible in the case that this is used.

[00204] By considering the result and reweighting each of the points in an iterative manner, the system may be modified so as to converge to the required geometric configuration of an acoustic quantity. This may be generally needed when the amplitudes are not consistent in the output.

[00205] This may be achieved by summing the acoustic quantities produced and reweighting the system to in the next iteration progress more closely to the result. For instance, the iterative reweighting for any sub-split volume of integral or single point of pressure becomes:

$$w_{t+1,i}(x, y, z) := w_{t,i}(x, y, z) \cdot \frac{p_{r,i}(x, y, z)}{p_{t,i}(x, y, z)},$$

where $t + 1$ refers to the next coarse iteration of the technique, t the current iteration, p_r the desired pressure (which may be complex) and the denominator the total pressure obtained by the current iteration. The fraction may be rendered as only real magnitudes in order to only affect the amplitudes of the resulting region and not the phases.

[00206] A Bessel beam may be created by expressing a weighting function, $w(x, y, z)$, in the direction of beam travel and creating unit pressure while modifying the phase angle forward as the beam moves forward. This may be sampled with point samples, so long as the angle traversed by each subsequent point is small enough to sufficiently sample the phase along the beam axis.

[00207] The sampling of the phase must importantly have more than two phase samples for every 2π in angle that it moves. This is because if the increment along the path is π , the direction of the wave travel is undecidable. This also means in the majority of cases, the phase must increment forward such that it completes a 2π radian rotation for every wavelength of the monochromatic frequency distance it covers. This is an approximation to a path integral of the form described in the control region definition.

[00208] The transducers, as a result of maximizing the integral written in this way, will in the case of a simple beam progressively defined in a straight line moving away from the array, generate an optimal Bessel-like beam due to its theoretical ability to generate optimal pressure in this configuration. Shown in Figure 20 are cuts through the field at the edge of the

near field region parallel to the flat rectilinear transducer array, revealing the structure of the complex-valued transducer activation coefficients. A simple progressive phase beam results in a Bessel function approximation (left 2000) and an under-sampled helical beam results in a Bessel-like function with dipole behavior (right 2010), where the dipole creates two structures that activate in a manner wherein each opposes the other in phase. The beam may not be an exact Bessel beam due to the effects of limited transducer resolution, finite array size and other physical considerations which the optimization will need to work around, but may perform more effectively than a simple evaluation of a Bessel beam function due to the optimization taking into account the physical limitations of a real transducer array. On a large array, it follows the Bessel function closely, Figure 21 shows a graph 2100 of a Bessel function plotted alongside pressure when considered radially from the center of the array in Figure 20.

[00209] A further property of this system is that transducers in any position and orientation, with any generated monochromatic field function may be used. This is a significant advantage in Bessel beam generation over a sampling method, which would by definition function only in the planar array case.

[00210] By creating multiple out-of-phase paths, phase shifts and phase singularities may be created implicitly by winding helical waves around a path. This path may also be curved, yielding a large degree of configurability, while using the simple quadratic optimization embedded in the eigensystem approach. This is achieved by inducing rotations in the beam through the definition of separate phase windings in many parallel beam lines that are closer than the beam width but differ in phase between each other. Then, by setting the nearby phases to cancel each other in the central line of the beam, a high-order Bessel beam may be formed, as shown in Figure 22. In Figure 22 more than two samples of the progressive phases around a circle (thus, phases in a helical arrangement) causes the formation of Bessel function with angular momentum as seen on the left 2210. By using more than 2π radians for specifying the phases around the helix, higher order Bessel beams may be generated such as the 4π beam on the right 2220.

[00211] These lines of phase may be sampled using points or any other integral approximation as before, as these are examples of how disjoint regions may be defined as a single regional optimization, wherein each may be added to the optimization region by setting the weighting function to non-zero only in the volume to be optimized. High-order Mathieu-type elliptic helical beams may be created by introducing incremental phase beams incremental in phase and spaced around an ellipse. As shown in Figure 23, elliptical helical

beams that have similar properties to high-order Bessel beams, such as Matheiu-type beams, may be created by parametrizing the circumference of an ellipse in phase. Determining a proportion of the $2n\pi$ radians and allocating a phase angle proportional to the ellipse chord circumference generates the appropriate configuration for the optimisation of the beam. The left image 2310 shows a high order Bessel reshaped into an ellipse, nearing a Matheiu beam. The right image 2320 shows an approximate Matheiu beam with cross-linked helical structures beginning to form in the central region. Both are based on a phase circle of 12π radians, creating a helical wavefront.

[00212] The Bessel family of beams that are created in this manner may be modified to produce more exotic beams that exist in more restrictive forms in the literature. By creating arbitrary paths in the field traversed by physically plausible phase functions that move forward in a way appropriate to the traversal of waves, a variety of results may be reproduced, but due to the optimization implied by the eigensystem in a form that is more general and lower in initial bandwidth cost than before, as the beam need only be expressed by the integral.

[00213] Taking a curved path and moving the phase forward may produce a curved beam, generally used to maneuver beams around obstacles produced by a 45° circle segment followed by a path that pushes the beam formed out to infinity. Shown in Figure 24 is an image 2400 of a bending beam, formed from a phase function that follows a 45 degree segment of a circle connected to a continuation along a line, using the coarse iteration method to equalize acoustic pressures along the beam. The transducer array is situated along the lower edge of the image.

[00214] A further approach involving multiple phase paths may create a “bottle beam”, a beam with a portion cut from the center, which can facilitate the levitation of large objects. This has been generated in Figure 25 which shows an image 2500 of twinning two Gaussian curves along which the phase has been made to follow. The bottle beam produced by creating phase functions follow Gaussian curves using the coarse iteration method to equalize acoustic pressures along the beam. The transducer array is situated along the lower edge of the image. The dotted nature of the lower region suggests that the phase velocity input to the optimization is too slow, and that the beam may be improved by using a faster phase velocity for the initial section. Portions of curves may also be used, such as in Figure 26, since if the object is sufficiently rigid, only some sections of the space need to be made high acoustic pressure regions. Figure 26 shows an image 2600 of an incomplete bottle beam formed from

portions of a circle defined in progressive phase that may be used to levitate even larger objects by optimizing the region of high acoustic pressure in a non-uniform manner. It is not necessary to spread the forces exerted on the large object evenly if the object or matter to be levitated is assumed to have enough cohesion to not flow, extrude or leak through the low acoustic pressure regions.

[00215] A Bessel-like beam may be created by moving a path forward faster than the phase velocity. In this way, a beam will be optimized for with a faster phase velocity than usual, creating a beam that appears to have longer wavelength and by virtue of the phase velocity travel faster than the speed of sound, as shown in Figure 27. Shown in Figure 27 is an image 2710 of a wave front that travels faster than the speed of sound, which may be created by implementing a phase path that moves through angle in a way that paradoxically generates a longer wavelength with the same frequency. The implicit optimization infers that the main beam is composed of intersecting beams, giving rise to the specified increase in phase velocity. This can be seen from the image 2720 on the right, where the wavelength is manipulated to be approximately 40% longer at the same frequency for the initial section of the beam. A snaking beam may be constructed by sampling progressive phases along a sine curve. Shown in Figure 28 are examples 2810 wherein a short section of the sine curve is defined, followed by 2820 of a longer two-dimensional snaking beam constructed from progressive phases of pressure following a sine curve above a phased array. These show that the extension of the snaking sine wave beam by simply adding further sections of progressive phase causes the beam to be naturally extended as far as required (and allowed by the physical arrangement of transducers). By creating the snaking beam in all three dimensions to make a helically snaking beam, a beam can be produced that may be used in a manner similar to an Archimedes screw system to apply forces, potentially in the direction perpendicular or opposite to the emission of the waves, to solid objects in the field that are much larger than a wavelength. This gives a different mechanism whereby a beam may be used to levitate objects. Figure 29 shows a left image 2910 detailing a section in the x - z plane and a right image 2920 showing a section of the beam in the x - y plane. Using this technique of following a path in a circle winding in the x and y directions, but which is also progressive in z , a true helical formation can be created. This could be used to levitate and apply forces in a tractor beam-like manner to objects much larger than a wavelength.

[00216] If the eigensystem matrix is defined by a single control region, then the matrix is of rank 1 and formed by the dyadic product of the complex-valued vector $\mathbf{m}$ with itself only,

defining the matrix as $\mathbf{f}^H \mathbf{f}$. In this specific case, the single eigenvector of the system is known immediately and is $\mathbf{f}$, meaning $\mathbf{f} = \mathbf{x}$. By evaluating the objective function in this case, the objective function specifies itself the coefficients of the complex-valued transducer activations needed. This is defined by simply evaluating the control region function for each transducer, which becomes the activation coefficient.

[00217] As before, due to the quadratic form, evaluating the effectiveness through a re-weighting on each iteration may also be pursued. This is achieved by evaluating the pressure at each point or the integral of pressure over each sub-region and dividing through by the desired pressure, generating a new function value to update $w_{t,i}(x, y, z)$, the weight for the sub-region i . The single eigenvector which is also $\mathbf{f}$ may then be re-evaluated, generating an equivalent approach to the repeated iteration over the eigensystem from before, but at a greatly reduced computational cost.

