

DYNAMIC CAPPING WITH VIRTUAL MICROPHONES

PRIOR APPLICATIONS

[0001] This application claims the benefit of U.S. Provisional Patent Application No. 62/914,502 filed on October 13, 2019, which is incorporated by reference in its entirety.

[0002] The prior application U.S. Application No. 15/960,113 filed on April 23, 2018 is incorporated by reference in its entirety.

[0003] The prior application U.S. Application No. 62/507,822 filed on May 18, 2017 is incorporated by reference in its entirety.

FIELD OF THE DISCLOSURE

[0004] The present disclosure relates generally to add a virtual microphone to an ultrasound phased array system.

BACKGROUND

[0005] A goal of this disclosure is to produce an estimate of the acoustic pressure from an ultrasound phased array that reasonably matches the measurement of a stationary or slow-moving microphone at a similar location.

[0006] There are methods that detail ways to calculate instantaneous pressure or intensity or other metrics in the field. Here, a series of algorithms efficiently use computational resources to calculate time-averaged metrics. These are useful for determining and regulating hot spots and higher-than desired pressure.

SUMMARY

[0007] Estimating the field strength from an ultrasonic phased array can be done via a processor by summing the contribution of each transducer to the point of interest. This contribution is already calculated when creating a converging spherical wave. This calculation can be reused to add a virtual microphone to the

system. By monitoring this microphone and moving it along with new focus points, a robust system of field estimates and regulation may be established.

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] 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, serve to further illustrate embodiments of concepts that include the claimed invention and explain various principles and advantages of those embodiments.

[0009] Figure 1 shows a flowchart of a dual-mic arrangement.

[0010] Figure 2 shows a flowchart of an N-mic, N-average arrangement.

[0011] 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.

[0012] 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

[0013] I. Virtual Microphones

[0014] The pressure output of an n -sized array of these transducers at point $\mathbf{x}$, relative to the center of the array, can be then written as,

$$P_{tot}(\mathbf{x}) = \sum_n X_n P_n(\mathbf{x} - \mathbf{y}_n), \quad (1)$$

where $\mathbf{y}_n$ is the offset of each transducer relative to the center of the array, P_n is a function/model which gives the complex pressure output at given vector relative to the transducer, and X_n represents a complex activation coefficient for each transducer. A subscript n is included in the pressure function to allow for potentially different transducers within an array. The activation coefficient can elicit a phase shift, amplitude change, or both and manipulating these coefficients controls the acoustic field. In a real system this is interpreted into an amplitude and phase to drive each transducer.

[0015] An important field from an array for haptics is that of a spherical wave front converging on a focus point. If we compose a list of pressure values from each transducer at point $\mathbf{x}$ relative to the center of the array, $A_n = P_n(\mathbf{x} - \mathbf{y}_n)$, the trivial activation solution is:

$$X_n = \frac{A_n^*}{\sum_k |A_k|^2} b ,$$

where b is the desired complex pressure at the focal point. If we declare $A_n = \mathbf{A}$ as a row-vector and $X_n = \mathbf{X}$ a column vector, another way to write equation 1 is:

$$\mathbf{A} \cdot \mathbf{X} = b.$$

[0016] Calculating the row-vector $\mathbf{A}$ is part of constructing this field solution. When the focal point is moved, the system produces a new $\mathbf{A} = \mathbf{A}'$, a new $\mathbf{X} = \mathbf{X}'$, and the system moves on with making new focal point locations. The key here is to recognize the significance of $\mathbf{A} \cdot \mathbf{X}'$, the old row-vector multiplied by the current activation vector. The output of this multiplication is an estimate of the pressure at the *old* focus location with the *new* activation coefficients. With one extra vector multiply, the system can estimate of the pressure at a specific point in space (in this case an older focus location) produced by the current field it is rendering. By storing the old $\mathbf{A}$ and performing $\mathbf{A} \cdot \mathbf{X}'$ each time the activation coefficients are updated we have a field estimate (virtual microphone or for

brevity ‘mic’) at the previous focus location with a minimum of extra computation needed.

[0017] Unlimited virtual microphones can be created – either at previous focus locations or using $A_n = P_n(\mathbf{x} - \mathbf{y}_n)$ to interrogate a new $\mathbf{x}$, unrelated to past fields. Every new microphone, however, requires a dot product vector-by-vector multiplication each time its pressure is to be updated. Therefore, it is not computationally efficient to investigate every point in the field. For a mid-air haptic application, it is most useful to be near focal points. Haptic curves only produce a sensation when they are repeated many times—this gives a guide as to where possible high-pressure points are = along points that have been visited before.