[00218] To achieve this with the particle velocity of the medium, we can use a test of the particle velocity, which amounts to changing the direction until a maximum is reached:

$$w_{t+1,x,i}(x, y, z) := w_{t,x,i}(x, y, z) \cdot \frac{u_{r,x,i}(x, y, z)}{u_{t,x,i}(x, y, z)},$$

although the pressure may also be monitored using the previous technique to create a hybrid method. This hybrid technique involves specifying the weighting function and update in pressure as before, but with a control region that is effectively:

$$f_v(x, y, z) = \iiint_{\Omega} \overline{w(x, y, z)} \cdot \mathbf{v}(x, y, z) \cdot \mathbf{u}(x, y, z) dx dy dz,$$

yielding a scalar but changeable and complex-valued $w(x, y, z)$ paired with a unchanging weighting normal vector $\mathbf{v}$ and particle velocity $\mathbf{u}$. This approach can optimise to find a satisfactory pressure but without changing the direction of the wave that is sought. This is shown, alongside specific corrections described in the next section, in Figure 28.

[00219] II. Wavefront filtering for directional control points and regions

[00220] A further technique necessary to the creation of directional output is to discard transducer data, set it to zero, in the case where the transducer output wavefront $\hat{\delta}$ and the normal vector or particle velocity vector $\mathbf{u}$ are opposite, so satisfy:

$$\begin{bmatrix} \hat{\delta}_{x,i}(x, y, z) \\ \hat{\delta}_{y,i}(x, y, z) \\ \hat{\delta}_{z,i}(x, y, z) \end{bmatrix} \cdot \begin{bmatrix} u_{x,i}(x, y, z) \\ u_{y,i}(x, y, z) \\ u_{z,i}(x, y, z) \end{bmatrix} < 0,$$

for the point of the region under consideration.

[00221] Figure 30 shows a left image 3010 of two velocity vector samples defined further from the transducer array and a right image 3020 of two velocity vector samples defined close to the array that have had transducers whose wavefronts are in the direction opposite to the defined control direction removed. Using the velocity weighting function specified in these two disjoint regions makes up a single region in each case, wherein the region is optimizing for directional vectors representing particle velocity of the medium in two separate points.

[00222] Note that in trying to find the maximum amount of velocity in the x direction does not necessarily mean that the resulting particle velocity direction will be in entirely in x . Whereas Figure 30 shows the result of filtering out conflicting transducers that offer wavefronts that are opposed to the desired direction, shown in Figure 31 is a left image 3110 and a right image 3120 that are the result in which the transducer wavefront directions have not been filtered using the wavefront criterion. Conflicting transducers are in the case of Figure 31 free to interfere, whose pathology is apparent in the form of a standing wave structure horizontally. This is because the wavefront direction conveys directional information that is missing from the particle velocity, the particle velocity does not yield information pertaining to which direction the wave originated from, enabling unwanted interference.

[00223] This is also important for the simple case of a control point with normal and not just a control region. The case of a control point is effectively a simplified control region wherein it is approximated with a single point sample of particle velocity and no weighting factor. In either case, not filtering out ‘unhelpful’ transducers yields a similar result with elements of the effects of a ‘standing wave’ phenomenon. This is because a particle velocity vector with complex components cannot determine the difference between waves travelling in different directions: a wave travelling left-to-right with phase 0 is in this model for a point sample mathematically identical to a right-to-left travelling wave with phase π . As a result, if this is not dealt with correctly in the generation of the control point or region basis, the optimization will attempt to use both, generating a solution with standing wave components that do not contribute meaningfully, creating extra nulls, regions with no field, that are unwanted. This is an important point, as control points entered as basis functions into linear optimizations, eigensystems or even simple formula without this filtering step will exhibit these artefacts.

[00224] These techniques described may be integrated into a device to create an acoustic phased array system that is capable of producing arbitrarily configured regions and beams by

iterating a single eigensystem as defined here or in the simple case creating a single point or region. In the single point or region case an optimization may be applied to generate a simple formula which shows that the amplitudes and phases needed to drive the transducer elements are the amplitudes and phases expressed by evaluating the point or region. The iteration of the global method and repeated application through time may be changed and modified in a time dependent manner to produce a variety of haptic effects. This may be applied to both the desired field and the transducer coefficients derived, as they are functionally related.

[00225] III. Clauses Describing Embodiments of the Invention

1. A method of, on a monochromatic acoustic field:
 - Evaluating a distribution of acoustic pressure to a single value by integrating the product of the acoustic pressure field function with a weighting function, either of which function may be complex-valued and are derived from physical transducer properties.
 - Applying zero or more null points to be added to the null space of the matrix.
 - Using this as the objective function to be maximised in an optimisation, where the optimisation result takes the form of the eigenvectors of a matrix formed from cross-multiplied terms of the objective function for each transducer.
 - Summing one or more such matrices to generate an eigensystem where one or more eigenvectors produces the specified field.
2. A method of, on a monochromatic acoustic field:
 - Evaluating a distribution of acoustic particle velocity to a single vector value by integrating the dot product of a real three-dimensional weighting vector field with a particle velocity field with a weighting function, where either of which functions may have complex-valued components and are derived from physical transducer properties.
 - Applying zero or more null points to be added to the null space of the matrix.
 - Using this as the objective function to be maximised in an optimisation, where the optimisation result takes the form of the eigenvectors of a matrix formed from cross-multiplied terms of the objective function for each transducer.
 - Summing one or more such matrices to generate an eigensystem where one or more eigenvectors produces the specified field.
3. A method of, on a monochromatic acoustic field:

- Evaluating a distribution of acoustic particle velocity to a single vector value by integrating the dot product of a real three-dimensional weighting vector field with a particle velocity field with a weighting function, where either of which functions may have complex-valued components and are derived from physical transducer properties.
 - Setting to zero the contribution of transducers whose wavefronts at the point have an opposing direction to the normal specified at each evaluated sub-region of the integral.
 - Applying zero or more null points to be added to the null space of the matrix.
 - Using this as the objective function to be maximised in an optimisation, where the optimisation result takes the form of the eigenvectors of a matrix formed from cross-multiplied terms of the objective function for each transducer.
 - Summing one or more such matrices to generate an eigensystem where one or more eigenvectors produces the specified field.
4. A method of, on a monochromatic acoustic field:
- Evaluating a distribution of acoustic pressure to a single value by integrating the product of the acoustic pressure field function with a weighting function, either of which function may be complex-valued and are derived from physical transducer properties.
 - Applying zero or more null points to be added to the null space of the matrix.
 - Using this as the objective function to be maximised in an optimisation, where the optimisation result takes the form of the eigenvectors of a matrix formed from cross-multiplied terms of the objective function for each transducer.
 - Using the contribution of each transducer to describe the optimal phases and amplitudes for the transducer input signal.
5. A method of, on a monochromatic acoustic field:
- Evaluating a distribution of acoustic particle velocity to a single vector value by integrating the dot product of a real three-dimensional weighting vector field with a particle velocity field with a weighting function, where either of which functions may have complex-valued components and are derived from physical transducer properties.
 - Applying zero or more null points to be added to the null space of the matrix.

- Using this as the objective function to be maximised in an optimisation, where the optimisation result takes the form of the eigenvectors of a matrix formed from cross-multiplied terms of the objective function for each transducer.
 - Using the contribution of each transducer to describe the optimal phases and amplitudes for the transducer input signal.
6. A method of, on a monochromatic acoustic field:
- Evaluating a distribution of acoustic particle velocity to a single vector value by integrating the dot product of a real three-dimensional weighting vector field with a particle velocity field with a weighting function, where either of which functions may have complex-valued components and are derived from physical transducer properties.
 - Setting to zero the contribution of transducers whose wavefronts at the point have an opposing direction to the normal specified at each evaluated sub-region of the integral.
 - Applying zero or more null points to be added to the null space of the matrix.
 - Using this as the objective function to be maximised in an optimisation, where the optimisation result takes the form of the eigenvectors of a matrix formed from cross-multiplied terms of the objective function for each transducer.
 - Using the contribution of each transducer to describe the optimal phases and amplitudes for the transducer input signal.
7. The method of clauses 1 to 6, where further a device implements the method and uses the result to drive acoustic transducers.
8. The method of clauses 1 to 6, where further a device implements the method and uses the result as a portion of a linear basis set which is used to drive acoustic transducers.
9. The method of clauses 1 to 6 where an acoustic beam is produced
10. The method of clauses 1 to 6 where an acoustic volume is produced
11. The method of clauses 1 to 6 where an acoustic surface is produced
12. A method of, on an acoustic field:
- Evaluating an acoustic particle velocity to a single value at a point using the linear combination of fields derived from physical transducer properties.
 - Setting to zero the contribution of transducers whose wavefronts at the point have an opposing direction to the normal at the control point.

- Using this as a basis function in an optimisation problem for finding the phases and amplitudes for the transducer input signal.
13. A method of, on an acoustic field:
- Evaluating an acoustic particle velocity to a single value at a point using the linear combination of fields derived from physical transducer properties.
 - Setting to zero the contribution of transducers whose wavefronts at the point have an opposing direction to the normal at the control point.
 - Using this as the objective function to be maximised in an optimisation, where the optimisation result takes the form of the eigenvectors of a matrix formed from cross-multiplied terms of the objective function for each transducer, which is then finding the phases and amplitudes for the transducer input signal.
14. A method of, on an acoustic field:
- Evaluating an acoustic particle velocity to a single value at a point using the linear combination of fields derived from physical transducer properties.
 - Setting to zero the contribution of transducers whose wavefronts at the point have an opposing direction to the normal at the control point.
 - Using the contribution of each transducer to describe the optimal phases and amplitudes for the transducer input signal.
15. The method of clauses 1 to 14, wherein the transducers produce a haptic effect when the resulting structure is applied as an input signal and modulated.
16. The method of clauses 1 to 14, wherein the transducers produce a haptic effect when the resulting structure is applied as an input signal and is modified through time.