[0018] With this in mind, in one arrangement of this disclosure, we update the location of the virtual microphone as the array focuses to new locations. In this case we ‘move’ the mic to the current focus location by overwriting the stored $\mathbf{A}$ with the current calculated $\mathbf{A}'$. Any moving averages (explained below) are retained and proceed with the new location and associated row-vector. The decision of when-to-move can involve threshold pressure, an external flag, or any other signal. One example is a simple count-down of timesteps, here deemed Δn . Rather than a fixed value after every move, including a random component prevents locking into a location if the field is repeating locations with a consistent frequency. If Δn is a fixed value then loops which are an integer multiple of Δn cause the moving mic to be locked into only a few points, possibly missing higher-pressure regions of a repeated curve. To remedy this, Δn can be composed of a fixed minimum- n (n_{fixed}) and an additional random-length (n_{random}). By making n_{random} equal to or larger than n_{fixed} , the moving mic can equally sample curves of arbitrary length.

[0019] II. Moving Averages

[0020] A moving average is a calculation which smooths input. This is essential to even out random noise but is also useful to develop metrics for sinusoidal signals. In acoustics, for instance, calculating the energy contained in a

monochromatic wave is achieved by averaging the instantaneous pressure squared over the course of one period. For a complex signal composed of many frequencies, a single, repeating period is long or does not exist. Calculating instantaneous pressure-squared would over-estimate the energy in the sound field. To compensate for this, we can use moving averages to include adjacent values and develop a metric relevant to the energy in the acoustic wave.

[0021] One method to calculate a moving average is called a “box filter”. This is implemented by averaging a series of points in a signal. As new points are acquired, old points are forgotten, and a new average is calculated. Another way to look at this is with convolution. Convolution takes a kernel, or a series of weights, and multiplies those weights by the input signal, where each point in the kernel is multiplied by the input based delayed equally to its position in the kernel array, and then performs a sum as its output. The “box filter” moving average kernel, for instance, is simply a series of equal-valued variables equal to $1/n$ where n is the size of the kernel.

[0022] Another method to calculate a moving average is a recursive implementation whereby the last calculated value of the average is used, along with new data, to calculate subsequent points. One particular recursive calculation which is relatively easy to implement is an exponentially weighted rolling average. This method is described by,

$$Avg_n = (1 - \alpha) Avg_{n-1} + \alpha x, \quad (2)$$

where Avg_n is the current rolling average, Avg_{n-1} is the calculated average from the previous iteration, x is the new input, and α is a constant. The constant α determines the exponential weighting in units of the sampling rate.

[0023] A particular metric often used in acoustics is known as sound-pressure-level (SPL), expressed in decibels (re 20 microPascals). This is calculated through the square root of the exponentially weighted moving average value of pressure squared (rolling root-mean-squared (RMS) value). Time constants used for the

exact calculation vary from a few milliseconds to several seconds depending on the application and engineering specification.

[0024] This metric can be tracked using virtual microphones by keeping a rolling average of the pressure-squared value (mean-squared) and taking the root when a decibel value is needed. This value can be used to regulate (attenuate) the array output if desired. Alternatively, pressure-squared units can be tracked and used as the regulatory value. Since the square-root is a single-valued function, the resulting SPL value will be correctly bounded.

[0025] Calculating the exponentially weighted average of P-squared (squared pressure output) ($\overline{P^2}$) using discrete time-steps can be accomplished with a recursive algorithm,

$$\overline{P^2}(n) = \alpha P^2 + (1 - \alpha) \overline{P^2}(n - 1),$$

where $\overline{P^2}(n - 1)$, is the output from the previous timestep, P^2 is the square of the pressure at the current timestep, and α is a constant between zero and one. The constant α represents the exponential time-constant in timestep units. For instance, a 1-second integration constant running at 40kHz would be given by $\alpha = \frac{1}{40000}$. This would average points 1-second in the past by $1/e$ relative to the current point. Points 2-seconds in the past would be weighted by $(1/e)^2$ and so on.

[0026] The same approach may be taken with an energy metric, which is based on the simulation of acoustic intensity. The calculation of an energy metric involves going back to the initial computation of $\mathbf{A} \cdot \mathbf{X}'$. Each component of the vector particle velocity of the medium may be calculated for a single source by a multiplication of the pressure with the relevant component of the wavefront normal vector divided through by the acoustic impedance. As each element is modelled as a spherical wave source the wavefront normal vector is simply:

$$\hat{\mathbf{n}}_n = \frac{\mathbf{x} - \mathbf{y}_n}{\|\mathbf{x} - \mathbf{y}_n\|},$$

As A_n is $P_n(\mathbf{x} - \mathbf{y}_n)$, the vector particle velocity of the medium neglecting the division by the constant acoustic impedance of air, may be calculated for each element as $\hat{\mathbf{n}}_n P_n(\mathbf{x} - \mathbf{y}_n)$. If these normal directions are retained alongside $\mathbf{A}$, they may be incorporated into the row-vector as $\mathbf{A}_{n,x}$, $\mathbf{A}_{n,y}$ and $\mathbf{A}_{n,z}$. As these are also linear quantities these may be summed and reconstructed as $\mathbf{A}_{n,x} \cdot \mathbf{X}'$, $\mathbf{A}_{n,y} \cdot \mathbf{X}'$ and $\mathbf{A}_{n,z} \cdot \mathbf{X}'$. The acoustic intensity vector (again neglecting constant terms involving acoustic impedance), whose magnitude describes the total energy of the acoustic wave at a point, can then be written as:

$$\mathbf{I} = \mathbf{A} \cdot \mathbf{X}' \cdot \begin{bmatrix} \mathbf{A}_{n,x} \cdot \mathbf{X}' \\ \mathbf{A}_{n,y} \cdot \mathbf{X}' \\ \mathbf{A}_{n,z} \cdot \mathbf{X}' \end{bmatrix},$$

[0027] This provides a measurement of the total energy of the wave which may be used in the place of pressure. As the energy has units proportional to P-squared, the rolling average of $\mathbf{I}$ per equation 2 represents a true energy metric. This is in contrast to the P-squared which for SPL requires a square-root to give a root-mean-squared (RMS) value.

[0028] For situations with lower sample rate (or when the mic averages are not updated every cycle) α needs to be adjusted to represent the new value. This can be done on the fly if necessary, with α being adjusted each time that the average is updated.

[0029] III. Dual Mic Arrangement

[0030] There exists a class of curves with self-crossings (or crossings from other points/grating lobes) where the moving-mic arrangement presented above can mis-represent the highest-pressure point in the path by as much as 6dB. To remedy this, we can divide the problem into two separate virtual mics – one which always moves as above, and another which locks onto a given point until the moving mic has found a higher-pressure location. In this way, the moving mic, from here called the “seek” mic, samples the path until a hotspot has been found, at which point the locking mic, from here called the “regulation” mic, moves to that point. As long as that point is the “hottest” in the field, it will stay there and

correctly estimate the maximum pressure one could measure. Only when the path has changed (and stopped contributing to the “hotspot”) will the seek mic be able to pull the regulation mic to a new location.

[0031] The seek mic is able to detect hotspots quickly because it will be using a larger alpha for its rolling average. Functionally, this creates a situation where its rolling average is more weighted towards the recent past relative to the small alpha. A hotspot occurs when the motion of a focus point is slowed, or the array focus returns near the virtual mic after a short period. With the larger alpha, seek mic approaches the true (long-term) P-squared value more quickly than a regulation-alpha mic would and sees a spike at the hotspot. A regulation-mic move is initiated when the large-alpha P-squared exceeds the regulation-alpha P-squared of the regulation mic by some margin. The regulation mic’s P-squared is then set to the seek-mic P-squared, basically setting a new (higher) bar for further movement. If the regulation mic did indeed get placed on a hotspot then its small-alpha P-squared will continue to increase and be out of reach of the large-alpha P-squared of the seek mic until the system changes paths. If the move did not put the regulation mic on a hotspot then its P-squared will decay until it is moved again.

[0032] In this system we still need two more rolling averages. The first is a small-alpha P-squared which is not adjusted when the regulation mic is moved. This is necessary because without it we do not have an estimate which represents the pressure that would be measured by a stationary mic. The regulation-alpha P-squared which is overwritten by the large-alpha seek mic P-squared changes much more quickly and only serves as a comparison for a movement decision. By keeping a second regulation-alpha P-squared which is not overwritten, the system has an estimate of the pressure with which to dictate regulation.