[00226] 6. EVALUATION OF COMPLEX-VALUED TRANSCENDENTAL FUNCTIONS, MULTIPLICATION AND DIVISION IN HARDWARE

[00227] By aiming to convert between $r + i\theta \leftrightarrow 2^r(e^{\pi/2})^{i\theta}$ instead of $r + i\theta \leftrightarrow e^{r+i\theta}$, the BKM algorithm may be reengineered to yield a greatly simplified technique which is more useful in physical hardware implementations.

[00228] I. Introduction

[00229] There exist many techniques for the computation of trigonometric and other special functions in hardware. One such method is the BKM method, which exchanges complex logarithms for complex exponentiations and vice versa, generating these special functions as a side effect of moving from one representation to another. Of special interest are those

dealing with complex-valued functions as situations in which such functions are found to be necessary are encountered often.

[00230] An efficient method to move between logarithm and exponential forms enables multiplication and division to be carried out in the logarithmic space as simple addition and subtraction. Due to the nature of the process this can also be exploited to produce transcendental functions related to the exponential easily. Since this can be generated from a single algorithm, this results in reduced implementation size, since only one element requires implementation, and reduced development time, as only one method need be implemented and optimized.

[00231] An existing algorithm, BKM, implements this approach, but includes three disadvantageous requirements which are almost or completely removed in the proposed method:

[00232] Hardware division among other difficult to implement functions to pre- and post-processing steps on the data input and output.

[00233] Irrational values and factors that create rounding issues for efficient implementations.

[00234] Extra operations that are required to integrate with floating point types.

[00235] The next section will describe how the proposed method has been designed to be free of these drawbacks.

[00236] II. Motivation and comparative range reduction

[00237] The process is centered on Euler's formula, $e^{i\theta} = \cos \theta + i \sin \theta$, but where e^r is a further coefficient as the formula has a complex generalization which yields, $e^{r+i\theta} = e^r(\cos \theta + i \sin \theta)$. The complex function which the BKM algorithm embodies is exactly this, making the implemented process $r + i\theta \leftrightarrow e^{r+i\theta}$. The algorithm heavily uses range reduction techniques to fit the complex value $r + i\theta$ into a range where convergence may be guaranteed. This is not a function that has a neat implementation, to the extent that it is likely easier to evaluate using the proposed method by pre- or post-processing with an irrational scaling factor to achieve $r + i\theta \leftrightarrow e^{r+i\theta}$. In many cases, this irrational scaling factor may be eliminated to provide a workflow whose inputs and outputs are intended to be multiples of physically or computationally meaningful rational numbers.

[00238] When processing a complex exponentiation, the exponentiation may be split into a real part and an imaginary part, which use the same input data but produce different results. This is in some ways similar to a Cartesian-to-polar transform.

[00239] Considering just the real part of the logarithm to be exponentiated, this contains all the modulus information, similar to the r component of the Cartesian-to-polar transform. This is traditionally the irrational constant e raised to the power of the real part. To feed it into the algorithm this first has to undergo range reduction, meaning that we must find the remainder, dividing modulo the irrational value $\ln 2$ such that it fits inside the convergence domain. This involves a hardware division (the 1994 paper recommends the SRT division technique), and so this adds a highly expensive technique in terms of the logic required to implement a division alongside the algorithm.

[00240] The premise of the proposed algorithm is that powers of two should be used to make the range reduction straightforward. With this motivation, the real part should be interpreted as two raised to the given power which is the real part, so an exponentiation with base two. The whole number of the real logarithm input then may refer to bit shifts on the output and only the fractional part of the relative binary logarithm input needs to fit in the domain of convergence. The integer part of the real component of the complex logarithm is then taken by round-to-nearest and exponentiated separately in this way.

[00241] Considering just the imaginary part of the logarithm to be exponentiated, this contains all the argument or angle information, similar to the θ component of the Cartesian-to-polar transform. This is also traditionally raised to the power e , but e is not the problem here. The problem is π this time, as the imaginary logarithm part can be construed as an angle, but to restrict the domain of convergence we again have to find the remainder from an integer fraction of 2π so that rotational symmetry may be exploited. This involves yet more division processes that are expensive in terms of logic size.

[00242] If instead of the base e logarithm, base two could be used. This obtains the expansion $2^{i\theta} = \cos(\theta \ln 2) + i \sin(\theta \ln 2)$. This is yet more unwanted irrational values, however, if we specifically use the base $e^{2\pi}$ we can use θ to express the number of rotations. Unfortunately, using this base leads to a method that does not converge. The best trade-off between technique and convergence range yields base $e^{\frac{\pi}{2}}$. This means each integer value represents a quadrant. Each multiple of four is a whole rotation. Bitwise operations may then be used to extract the number of rotations. The number of quadrants may be applied to the output as a switching of real and imaginary numbers and signs. Again, this yields a technique wherein

the integer part of the imaginary component of the complex logarithm is taken by round-to-nearest and processed separately in this way.

[00243] This leaves the fractional part of the logarithm still to be exponentiated in the range $R = [-0.5, +0.5) + i[-0.5, +0.5)$ as shown in the right panel 3220 of Figure 32, which fits into the range of convergence for the main algorithm. In Figure 32, errors in the exponential plotted for the input to the logarithm-to-exponential process in the region $[-2.0, +2.0)$, where the right-hand side is positive real and the lower part of the image positive imaginary. With 50 fractional bits and a similar number of iterations, the black area represents values that have converged to an accuracy greater than 2^{-45} . The left panel 3210 shows the domain of convergence where gray is mapped to errors greater than 2^{-45} . The right panel 3220 shows the final required convergence region $R = [-0.5, +0.5) + i[-0.5, +0.5)$ in white.

[00244] Discussion of the main algorithm for the fractional part of the logarithm input is described later.

[00245] The reverse process is also achieved, taking the value with its exponentiated real and imaginary parts and reducing it to its original logarithms. To begin, we find the quadrant, offset by 45° , that the complex value belongs to. This is achieved by determining which component has the largest absolute value; whether it is real or imaginary and its sign. This locates which of these quadrants offset in angle the value belongs to. This is then used to compute a rotation to move the values into the first quadrant whose area is centered on the positive part of the real line, as shown in the center panel 3320 of Figure 33. Figure 33 shows errors in the exponential plotted for the input to the exponential-to-logarithm process in the region $[-2.0, +2.0)$, where the right-hand side is positive real and the lower part of the image positive imaginary. With 50 fractional bits and a similar number of iterations, the black area represents values that have converged to an accuracy greater than 2^{-45} . The left panel 3310 shows the domain of convergence where gray is mapped to errors greater than 2^{-45} . The middle panel 3320 shows the quadrant centered on the real line using inverted grayscale. The right panel 3330 shows the final required convergence region $R = [+0.5, +1.0) + i[-\Re(R), +\Re(R))$ in white.

[00246] The imaginary part of the logarithm output is pre-loaded with the equivalent imaginary integer part of the output complex logarithm. Due to the 45° offset, a further test would be necessary for the complex branch cut to fall on the negative part of the real line, but if a two's complement system wrap around is used this may be ignored, as given the

necessary bits the imaginary part of the logarithm measured in quadrants can be easily restricted to lie in the interval $[-2.0, 2.0)$.

[00247] At this point as it is certain that the value lies in the first quadrant, the real part is known to be greater than or equal to the imaginary part in absolute value and the real part is known to be positive. To find the integer part of the real part of the logarithm, we can count the leading zeroes in the fixed-point representation. The value of the subtraction of this count from the number of bits in the fixed-point representation is effectively

$\lfloor \log_2 \Re(z_{\text{input, rotated}}) \rfloor + k$, where k depends on the representation of the bits and so is a representation of the integer part of the binary logarithm, which can be immediately added onto the real part of the logarithm register. The existing data, both real and imaginary, can then be shifted up by the number of bit places required to cause the real part to lie in the interval $[+0.5, +1.0)$, as shown in the right panel of Figure 33, which is determined by the initial count leading zeroes and the k that depends on the representation of the bits. It should be noted that while this may not be the nearest integer representation of the output real logarithm, it is integer and brings the remainder of the process into the region of convergence with the correct format.

[00248] At the end of this process, the integer part of the output complex logarithm has been determined, leaving the fractional part of the output undetermined and the input range $R = [+0.5, +1.0) + i[-\Re(R), +\Re(R))$, which fits into the range of convergence for the main algorithm. Discussion of the main algorithm for the substantially fractional part of the logarithm output will be described below.

[00249] The result is that the function $r + i\theta \leftrightarrow 2^r(e^{i\pi/2})^{i\theta}$ is the process now modelled as can be seen from Figure 34. In Figure 34, plotted errors in the exponential for the exponentiation 3410 and logarithm process 3430 domains for the sufficiently fractional parts only, alongside the range reduced counterparts that make up the final composite values for the exponentiation 3420 and logarithm 3440 processes. The noisy gray denotes areas where correct convergence causes small numerical error to dominate the error measurement, showing that the convergence has been achieved everywhere for exponentiation 3420 and logarithm 3440.

[00250] Changing the base of the result still means it has all the properties of logarithms and exponents, meaning that the real parts of logarithms can be added and subtracted to provide exponentiated multiplication and division. Equally, the angles behave the same, only that

they are represented as integer multiples of a quadrants, so 4 to a rotation, a value easily manipulated in hardware logic. In their logarithm form, adding and subtracting the imaginary values is the same as rotating clockwise and anti-clockwise around zero in the complex-valued exponentiated space. For the complex values taking their combined real and imaginary parts irrespective of the fact that they are composed of different bases, adding and subtracting the real and imaginary parts, can used to implement complex multiplication and division.