[0033] The last P-squared needed is calculated using the seek mic location (the output of its pressure estimate) but is averaged using the small-alpha regulation mic time constant. This effectively gives the moving-mic P-squared average achieved using only a single moving virtual microphone. For a certain class of

curves this value represents a better estimate of the path pressure than the regulation mic P-squared.

[0034] Turning to Figure 1, shown is a diagram 1300 of the decision tree for a dual mic arrangement. A new time step begins 1302 following by calculating new array parameters 1303. These parameters are then provided to the vectors 1317 (row vector **A** 1316 and column vector **X** 1315). The row vector **A** 1316 is provided to overwrite **A_seek** with **A** and reset move counter with a fixed minimum-n (n_{fixed}) plus additional random-length (n_{random}) 1311. The column vector **X** 1315 is provided to a group of **A** vectors 1320 (**A_seek** 1321 and **A_reg** 1322).

[0035] Calculating new array parameters 1303 are also provided to estimate new microphone pressures 1304, which is provided to the group of **A** vectors 1320. **A_seek** 1321 is provided to an avg group 1323 consisting of seek avg fast 1324 and seek avg slow 1325. **A_reg** 1322 is provided to a reg group 1328 consisting of reg avg for movement 1326 and reg avg unmoving 1327. The seek avg slow 1325 and reg avg unmoving 1327 are compared and the largest value is taken for regulation 1329.

[0036] The **A_seek** 1321 is sent to overwrite **A_reg** with **A_seek** and overwrite Reg avg for movement with seek avg fast 1312 and is then sent to **A_reg** 1322.

[0037] The seek avg fast 1324 is sent to overwrite **A_reg** with **A_seek** and overwrite Reg avg for movement with seek avg fast 1312 and is then sent to the reg avg for movement 1326.

[0038] The values from seek avg fast 1324 and reg avg for movement 1326 are compared to determine which one is greater (with margin) 1313; the result is sent to overwrite Reg avg for movement with seek avg fast 1312 and then to Reg avg for movement 1326.

[0039] The values of estimate new microphone pressures 1304 are provided to update the moving averages 1305, which is provided to: 1) the avg group 1323;

and the move decision for Reg mic 1306 and then to the move decision for Seek mic 1307.

[0040] The counter is moved from the last step 1308 and then decremented 1309. If the counter is less than or equal to zero 1310, an instruction is sent to overwrite A_{seek} with A and reset move counter with a fixed minimum- n (n_{fixed}) plus additional random-length (n_{random}) 1311.

[0041] Rolling averages from the last time step 1314 are provided to the avg group 1323 and the reg group 1328.

[0042] In some implementations of the system there exists time where the seek microphone might be calculating an average but due to propagation delay, the regulation mic might not be able to move to that location for several time steps. In that case, one solution is to stop averaging the seek mic into the moving average when the regulation mic is not able to move to that location. This introduces a “blind spot” for those particular points (especially if they happen on a regular basis). Since the movement of the seek mic is adjusted with a random factor, however, even a regular ‘blind spot’ should be distributed throughout any regular curve and not represent a hole in the measurement. Excessive lack of measurement represents degraded performance for finding hot-spots and should be minimized.

[0043] IV. N-Mic Arrangement

[0044] In order to find a hotspot, the seek mic still needs to randomly land on that point. For some design parameters (peak pressure possible, regulation pressure, Δn , etc) this can take a considerable amount of time. This time can be reduced by adding more seek mics for a given focus point. Each keeps its own P-squared values and move counters. At every time step, the regulation mic would compare its P-squared to all of the seek mics and move if any of them best it. With more points being checked simultaneously, a hotspot will be found more quickly.

[0045] Ideally, every seek mic is measuring at a different physical location. Practically, this is difficult to guarantee because the physical location of the mic is

not stored and even if it were, the comparison to every other mic takes computation. One solution is to simply guarantee that no two mics are moved on the same timestep. If two move counters expire at the same time, move one mic and add some value to the move counter of the other. This can be 1 at minimum or some other (possibly random) value based upon the number of other mics in the system.

[0046] In addition to mics, the system may include any number of averages (with different alphas) per microphone. This could be used, for example, to switch to different time-constants for regulation on the fly. Some could be overwritten for a move while others are held. This yields a more flexible system with regards to possible fluctuations. For instance, a short-time alpha reacts more quickly to a hot-spot but also can randomly reach a peak in an instance where it randomly lands on a spurious crossing of another focal point. By including several averages of differing time-constant this could give the designer the possibility of ignoring faster-responding values. The extra averages could also be of different metrics, P -squared and I being examples. Move decisions could then involve multiple metrics, for instance.