[00251] These results may be trivially extended to floating-point representations as in these formats the required pre-processing is already mostly achieved by virtue of the mantissa and exponent form. If the values in base e are absolutely required and most of the time they are not, then the real part of the output logarithm may be multiplied by $\ln 2$ and the imaginary part by $\frac{\pi}{2}$. Equally, to take in logarithm data in base e format, the values can be taken as reciprocals $\left(\frac{1}{\ln 2}, \frac{2}{\pi}\right)$ and multiplied component-wise to convert, before engaging in the initial range reduction to exponentiate with this method.

[00252] III. General procedure

[00253] The new algorithm proceeds in a similar way to the BKM technique to find or use the fractional part of the logarithm. There is a set of values $d_n \in \{0, 1, i, -1, -i, 1 + i, 1 - i, -1 + i, -1 - i\}$ from which d_n is chosen on each algorithm step n . Choices of which d_n to use are made by comparing a truncated number with a set bit count which may be a subset of the bits from the input registers. Once a d_n is selected, look-up tables of precalculated logarithms are required to balance the $1 + d_n 2^{-n}$ (bit shifting (potentially with rounding applied on the most significant unused bit) and addition/subtraction arithmetic) multiplications applied to the registers holding exponentiated values with the equivalent logarithms to be subtracted from the registers holding logarithm values. Depending whether the process is to find a logarithm or to compute exponentiation, there may be extra shifting, extra arithmetic or both in order to arrange the system into a convergent configuration. At the end of some N iterations, the desired value, either exponential or logarithm, is found to within an accuracy that depends on the number of iterations N and the number of bits used for the representation.

[00254] IV. Look-up tables of logarithms

[00255] The proposed method differs from existing techniques in that these logarithms must now reflect the exponentiation as $2^r (e^{\pi/2})^{i\theta}$. Due to this, the real part of the logarithm table

must now be preloaded with $\log_2 \|1 + d_n 2^{-n}\|$, and the imaginary part $2/\pi \arg(1 + d_n 2^{-n})$ or $\Im(\log_{e^{\pi/2}}(1 + d_n 2^{-n}))$, for $n = 1, \dots, N$, in the desired fixed point format.

[00256] Although we have chosen 2 and $e^{\pi/2}$ for the logarithm base, it is likely that a different choice of base may give an algorithm for the logarithm system that does not have a different first iteration. These values have been chosen because they represent a design that may achieve both processes using the same hardware. If a separate implementation using distinct hardware is to be used for the logarithm mode as from the exponentiation mode, then this may not be the best choice of base.

[00257] V. Method for the fractional part of the logarithm

[00258] The core method for the new algorithm largely follows from those used by the BKM process, with some new key differences. The comparison thresholds have been derived anew to accommodate the new bases for the logarithms and exponentiation and the changed domain of convergence required for the new technique. There is also a slightly modified comparison threshold for the first iteration of the logarithm process to achieve convergence within the required domain.

[00259] VI. Exponentiation mode

[00260] The exponentiation mode described in this section may also be modified to provide a complex multiplication with the exponential value. If this is to be achieved, this must be preloaded before the range reduction steps if the output is to be correct. It should also be noted that this would replace the output exponentiation value and so should be not used if this value is required. It is also feasible to store and wait to apply the solution from the integer parts of the logarithm (the value $z_{\text{integer, output}}$) to the end of the process. This may reduce the storage required for the intermediate registers in which the processing occurs, although this should be weighed against the extra requirements of storage needed for the integer parts of the solution.

[00261] Alternatively, multiplication with the final exponentiation value may be achieved in parallel by creating extra registers and ensuring that equivalent operations occur in these extra registers. In this way, the exponentiation process may complex multiply the output exponential value with almost arbitrarily many other complex values.

[00262] Assuming the fractional part of the input logarithms to be the input, the algorithm for the domain of convergence $z_{\text{input}} \in \mathbb{R} = [-0.5, +0.5) + i[-0.5, +0.5)$ is:

[00263] 1. Assuming there are four basic registers, labelled $\Re(z_{\log})$, $\Im(z_{\log})$, $\Re(z_{\exp})$ and $\Im(z_{\exp})$. Alongside, there are two extra slave multiplication registers $\Re'(z_{\exp})$ and $\Im'(z_{\exp})$ to demonstrate how the method operates when used for auxiliary complex multiplication. The initial values of these registers are:

$$\begin{aligned}\Re(z_{\log}) &:= 2\Re(z_{\text{input}}), \\ \Im(z_{\log}) &:= 2\Im(z_{\text{input}}), \\ \Re(z_{\exp}) &:= \Re(z_{\text{integer input, output}} \times z_{\text{premultiply}}), \\ \Im(z_{\exp}) &:= \Im(z_{\text{integer input, output}} \times z_{\text{premultiply}}),\end{aligned}$$

where $z_{\text{premultiply}} := 1.0$, if there are no requirements for pre-multiplication. The slave multiplication registers may also be similarly constructed with:

$$\begin{aligned}\Re'(z_{\exp}) &:= \Re(z_{\text{integer input, output}} \times z_{\text{premultiply}'}), \\ \Im'(z_{\exp}) &:= \Im(z_{\text{integer input, output}} \times z_{\text{premultiply}'}).\end{aligned}$$

[00264] 2. Iterate through the values $1, \dots, N$ as the index n :

[00265] 3. Truncate $\Re(z_{\log})$ to form $\Re(z_{\log, \text{test}})$ such that it has three bits; one sign bit and two integer bits in two's complement such that the range is $[-4.0, +4.0)$ with the smallest change being 1.

[00266] 4. Truncate $\Im(z_{\log})$ to form $\Im(z_{\log, \text{test}})$ such that it has three bits; one sign bit, one integer bit and one fraction bit in two's complement such that the range is $[-2.0, +2.0)$ with the smallest change being 0.5.

[00267] 5. Test the 3-bit values to determine d_n :

$$\begin{aligned}\Re(d_n) &:= \begin{cases} -1, & \text{if } \Re(z_{\log, \text{test}}) < -0b01 \text{ (or } < -1.0), \\ +1, & \text{if } \Re(z_{\log, \text{test}}) > +0b00 \text{ (or } \geq +1.0), \\ 0, & \text{if neither,} \end{cases} \\ \Im(d_n) &:= \begin{cases} -i, & \text{if } \Im(z_{\log, \text{test}}) < -0b01 \text{ (or } < -0.5), \\ +i, & \text{if } \Im(z_{\log, \text{test}}) > +0b00 \text{ (or } \geq +0.5), \\ 0, & \text{if neither.} \end{cases}\end{aligned}$$

[00268] 6. Apply the multiplication to the exponential registers using shifting and arithmetic to compute:

$$\Re(z_{\exp, n}) := \Re(z_{\exp, n-1}) +$$

$$\begin{aligned}
& + \begin{cases} -2^{-n}\Re(z_{\text{exp}, n-1}) = -\text{sra}(\Re(z_{\text{exp}, n-1}), n), & \text{if } \Re(d_n) = -1, \\ +2^{-n}\Re(z_{\text{exp}, n-1}) = +\text{sra}(\Re(z_{\text{exp}, n-1}), n), & \text{if } \Re(d_n) = +1, \\ 0, & \text{if neither,} \end{cases} \\
& + \begin{cases} +2^{-n}\Im(z_{\text{exp}, n-1}) = +\text{sra}(\Im(z_{\text{exp}, n-1}), n), & \text{if } \Im(d_n) = -i, \\ -2^{-n}\Im(z_{\text{exp}, n-1}) = -\text{sra}(\Im(z_{\text{exp}, n-1}), n), & \text{if } \Im(d_n) = +i, \\ 0, & \text{if neither.} \end{cases}
\end{aligned}$$

And:

$$\begin{aligned}
& \Im(z_{\text{exp}, n}) := \Im(z_{\text{exp}, n-1}) + \\
& + \begin{cases} -2^{-n}\Im(z_{\text{exp}, n-1}) = -\text{sra}(\Im(z_{\text{exp}, n-1}), n), & \text{if } \Re(d_n) = -1, \\ +2^{-n}\Im(z_{\text{exp}, n-1}) = +\text{sra}(\Im(z_{\text{exp}, n-1}), n), & \text{if } \Re(d_n) = +1, \\ 0, & \text{if neither,} \end{cases} \\
& + \begin{cases} -2^{-n}\Re(z_{\text{exp}, n-1}) = -\text{sra}(\Re(z_{\text{exp}, n-1}), n), & \text{if } \Im(d_n) = -i, \\ +2^{-n}\Re(z_{\text{exp}, n-1}) = +\text{sra}(\Re(z_{\text{exp}, n-1}), n), & \text{if } \Im(d_n) = +i, \\ 0, & \text{if neither.} \end{cases}
\end{aligned}$$

Do the same to any auxiliary registers such as $\Re'(z_{\text{exp}})$ and $\Im'(z_{\text{exp}})$ to apply the multiplication process to these also.