[0047] Turning to Figure 2, shown are possibilities of a flowchart 1400 of an N-mic, N-average arrangement.

[0048] A new time step begins 1401 following by calculating new array parameters 1402. These parameters are then provided to the vectors 1408 (row vector A 1409 and column vector X 1407). The row vector A 1409 is provided to overwrite A_{seek} with A and reset move counter with a fixed minimum- n (n_{fixed}) plus additional random-length (n_{random}) 1421. The column vector X 1315 is provided to a group of A vectors 1412 (one or more A_{seek} and a A_{reg}).

[0049] Rolling averages from the last time step 1410 are sent to a group of seek avg 1414 and a group of reg avg 1413.

[0050] Calculating new array parameters 1402 are also provided to estimate new microphone pressures 1403, which is provided to the group of A vectors 1412. Data from the group of A vectors 1412 are sent to a group of seek avg 1414. The

data from the group of seek avg 1414 and a group of reg avg 1413 are sent to take appropriate value for regulation 1416.

[0051] The data from the group of seek avg 1414 are sent to overwrite appropriate averages if specified and overwrite A_reg with the A_seek which triggered the move 1417. Data from the group of A vectors 1412 are also sent to overwrite appropriate averages if specified and overwrite A_reg with the A_seek which triggered the move 1417. This is then sent to the group of A vectors 1412 and the group of reg avg 1413.

[0052] In addition, data from the group of seek avg 1414 and the group of reg avg 1413 are compared 1415 and the greater value with margin is sent to overwrite appropriate averages if specified and overwrite A_reg with the A_seek which triggered the move 1417.

[0053] The values of estimate new microphone pressures 1403 are provided to update the moving averages with many averages per mic 1404, which is provided to: 1) the group of seek avg 1414; and 2) the move decision for Reg mic 1405 and then to the move decision for each mic 1406.

[0054] The counter is moved for each seek from the last step 1418 and then decremented 1419. If the counter is less than or equal to zero 1420, an instruction is sent to overwrite A_seek with A and reset move counter with a fixed minimum-n (n_fixed) plus additional random-length (n_random) 1421. In addition, if more one mic is ≤ 0 , add a specified value to all counters but one 1422.

[0055] V. Multiplexing

[0056] Phased arrays are capable of multiplexing spherical waves to create multiple simultaneous focus points. For this type of arrangement, each focus point could have one or more of its own, independent virtual microphones. In this case, the row-vector for each point would be considered separately from the perspective of mic-storage.

[0057] For regulation, one option is to use the output of the highest rolling pressure average to be as a global regulating maximum. Alternatively, each focus

point could be regulated separately. Yet another option is to have every focus point have its own seek mic but only use one regulation mic. In this case the regulation mic would move to the location of the highest seek mic. In this case, focus points would all be regulated to the global maximum, as the regulation mic has no knowledge of which focus points are contributing to its average.

[0058] VI. Control

[0059] One use of the virtual microphone measurement is to attenuate the output of the array to satisfy a user-specified maximum SPL level. In one arrangement, the array can simply check against one of the microphones (such as the regulatory microphone) and adjust the output so that the next cycle is attenuated enough to correct the rolling average. This would be done by solving (2) for x for the desired Avg_n and then using that value for output. This, unfortunately, can lead to sharp changes in the pressure field and can result in undesired audible sound.

[0060] One smoother method to feed the measured pressure back into the output is to set the pressure targets based upon the pressure goal divided by the highest virtual microphone pressure and let the value of 1 represent the pressure goal. In this way, when the measured pressure is far below the pressure goal, the output is free to exceed it. When the pressure is equal to the pressure goal, it sustains the output. When it is too high, the output is attenuated. This regulation approach yields smoother results as the virtual microphone pressure is naturally smoothed by its own rolling average.

[0061] Yet another method is to use a standard proportion-integral-differential (PID) controller. PID is a standard real-time control scheme used across all of control systems for decades. It is simple, flexible, and computationally efficient. In this context, we can use a simple PD (no integral term) controller to adjust the capping value based upon the P-squared value returned by the virtual microphone schemes presented above. We do not need the integral term as the rolling average accomplishes an integral-like effect without the downside of wind-up. Coefficients in the PID can be adjusted to balance speed of response with overshoot and audibility. A properly implemented PID controller approaches the

limiting value at a controlled rate and does not ring excessively, thereby limiting unwanted audio.