[00269] 7. Subtract the corresponding entry in the logarithm tables from the registers:

$$\begin{aligned}
& \Re(z_{\log, n}) := \Re(z_{\log, n-1}) - 2^n \log_2 \|1 + d_n 2^{-n}\|, \\
& = \Re(z_{\log, n-1}) - \text{sll}(\text{table}_{\Re}[d_n, n], n), \\
& \Im(z_{\log, n}) := \Im(z_{\log, n-1}) - 2^n \left(2/\pi \arg(1 + d_n 2^{-n})\right), \\
& = \Im(z_{\log, n-1}) - \text{sll}(\text{table}_{\Im}[d_n, n], n).
\end{aligned}$$

[00270] 8. Shift the logarithm registers over to the left by one place, doubling the value of their contents.

$$\begin{aligned}
& \Re(z_{\log, n}) := \text{sll}(\Re(z_{\log, n}), 1), \\
& \Im(z_{\log, n}) := \text{sll}(\Im(z_{\log, n}), 1).
\end{aligned}$$

[00271] 9. Return to step 2 for the next iteration, until N is reached, at which point the registers will contain their final values:

$$\begin{aligned}
& \Re(z_{\text{exp}, N}) := \Re\left(z_{\text{integer input, output}} \times z_{\text{premultiply}} \times 2^{\Re(z_{\text{input}})} (e^{\pi/2})^{i\Im(z_{\text{input}})}\right), \\
& \Im(z_{\text{exp}, N}) := \Im\left(z_{\text{integer input, output}} \times z_{\text{premultiply}} \times 2^{\Re(z_{\text{input}})} (e^{\pi/2})^{i\Im(z_{\text{input}})}\right),
\end{aligned}$$

[00272] And:

$$\Re'(z_{\text{exp}}, N) := \Re\left(z_{\text{integer input, output}} \times z_{\text{premultiply}} \times 2^{\Re(z_{\text{input}})} (e^{\pi/2})^{i\Im(z_{\text{input}})}\right),$$

$$\Im'(z_{\text{exp}}, N) := \Im\left(z_{\text{integer input, output}} \times z_{\text{premultiply}} \times 2^{\Re(z_{\text{input}})} (e^{\pi/2})^{i\Im(z_{\text{input}})}\right).$$

[00273] Having appreciated the form of the process, it is easy to find other testing procedures that are convergent, even sometimes in the required domain, by forming

$\Re(z_{\text{log, test}})$, $\Im(z_{\text{log, test}})$ or both using different number of bits or different comparison values, although we have endeavored to reduce complexity by specifying the required value tests in the simplest known form.

[00274] VII. Logarithm mode

[00275] The logarithm mode described in this section may also be modified to provide a complex division with the input value. If this is to be achieved, it can only occur in parallel by creating extra registers and ensuring that equivalent operations occur in these extra registers. This contrasts with the auxiliary multiplication, where the original register could be overloaded, which cannot be achieved here because the modification of the exponentiation register in this mode would prevent convergence of the algorithm. However, using auxiliary registers can circumvent this, allowing the logarithm process to, if desired, produce complex-valued division of almost arbitrarily many other complex values with the value input to this process as denominator.

[00276] 1. Assuming the fractional part of the output logarithms to be the output, the algorithm for the domain of convergence $z_{\text{input}} \in \mathbb{R} = [+0.5, +1.0) + i[-\Re(R), +\Re(R))$ is:

[00277] 2. Assuming there are four basic registers, labelled $\Re(z_{\text{log}})$, $\Im(z_{\text{log}})$, $\Re(z_{\text{exp}})$ and $\Im(z_{\text{exp}})$. Alongside, there are two extra slave division registers $\Re'(z_{\text{exp}})$ and $\Im'(z_{\text{exp}})$ to demonstrate how the method operates when used for auxiliary complex division. The initial values of these registers are:

$$\begin{aligned}\Re(z_{\text{log}}) &:= \Re(z_{\text{integer output, output}}), \\ \Im(z_{\text{log}}) &:= \Im(z_{\text{integer output, output}}), \\ \Re(z_{\text{exp}}) &:= 2(\Re(z_{\text{input}}) - 1.0), \\ \Im(z_{\text{exp}}) &:= 2\Im(z_{\text{input}}),\end{aligned}$$

The slave division registers may also be similarly constructed with:

$$\Re'(z_{\text{exp}}) := \Re(z_{\text{numerator}} \div z_{\text{integer output, input}}),$$

$$\mathfrak{I}'(z_{\text{exp}}) := \mathfrak{I}(z_{\text{numerator}}' \div z_{\text{integer output, input}}).$$

Noting that neither the doubling or the -1.0 is applied to the registers $\mathfrak{R}'(z_{\text{exp}})$ and $\mathfrak{I}'(z_{\text{exp}})$.

[00278] 2. Iterate through the values $1, \dots, N$ as the index n :

[00279] 3. Truncate $\mathfrak{R}(z_{\text{exp}})$ to form $\mathfrak{R}(z_{\text{exp, test}})$ such that it has five bits; one sign bit, one integer bit and three fraction bits in two's complement such that the range is $[-2.0, +2.0)$ with the smallest change being 0.125.

[00280] 4. Truncate $\mathfrak{I}(z_{\text{exp}})$ to form $\mathfrak{I}(z_{\text{exp, test}})$ such that it has five bits; one sign bit, two integer bits and two fraction bits in two's complement such that the range is $[-4.0, +4.0)$ with the smallest change being 0.25.

[00281] 5. Test the two 5-bit values (with modified comparisons for the first iteration) to determine d_n :

$$\mathfrak{R}(d_n) := \begin{cases} +1, & \text{if } \mathfrak{R}(z_{\text{exp, test}}) < -0b0110 \text{ (or } < -0.75), \\ \text{or if } n = 1 \wedge \mathfrak{R}(z_{\text{exp, test}}) \leq -0b0110 \text{ (or } < -0.625), \\ -1, & \text{if } \mathfrak{R}(z_{\text{exp, test}}) > +0b0001 \text{ (or } \geq +0.25), \\ 0, & \text{if neither,} \end{cases}$$

$$\mathfrak{I}(d_n) := \begin{cases} +i, & \text{if } \mathfrak{I}(z_{\text{exp, test}}) < -0b0010 \text{ (or } < -0.25), \\ \text{or if } n = 1 \wedge \mathfrak{I}(z_{\text{exp, test}}) \leq -0b0010 \text{ (or } < -0.125) \\ -i, & \text{if } \mathfrak{I}(z_{\text{exp, test}}) > +0b0001 \text{ (or } \geq +0.25), \\ \text{or if } n = 1 \wedge \mathfrak{I}(z_{\text{exp, test}}) \geq +0b0001 \text{ (or } \geq +0.125) \\ 0, & \text{if neither.} \end{cases}$$

[00282] 6. Apply the multiplication to the exponential registers using shifting and arithmetic to compute:

$$\mathfrak{R}(z_{\text{exp, n}}) := \mathfrak{R}(z_{\text{exp, n-1}}) + \begin{cases} -2^{-n}\mathfrak{R}(z_{\text{exp, n-1}}) = -\text{sra}(\mathfrak{R}(z_{\text{exp, n-1}}), n), & \text{if } \mathfrak{R}(d_n) = -1, \\ +2^{-n}\mathfrak{R}(z_{\text{exp, n-1}}) = +\text{sra}(\mathfrak{R}(z_{\text{exp, n-1}}), n), & \text{if } \mathfrak{R}(d_n) = +1, \\ 0, & \text{if neither,} \end{cases}$$

$$+ \begin{cases} +2^{-n}\Im(z_{\text{exp}, n-1}) = +\text{sra}(\Im(z_{\text{exp}, n-1}), n), & \text{if } \Im(d_n) = -i, \\ -2^{-n}\Im(z_{\text{exp}, n-1}) = -\text{sra}(\Im(z_{\text{exp}, n-1}), n), & \text{if } \Im(d_n) = +i, \\ 0, & \text{if neither.} \end{cases}$$

And:

$$\begin{aligned} \Im(z_{\text{exp}, n}) &:= \Im(z_{\text{exp}, n-1}) + \\ &+ \begin{cases} -2^{-n}\Im(z_{\text{exp}, n-1}) = -\text{sra}(\Im(z_{\text{exp}, n-1}), n), & \text{if } \Re(d_n) = -1, \\ +2^{-n}\Im(z_{\text{exp}, n-1}) = +\text{sra}(\Im(z_{\text{exp}, n-1}), n), & \text{if } \Re(d_n) = +1, \\ 0, & \text{if neither,} \end{cases} \\ &+ \begin{cases} -2^{-n}\Re(z_{\text{exp}, n-1}) = -\text{sra}(\Re(z_{\text{exp}, n-1}), n), & \text{if } \Im(d_n) = -i, \\ +2^{-n}\Re(z_{\text{exp}, n-1}) = +\text{sra}(\Re(z_{\text{exp}, n-1}), n), & \text{if } \Im(d_n) = +i, \\ 0, & \text{if neither.} \end{cases} \end{aligned}$$

[00283] Do the same to any auxiliary registers such as $\Re'(z_{\text{exp}})$ and $\Im'(z_{\text{exp}})$ to apply the division process to these.

[00284] 7. Add d_n to $\Re(z_{\text{exp}, n})$ and $\Im(z_{\text{exp}, n})$:

$$\begin{aligned} \Re(z_{\text{exp}, n}) &:= \Re(z_{\text{exp}, n}) + \Re(d_n), \\ \Im(z_{\text{exp}, n}) &:= \Im(z_{\text{exp}, n}) + \Im(d_n). \end{aligned}$$

Do not add this to $\Re'(z_{\text{exp}})$ and $\Im'(z_{\text{exp}})$.

[00285] 8. Shift the exponential registers over to the left by one place, doubling the value of their contents.

$$\begin{aligned} \Re(z_{\text{exp}, n}) &:= \text{sll}(\Re(z_{\text{exp}, n}), 1), \\ \Im(z_{\text{exp}, n}) &:= \text{sll}(\Im(z_{\text{exp}, n}), 1). \end{aligned}$$

Do not double the values of $\Re'(z_{\text{exp}})$ and $\Im'(z_{\text{exp}})$.

[00286] 9. Subtract the corresponding entry in the logarithm tables from the registers:

$$\begin{aligned} \Re(z_{\log, n}) &:= \Re(z_{\log, n-1}) - 2^n \log_2 \|1 + d_n 2^{-n}\|, \\ &= \Re(z_{\log, n-1}) - \text{sll}(\text{table}_{\Re}[d_n, n], n), \\ \Im(z_{\log, n}) &:= \Im(z_{\log, n-1}) - 2^n \left(2/\pi \arg(1 + d_n 2^{-n}) \right), \\ &= \Im(z_{\log, n-1}) - \text{sll}(\text{table}_{\Im}[d_n, n], n). \end{aligned}$$

[00287] 10. Return to step 2 for the next iteration, until N is reached, at which point the registers will contain their final values:

$$\begin{aligned}\Re(z_{\log, N}) &:= \log_2 \|\Re(z_{\text{input}}) + i\Im(z_{\text{input}})\|, \\ \Im(z_{\log, N}) &:= 2/\pi \arg(\Re(z_{\text{input}}) + i\Im(z_{\text{input}})),\end{aligned}$$

[00288] And:

$$\begin{aligned}\Re'(z_{\text{exp}, N}) &:= \Re(z_{\text{numerator}'} \div (z_{\text{integer output, input}} \times z_{\text{input}})), \\ \Im'(z_{\text{exp}, N}) &:= \Im(z_{\text{numerator}'} \div (z_{\text{integer output, input}} \times z_{\text{input}})).\end{aligned}$$

[00289] Having appreciated the form of the process, it is possible to find other testing procedures that are convergent, even sometimes in the required domain, by forming $\Re(z_{\text{exp, test}})$, $\Im(z_{\text{exp, test}})$ or both using different number of bits or different comparison values, although we have endeavored to reduce complexity by specifying the required value tests in the simplest known form.

[00290] VIII. Redundant number systems

[00291] The algorithm, like the BKM method, is ideal for supporting techniques that slow the local flow of data, allowing for integration into high speed logic implemented in hardware. Using a redundant number system scheme such as carry-save for the arithmetic steps, coupled with the reduced number of bits presented for comparison on each iteration allows for an algorithm such as is presented here to have a high-speed, low-area implementation that is ideal for inclusion into integrated circuit designs. Benefits of such an implementation would include being able to reach higher clock speeds, while scaling well to operands with larger numbers of bits, making this design ideal for fabricating low-cost, modern solutions to problems involving manipulation of complex values.

[00292] For ease of intuition, rounding has been neglected in the foregoing description of the algorithm and should not be interpreted as an implication that correct rounding has not been considered in this invention. It should be clear that capitalizing on the most significant bit not directly used in any arithmetic operation to achieve a result with reduced error through the application of rounding, when applied at each step may increase the accuracy of the results described herein.

[00293] IX. Higher-radix approaches

[00294] In a similar way to other algorithms that evaluate in a per-bit fashion, the steps may be merged to create a method that computes multiple bits at once. This is termed a higher radix algorithm, as the output is treated as having a different base than 2, meaning that one digit from the new higher base is calculated per step. The drawback of this is that the choices for the look-up tables and shifts would also have to be merged, creating a larger set of choices for the algorithm to work through per iteration.

[00295] X. Real algorithm

[00296] Many systems may benefit from using just the real part of the algorithm described herein. For instance, the sigmoidal logistic function, often used by neural network algorithms, may be generated by first determining the real 2^x or if pre-multiplied e^x using the exponentiation process. Then, adding one and dividing a preset constant by this value is possible, or continuing to use the E-mode followed by L-mode described by the original paper may produce the logistic function as:

$$L(x) = \frac{k}{1 + e^{-(ax+b)'}}$$

[00297] but the numerator k and further modifiers a and b remain configurable.

[00298] XI. Showing convergence

[00299] In Figure 35, convergence of the exponential and logarithm methods is shown by the graph 3500. Figure 35 shows testing each possible value in the necessarily convergent domain yields the worst case $\log_2$ error graphed here. The graph summarises the worst case behaviour from over one trillion complex test values (each real and imaginary value in the fractional domain that is remapped into has been tested for the number of iterations shown). To make the graphs comparable, the logarithmic result was translated into its exponentiated form before being tested for error. This ensures that the method is convergent for this many significant bits. This error may be reduced further by adding full rounding into the multiplication. The graph shows that as the precision and iterations increase, the error decreases similarly, showing convergence.

[00300] This graph 3500 show the absolute error of the worst case from all possible cases in the domain of convergence shrinking at a similar rate to the number of bits in the tables and iterations in the process. This reveals that the algorithm successfully converges linearly in bits of precision to the correct solution for all possible correct inputs for this number of bits of fixed point data.

[00301] XII. Clauses Describing Embodiments of the Invention

1. A method for the simultaneous conversion of an input 'logarithm' complex number into a resulting 'exponential' complex number where:
 - a. the modulus of the resulting complex number consists of two raised to a real power of the real input complex number part and;
 - b. the argument of the resulting complex number consists of two raised to an integer power rotations multiplied by the real value of the imaginary input complex number part.
2. A method for the simultaneous conversion of an input 'exponential' complex number into a resulting 'logarithm' complex number where:
 - a. the real binary logarithm of the real modulus of the input complex number forms the real part of the resulting complex number and;
 - b. the real number of two raised to an integer power rotations expressed in the argument of the input complex number forms the imaginary part of the resulting complex number.
3. The method of clauses 1 or 2 wherein the conversion is achieved substantially through bit shifting, comparisons and arithmetic on a binary representation of the complex number.
4. The method of clauses 1 wherein a reduction in input and output range is performed by processing values corresponding substantially to a rounding of the components of the input complex number separately to a further part.
5. The method of clause 2 wherein a reduction in input and output range is performed by processing values corresponding substantially to a rounding of the components of the output complex number separately to a further part.
6. The method of clause 1 wherein the output complex value may be initialised with a value not equal to one so that the application of the iterations will modify a non-unity value to compute a complex multiplication by the value originally intended to be output.
7. The method of clause 1 wherein further complex values are instantiated such that operations applied to the output complex value may also be applied to these further values, yielding a multiplication of these further complex values by the complex value originally intended to be output.
8. The method of clause 2 wherein further complex values are instantiated such that operations applied to the input complex value may also be applied to these further values, yielding a division of these further complex values by the complex value input.
9. The method of clause 1 wherein a redundant number system is used to implement the process of repeated arithmetic applied to the extant complex values.

10. The method of clause 2 wherein a redundant number system is used to implement the process of repeated arithmetic applied to the extant complex values.

11. A device designed to perform any method of clauses 1 through 10.

12. A data processor configured to perform any method of clauses 1 through 10.

[00302] 7. ACOUSTIC FIELDS BUILT FROM THE KRYLOV SUBSPACE OF A QUADRATIC PROBLEM

[00303] This disclosure describes a method to manipulate a quadratic problem when expressed as an eigensystem to produce a reduced rank basis set for inclusion into a linear system. By using a reduced rank basis an optimally information preserving set of (potentially reduced) basis functions may be automatically derived for inclusion into a complex-valued linear system to find the optimal complex-valued re-weightings of such basis functions to achieve a set of aims as described by a number of test functions.

[00304] The derivation of the reduced set of basis functions and how their inclusion into the appropriate stages of the implementation of an acoustic field linear optimization problem will be described.

[00305] I. Linear control regions

[00306] A linear control region, which is defined as an operator that performs an integral with an linear acoustic quantity, such as pressure, with a complex valued weighting function, $w(\chi)$, $\chi = \{x, y, z\}$:

$$f_p = \iiint_{\Omega} \overline{w(\chi)} \cdot p(\chi) dx dy dz,$$

may be used to evaluate an acoustic field. It should be noted that the complex conjugate is used to maximize function evaluation when $w(\chi)$ is paired with similar phase distributions in the acoustic quantity. Any approximation to this integral (such as sampling) can be used or indeed this may be used to approximate any linear operator over the field of acoustic pressure. A further integral may be constructed to optimize a control region using a vector particle velocity (wherein each vector component is complex-valued due to the harmonic motion of the field), given it is also linear, wherein:

$$f_v = \iiint_{\Omega} \overline{\mathbf{w}(\chi)} \cdot \mathbf{u}(\chi) dx dy dz,$$

with $\mathbf{u}$ the particle velocity of the medium, and $\mathbf{w}$ a three-dimensional vector field in which each dimension is complex-valued, which serves as a weighting function with the same

purpose as above. In this way, different patterns of acoustic quantities may be evaluated in a linear fashion, wherein a particular phase of the result will maximize the objective function when phases align with the ‘template’ created by the scalar or vector function in w .

[00307] It also should be noted that $w(\chi)$ may be largely zero (or comprised of zero vectors) to create a single evaluation of multiple seemingly disjoint regions. Any integral approximation (such as sampling) can be used as this formulation can extend to any linear operator over the field of a linear acoustic quantity.