[0062] In practice, the PID controller should use whichever is higher: the regulation mic P-squared (which is not being overwritten) or the seek mic P-squared which is using the small-alpha time constant.

[0063] VII. Additional Disclosure

[0064] 1. Coupling a moving-average calculation with field estimates.

[0065] 2. Using 2 virtual microphones in this arrangement with this relationship may represent the most computationally efficient manner to find hot spots in a quickly changing acoustic field.

[0066] 3. An ultrasonic phased array composed of

A plurality of transducers with known relative positions and orientations;

A complex activation coefficient representing the magnitude and phase of the driving signal to a transducer;

A point of interest;

Computing the transducer complex field at the point of interest;

Multiplying the transducer complex field by the activation coefficient to return an estimate of the field at that point;

Incorporating this field estimate into a rolling average;

[0067] Dependent claims:

1. Using a newly generated activation coefficient and multiplying by the older transducer complex field for a new estimate.

2. Using this new estimate for the rolling average.

[0068] VIII. Conclusion

[0069] While the foregoing descriptions disclose specific values, any other specific values may be used to achieve similar results. Further, the various features of the foregoing embodiments may be selected and combined to produce numerous variations of improved haptic systems.

[0070] 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.

[0071] 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.

[0072] 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, 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. An ultrasonic phased array comprising:

a plurality of transducers with known relative positions and orientations;

at least one complex activation coefficient representing a magnitude and phase of a driving signal to at least one of the plurality of transducers;

a point of interest;

a processor for:

1) computing a transducer complex field from at least one of the plurality of transducers at the point of interest;

2) multiplying the transducer complex field by an activation coefficient to return a field estimate at the point of interest;

3) incorporating the field estimate into a weighted average; and

4) moving the point of interest at a distance composed of a fixed minimum length and a random length;

wherein if the weighted average exceeds a pressure goal, the at least one complex activation coefficient has at least one of its real and imaginary components changed.

2. The ultrasonic phased array as in claim 1, wherein the field estimate is calculated using a squared magnitude.

3. The ultrasonic phased array as in claim 1, wherein the field estimate is calculated using a magnitude.

4. The ultrasonic phased array as in claim 1, wherein the field estimate is acted upon by a function before incorporating function output into a weighted average.

5. The ultrasonic phased array as in claim 1, wherein a plurality of versions of the transducer complex field are stored in the processor so that at least two of the

plurality of versions of the transducer complex field are used with the at least one complex activation coefficient to simultaneously estimate a plurality of points of interest.

6. The ultrasonic phased array as in claim 5, wherein a plurality of points of interest comprise a plurality of virtual microphone locations.

7. The ultrasonic phased array as in claim 6, wherein each of the plurality of virtual microphone captures its own weighted average.

8. The ultrasonic phased array as in claim 7, wherein at least one of the weighted average is calculated using a different function from others of the weighted averages.

9. The ultrasonic phased array as in claim 6, wherein at least one of the plurality of points of interest is changed when at least one weighted average exceeds a goal.

10. The ultrasonic phased array as in claim 1, wherein the transducer complex field comprises a previously calculated version of the transducer complex field.

11. The ultrasonic phased array as in claim 10, wherein the previously calculated version of the transducer complex field to be used is chosen based upon output of a weighted average from at least one virtual microphone.

12. The ultrasonic phased array as in claim 1, wherein the weighted average comprises an average of prior field estimates.

13. The ultrasonic phased array as in claim 1, wherein the weight average comprises a pressure goal divided by a highest virtual microphone pressure.
14. The ultrasonic phased array as in claim 1, wherein the activation coefficient is changed using a proportional-differential controller.
15. The ultrasonic phased array as in claim 14, wherein the proportion-differential controller adjusts the at least one complex activation coefficient value based on a squared pressure output value returned by at least one of a virtual microphone and the weighted average.
16. The ultrasonic phased array as in claim 14, wherein coefficients of the proportion-differential controller are adjusted to a controlled rate that limits unwanted audio.
17. The ultrasonic phased array as in claim 14, wherein the proportion-differential controller uses the greater of: 1) a regulation virtual microphone squared pressure output; or 2) a seek virtual microphone squared pressure output.
18. The ultrasound phased array as in claim 1, wherein the random length is not less than the fixed minimum length.