[00308] II. Quadratic problems and eigenproblems

[00309] A phased array may be described by a series of complex valued coefficients that controls each transducer. Driving an array of acoustic transducers at a fixed frequency with such an input $\mathbf{x}$ vector to maximise an acoustic quantity encoded in a matrix $\mathbf{M}$ may be expressed as a linear algebra problem, so that the intention is to:

$$\begin{aligned} &\text{maximize} && \mathbf{x}^H \mathbf{M} \mathbf{x} \\ &\text{subject to} && \mathbf{x}^H \mathbf{x} = \mathbf{1}, \\ &\text{and} && \mathbf{x} \in \mathbb{C}^n. \end{aligned}$$

[00310] This may be solved by taking the dominant eigenvector of the matrix $\mathbf{M}$, as the statement of the problem is also the definition of an eigenvector, here $\mathbf{x}$. Building the matrix $\mathbf{M}$ may be achieved by first considering the linear acoustic quantities emitted by each transducer as a basis function in the global space of the acoustic quantity in question. This is a way to achieve the maximization of pressure and other acoustic quantities at a single point, described by an absolute position, χ . However, using the described control region integral, the $q(\chi)$ that governs the acoustic quantities emitted into the field of each individual transducing element may be exchanged for $f(\chi)$, yielding the column vector:

$$\mathbf{f} = \begin{bmatrix} f_1(\chi) \\ \vdots \\ f_n(\chi) \end{bmatrix},$$

[00311] By invoking linearity, the complex activation values driving each of the transducers may be multiplied leading to the vector:

$$\mathbf{m} = \begin{bmatrix} x_1 \cdot f_1(\chi) \\ \vdots \\ x_n \cdot f_n(\chi) \end{bmatrix},$$

where n is the number of transducers. Given that this $\mathbf{m}$ is the result of a linear operator operating on a linear field, it may be Hermitian transposed and multiplied by itself to generate a quadratic form. Finally,

$$\mathbf{m}^H \mathbf{m} = \mathbf{x}^H \mathbf{M} \mathbf{x} = \mathbf{x}^H \mathbf{f} \mathbf{f}^H \mathbf{x}.$$

[00312] By modifying the weighting function $w(\chi)$ and solving for the eigenvectors of the matrix $\mathbf{M}$, a variety of different beams, surfaces and other acoustic phenomenon may be created and maximized, producing optimally actuated regions in the acoustic field.

[00313] The null space generation technique discussed previously (“Eigensystem solutions for quadratic problems in mid-air haptic systems”) may also be used to extend this method to describe acoustic levitation, as well as allowing for null steering and the formation of mid-air texture by introducing null points into the field at positions to describe structure. It should be noted that this process increases the rank of the matrix, so simplification is not possible should this be used.

[00314] Equally, multiple optimizations may be carried out by adding together $\mathbf{M}$ matrices for a number of control regions. This may also include matrices that represent control points. This generates a matrix that has increased rank, so some simplification methods are not possible in the case that this is used.

[00315] III. Reduced rank approximation to a high rank matrix

[00316] The key point is that a high-rank matrix may be intentionally created so that a reduced-rank approximation may be generated. For example, a system that has many hundreds of individual points may be automatically regrouped such that very few are described by independent basis functions, creating the realization of a basis set that appears to be higher dimensional than it is to increase fidelity, reduce complexity and allow basis functions to be created according to need automatically.

[00317] Existing approaches generate a new basis function for each control point, or control region. By creating an eigensystem for the optimization of each control point or region and summing them, or adding in nulls, the eigensystem gains rank. It is possible through the repeated process of adding basis function matrices and/or nulls, that a high degree of rank is obtained by the final matrix. The linear system, due to constraints on resources, may only be able to process at most n basis functions using at most m test functions. The high degree of rank of this optimization matrix may be in this case n' where n'' optimisation matrices for generating individual basis functions have been merged together through addition. Using this

approach, an otherwise impossible situation wherein $n' > n$ and/or $n'' > n$, may be reduced to a problem that is expressed in a reduced basis set with n basis functions that is solvable with this system that has finite processing capabilities.

[00318] This reduced basis set is found by computing the Krylov subspace which substantially occupies the same matrix sub-space as the n most dominant eigenvectors of the final matrix that is composed of multiple optimisations for basis functions. The Krylov subspace may be found using a standard technique such as the Arnoldi iteration, wherein a dominant eigenvector is found and then removed from the remaining matrix and then this is applied in a repeated manner until n such eigenvectors are found.

[00319] IV. Mixed approach

[00320] This approach whereby basis functions are extracted from high-rank matrices may be mixed and matched with the existing specific basis function approach. For instance, two control regions or points may be designated must-have and added into the basis set input to the linear system. Then, the remaining basis functions, given that they may be specified by quadratic optimizations and thus optimization matrices of their own, can be added to a single matrix which is then deconstructed using this Krylov sub-space method to find the most dominant $n - 2$ eigenvectors to make up the remaining number of the set of basis functions input to the linear system. In this way, the two basis functions that are specified up front are guaranteed as they are imported directly into the linear system, but the remaining functions of the set are approximated. This system may be taken even further, wherein a set number of basis functions may be specified and taken from disjoint sets of basis functions which are each described by an optimization matrix of high rank and then a set number of basis sets taken. A further technique that could be used to derive reduced basis functions is to continue to take basis functions from a matrix until a power threshold is crossed. This threshold may be different for each matrix made up from conjoined basis sets and therefore yield eigenvectors of different power, where the power of each eigenvector is described by the magnitude its corresponding eigenvalue, and in this way allow eigenvectors may be produced by selecting eigenvector that are calculated from each matrix based on the product of a real priority value and the corresponding eigenvalue magnitude.

[00321] 8. LINEAR SYSTEMS SOLUTION USING TEST FUNCTIONS

[00322] A method to 'test' whether the basis function set achieves certain goals that is defined separately from the basis function set would be of value. The quantitative evaluation

of whether these goals are achieved is supplied to the optimization and this is used to generate desired properties in the acoustic field that are defined by these test functions.

[00323] The test functions and their role in the complex-valued linear system optimization will be described.

[00324] I. Linear Control Regions

[00325] A linear control region, which is defined as an operator that performs an integral with an linear acoustic quantity, such as pressure, with a complex valued weighting function, $w(\chi)$, $\chi = \{x, y, z\}$:

$$f_p = \iiint_{\Omega} \overline{w(\chi)} \cdot p(\chi) dx dy dz,$$

may be used to evaluate an acoustic field. It should be noted that the complex conjugate is used to maximize function evaluation when $w(\chi)$ is paired with similar phase distributions in the acoustic quantity. Any approximation to this integral (such as sampling) can be used or indeed this may be used to approximate any linear operator over the field of acoustic pressure. A further integral may be constructed to optimize a control region using a vector particle velocity (wherein each vector component is complex-valued due to the harmonic motion of the field), given it is also linear, wherein:

$$f_v = \iiint_{\Omega} \overline{\mathbf{w}(\chi)} \cdot \mathbf{u}(\chi) dx dy dz,$$

with $\mathbf{u}$ the particle velocity of the medium, and $\mathbf{w}$ a three-dimensional vector field in which each dimension is complex-valued, which serves as a weighting function with the same purpose as above. In this way, different patterns of acoustic quantities may be evaluated in a linear fashion, wherein a particular phase of the result will maximize the objective function when phases align with the ‘template’ created by the scalar or vector function in w .

[00326] It also should be noted that $w(\chi)$ may be largely zero (or comprised of zero vectors) to create a single evaluation of multiple seemingly disjoint regions. Any integral approximation (such as sampling) can be used as this formulation can extend to any linear operator over the field of a linear acoustic quantity.

[00327] II. Basis function set

[00328] Given a set of n basis functions as input, m test functions may be used through a linear system to find the optimal linear combination of the n basis functions to achieve the desired objective evaluation of the m test functions. These m test functions may be take the form of linear control region evaluations. Using the forms of the objective f as given above these would be:

$$\omega(\mathbf{x}) = \sum_{q=1}^N x_q \bar{f}_q,$$

or they could be evaluated through the single control point definition as:

$$\omega(\mathbf{x}) = p(\chi) = \sum_{q=1}^N x_q \overline{p_q(\chi)},$$

where the overline denotes complex conjugate, χ is three-dimensional position, q parameterises all transducer elements and x is the transducer activation which is multiplied by the conjugate of the pressure field of the transducer. The equivalent form for a control point defined by particle velocity may also be used as an evaluator, albeit with a defined normal direction:

$$\omega(\mathbf{x}) = \mathbf{u}(\mathbf{x}; \chi) \cdot \hat{\mathbf{n}}(\chi) = \sum_{q=1}^N x_q \overline{\mathbf{u}_q(\chi) \cdot \hat{\mathbf{n}}(\chi)},$$

where any transducer q which is pointing in the wrong direction may be also omitted (have its contribution set to zero when $\hat{\mathbf{d}}(\chi) \cdot \hat{\mathbf{n}}(\chi) < 0$).

[00329] In both single control point definitions, the existing approach can be viewed as a special single point integral evaluation with unit weighting function case of the control region evaluation. The objective function here is then simply the acoustic quantity with no extra side effects.

[00330] III. Finding the dominant eigenvector

[00331] By evaluating each of the m basis functions using each of the n test functions, a rectangular matrix that is complex-valued is arrived at. As before, the dominant eigenvector denotes the combination of the m basis functions that best spans the space defined by the n test functions. However, as this is not necessarily a square matrix, the singular value decomposition is more appropriate, but then the only required eigenvector is the one that is dominant. A square Hermitian matrix may be constructed whose dominant eigenvector

represents the elements of dominant combination of the system basis functions. The eigensystem is then described by:

$$\begin{bmatrix} 0 & \begin{bmatrix} \omega_1(x_1) & \cdots & \omega_1(x_n) \\ \vdots & \ddots & \vdots \\ \omega_m(x_1) & \cdots & \omega_m(x_n) \end{bmatrix} \\ \begin{bmatrix} \overline{\omega_1(x_1)} & \cdots & \overline{\omega_m(x_1)} \\ \vdots & \ddots & \vdots \\ \overline{\omega_1(x_n)} & \cdots & \overline{\omega_m(x_n)} \end{bmatrix} & 0 \end{bmatrix} \begin{bmatrix} a_1 \\ \vdots \\ a_m \\ b_1 \\ \vdots \\ b_n \end{bmatrix} = \lambda \begin{bmatrix} a_1 \\ \vdots \\ a_m \\ b_1 \\ \vdots \\ b_n \end{bmatrix}$$

where $b_1, \dots, b_n$, describe coefficients whose phase generates a covering of the parameter space of the m test functions. The dominant eigenvector therefore has components $b_1, \dots, b_n$, which describe the best covering of the space achievable with a single combination vector, and whose phase therefore optimizes this covering.

[00332] It should be noted that even in the case that the number of tests functions is equivalent to the number of basis functions, so $m = n$, the existing eigensystem approach may not necessarily be applicable. This is due to the fact that the matrix may no longer be Hermitian and therefore the direct application of the Rayleigh-Ritz method to this matrix without completing the above step may cause the method to fail to converge to the dominant eigenvector.

[00333] The dominant eigenvector should be used for phase information in the event that the phase data is not specified or is incomplete. The eigenvector may, on the end of each iteration be, modified to strip the amplitude data, averaged to zero in phase, limited in phase progression from the previous solution in a separate time-step or have one set of phase solutions for some basis functions held with existing and constant phase relationships while the complement of this set is optimized with phases between themselves and the stationary set of existing phases. In this way, this rectangular matrix system may be used to achieve an effect similar to the existing eigensystem approach which is specialized for the situation in which the functions that generate the basis set are the same as the functions used to evaluate the objective of the basis set.

[00334] IV. Linear system solution

[00335] The complex-valued linear system now has a modified role, that of finding the set of basis functions that best replicate given values of the test functions. Here the linear system that must be used takes the form of a rectangular matrix:

$$\begin{bmatrix} \omega_1(\mathbf{x}_1) & \cdots & \omega_1(\mathbf{x}_n) \\ \vdots & \ddots & \vdots \\ \omega_m(\mathbf{x}_1) & \cdots & \omega_m(\mathbf{x}_n) \end{bmatrix} \begin{bmatrix} z_1 \\ \vdots \\ z_n \end{bmatrix} = \begin{bmatrix} x_1 \\ \vdots \\ x_m \end{bmatrix}.$$

[00336] As the matrix is rectangular, it may be that the number of the test functions is greater than the number of basis functions, resulting in a complex-valued least-squares system, or the number of test functions is less than the number of basis functions, resulting in a complex-valued least norm system. In either case, a standard decomposition may be used to find the solution vector, whose interpretation is that each component of the solution vector is the coefficient to be applied to each of the basis functions which cause the acoustic field when tested to express the given values of the test functions.

[00337] These test functions may express acoustic pressure or directional particle speed at a point, or a mixture of pressures and direction particle speeds across a region in space. Any combination of these functions may then make up the m test functions, which are solved for in the final step of the linear system.

[00338] Once the coefficients for the basis functions are known, the final values to actuating the transducers may be expanded as:

$$\mathbf{x} = \mathbf{A}\mathbf{z}^H = \sum_{j=1}^n \mathbf{x}_j \bar{z}_j$$

[00339] The input signals, give an assumption of a single monochromatic carrier frequency (which may be more or less valid due to changing circumstance) may be constructed by the application of components of this $\mathbf{x}$ to each transducer.

[00340] 9. CONCLUSION

[00341] The various features of the foregoing embodiments may be selected and combined to produce numerous variations of improved haptic systems.

[00342] In the foregoing specification, specific embodiments have been described. However, one of ordinary skill in the art appreciates that various modifications and changes can be made without departing from the scope of the invention as set forth in the claims below. Accordingly, the specification and figures are to be regarded in an illustrative rather than a restrictive sense, and all such modifications are intended to be included within the scope of present teachings.

[00343] The benefits, advantages, solutions to problems, and any element(s) that may cause any benefit, advantage, or solution to occur or become more pronounced are not to be construed as a critical, required, or essential features or elements of any or all the claims. The invention is defined solely by the appended claims including any amendments made during the pendency of this application and all equivalents of those claims as issued.

[00344] Moreover, in this document, relational terms such as first and second, top and bottom, and the like may be used solely to distinguish one entity or action from another entity or action without necessarily requiring or implying any actual such relationship or order between such entities or actions. The terms “comprises,” “comprising,” “has,” “having,” “includes,” “including,” “contains,” “containing” or any other variation thereof, are intended to cover a non-exclusive inclusion, such that a process, method, article, or apparatus that comprises, has, includes, contains a list of elements does not include only those elements but may include other elements not expressly listed or inherent to such process, method, article, or apparatus. An element preceded by “comprises ...a”, “has ...a”, “includes ...a”, “contains ...a” does not, without more constraints, preclude the existence of additional identical elements in the process, method, article, or apparatus that comprises, has, includes, contains the element. The terms “a” and “an” are defined as one or more unless explicitly stated otherwise herein. The terms “substantially”, “essentially”, “approximately”, “about” or any other version thereof, are defined as being close to as understood by one of ordinary skill in the art. The term “coupled” as used herein is defined as connected, although not necessarily directly and not necessarily mechanically. A device or structure that is “configured” in a certain way is configured in at least that way, but may also be configured in ways that are not listed.

[00345] The Abstract of the Disclosure is provided to allow the reader to quickly ascertain the nature of the technical disclosure. It is submitted with the understanding that it will not be used to interpret or limit the scope or meaning of the claims. In addition, in the foregoing Detailed Description, it can be seen that various features are grouped together in various embodiments for the purpose of streamlining the disclosure. This method of disclosure is not to be interpreted as reflecting an intention that the claimed embodiments require more features than are expressly recited in each claim. Rather, as the following claims reflect, inventive subject matter lies in less than all features of a single disclosed embodiment. Thus, the following claims are hereby incorporated into the Detailed Description, with each claim standing on its own as a separately claimed subject matter.

CLAIMS

We claim:

1. A method comprising:
 - i) producing an acoustic field from a phased-transducer array having known relative positions and orientations;
 - ii) defining a plurality of control points, wherein each of the plurality of control points has a known spatial relationship relative to the phased-transducer array; and
 - iii) optimizing for a creation of desired field values in a known direction in a physical field of a quantity generated by the phased-transducer array through superposition of wave phases and amplitudes resulting from input driving coefficients given by an optimization process, wherein the quantity is an acoustic vector quantity.
2. The method as in claim 1, further comprising:
generating a further optimization matrix using a sum of eigensystems.
3. The method as in claim 1, further comprising:
generating a further optimization matrix using at least one null point in at least one field.
4. The method of claim 1, wherein at least one of the desired field values comprises velocities of particles at points in a medium hosting an acoustic wave.
5. The method of claim 1, wherein at least one of the desired field values comprises momentum in an acoustic wave.
6. The method of claim 1, further comprising:
generating further optimization using first values to calculate a plurality of basis functions.
7. The method of claim 6, wherein the plurality of basis functions are used to realize at least one control point that specifies second values in the physical field, wherein the second values differ from the first values.

8. A method comprising:
- i) producing an acoustic field from a phased-transducer array having known relative positions and orientations;
 - ii) defining a plurality of control points, wherein each of the plurality of control points has a known spatial relationship relative to the phased-transducer array; and
 - iii) optimizing for maximization of an absolute magnitude of desired field values along a known direction in a physical field of a quantity generated by the phased-transducer array through superposition of wave phases and amplitudes resulting from input driving coefficients given by the optimization process, wherein the quantity is an acoustic vector quantity.
9. The method as in claim 8, further comprising:
generating a further optimization matrix using a sum of eigensystems.
10. The method as in claim 8, further comprising:
generating a further optimization matrix using at least one null point in at least one field.
11. The method as in claim 8, wherein at least one of the desired field values comprises acoustic intensity.
12. The method as in claim 8, wherein at least one of the desired fields comprises acoustic energy flux vectors.
13. The method of claim 8, wherein at least one of the desired field values comprises velocities of particles at points in a medium hosting an acoustic wave.
14. The method of claim 8, wherein at least one of the desired field values comprises momentum in an acoustic wave.

15. The method of claim 8, further comprising:
generating further optimization using first values to calculate a plurality of basis functions.
16. The method of claim 15, wherein the plurality of basis functions are used to realize at least one control point that specifies second values in the physical field, wherein the second values differ from the first values.
17. A method comprising:
- i) producing an acoustic field from a phased-transducer array having known relative positions and orientations;
 - ii) defining a plurality of control points, wherein each of the plurality of control points has a known spatial relationship relative to the phased-transducer array; and
 - iii) optimizing for substantial maximization of a weighted integral of at least one desired region of at least one physical field generated by the phased-transducer array through superposition of wave phases and amplitudes resulting from input driving coefficients given by an optimization process.
18. The method of claim 17, further comprising:
generating a further optimization matrix using a sum of eigensystems.
19. The method of claim 17, further comprising:
generating a further optimization matrix using at least one null point in at least one field.
20. The method of claim 17, wherein at least one of the desired regions of a field comprises a weighted acoustic intensity integral.

21. The method of claim 17, wherein at least one of the desired regions of a field comprises a weighted integral of acoustic energy flux vectors.

22. The method of claim 17, wherein at least one of the desired regions of a field comprises a weighted integral of acoustic pressure.

23. The method of claim 17, wherein at least one of the desired regions of a field comprises a weighted integral of velocities of particles at each point in a medium host to an acoustic wave.

24. The method of claim 17, wherein at least one of the desired regions of a field comprises a weighted integral of momentum in an acoustic wave.

25. The method of claim 17, further comprising:
generating a further optimization using first weighted integral values to calculate a plurality of basis functions.

26. The method of claim 25, wherein the plurality of basis functions are used to realize at least one control region that specifies second weighted integral values in the physical field, wherein the second weighted integral values differ from the first weighted integral values.