

Numerical optimization for loudspeaker and microphone arrays

C. Whitt, C. Moloney

Electrical and Computer Engineering
Memorial University of Newfoundland
St. John's, NL, Canada, A1B 3X5
{whitt,moloney}@engr.mun.ca

Abstract

Arrays of loudspeakers are often used for sound reinforcement in large concert halls, auditoria, houses of worship and outdoor events. Arrays of microphones have long been used in applications such as sonar and tomography, and are becoming more important for things like teleconferencing. Global optimization techniques such as Stochastic Region Contraction can be combined with previously developed software that predicts the performance of arbitrary arrays to find optimum arrays. Results of this study compare well with Berger and Silverman's earlier study on linear microphone arrays. Optimization of several other array configurations are presented.

1 Introduction

The use of arrays is well-established as a technique for improving the directional performance of transducers. Applications have traditionally included sonar, radar and radio transmission. Arrays also hold potential for acoustic applications such as teleconferencing systems and sound reinforcement.

In sound reinforcement applications there is a need for better directional performance from transducers, especially at low frequencies where horns and waveguides are impractical. Current loudspeaker array designs are mostly variations on a few well-known theoretical array types.

Ever increasing computing and signal processing power is making it practical to implement nearly arbitrary array configurations. Using DSP to tailor the signal sent to each individual transducer in an array is creating the possibility of new array designs that outperform current designs in specific situations. The challenge is to find the appropriate combination of position, time delay and gain for each transducer in an array to meet the individual performance goals for each system.

The first step to building the tools to optimize arrays is having a way to simulate array performance. With a general algorithm to simulate arrays it is possible to apply global optimization algorithms to find optimal values for array parameters such as transducer position, gain, and delay.

Berger and Silverman applied a technique called Stochastic Region Contraction (SRC) to microphone arrays to determine optimum configurations for a teleconferencing application [1]. They based their technique on the region contracting method first mentioned by Tang and Zheng [3]. Subsequent research by Alvarado[4] and others provided additional validation for the application of the SRC algorithm to acoustic array optimizations. The principles of microphone and loudspeaker arrays are the same, and in this work the SRC algorithm is applied to loudspeaker array design.

2 Acoustic Array Simulation

The total output of an array at any point can be calculated using superposition, using a formula correcting the one given in [2]:

$$R(\omega, c) = \sum_{i=1}^n e^{\frac{j\omega d_i}{c}}$$

where

$$d_i = [(x_r - x_i)^2 + (y_r - y_i)^2 + (z_r - z_i)^2]^{1/2} + \delta_i \cdot c;$$

where ω is the frequency of interest, c is the speed of sound in air, n is the number of array transducers, $x_r, y_r, z_r, x_i, y_i, z_i$, are the x, y , and z coordinates of the current observation point and the i th array element, respectively. The delay applied to the i th array element is δ_i .

This function is evaluated for as many observation points around the array as required to develop an adequate picture of the array's performance. The observation points are usually oriented in a circle about the array, to allow easy polar plotting of the array response. Each observation point can represent a sound source or a listener. If an array is a loudspeaker array, the function evaluated at the observation point represents the sound heard at that observation point from the array. If the array is a microphone array, the function evaluated at the observation point represents the electrical output of the array due to a source located at the observation point.

This simple formulation is very general and has the advantage of avoiding small-angle approximations. The distance from each source to each array element is separately computed, instead of assuming parallel wavefronts and calculating an offset based on the angle of incidence and the separation of array elements. This is important because it cannot be assumed that the observation point will be in the far field for very low frequencies with long wavelengths. At very low acoustic frequencies wavelengths are significant even when compared to the dimensions of large concert halls. This generality makes the equation valid in both the near and far field of an array's response.

3 Stochastic Region Contraction

Global optimization is a difficult problem without specific information about the character of the function to be optimized. Problems may often have large numbers of variables and many, many local minima. Stochastic techniques such as simulated annealing are very general but often very expensive in terms of number of function evaluations. Stochastic Region Contraction is intended to solve a restricted set of global optimization problems more efficiently than techniques like simulated annealing.

To successfully apply the SRC method, the cost function (or merit function) should meet the following conditions [1], [4]:

1. the function has a small number of large valleys, with perhaps a large number of small valleys superimposed on them;
2. the function has a strong global minimum;
3. the number of independent variables is relatively small (less than 100);
4. any variables which are quantized have a relatively large number of distinct possible values;
5. the desired uncertainty for a variable is small relative to the search range of that variable.

The SRC method operates by gradually reducing the search range of each independent variable. At each iteration candidate solutions are randomly chosen from the solution space. The cost function is evaluated for each candidate solution, and only solutions which are better than the mean of the previous iteration are kept.

Once a sufficient number of new candidates have been found that are potentially better solutions, the size of the solution space is updated. To update the solution space the best solutions are selected from the current set of candidate solutions. The number of solutions selected at this point is an internal parameter of the SRC algorithm.

The range of each independent variable is updated to only include the best solutions found so far, plus a small marginal zone. This usually results in a contraction of the solution space. It is possible for the region to occasionally expand if a good candidate solution is subsequently found in the marginal region. In no case is the solution space allowed to expand beyond the initial bounds.

The mean fitness of the best candidates is stored for use in the next iteration. Any existing candidate solutions that are better than the mean are automatically kept to the next iteration. This process is repeated until the stopping condition is met. The stopping condition can be either a specified value for the cost function, a specific volume of the solution space, or a fixed number of iterations. Formal presentations of the SRC algorithm can be found in [1] and [4].

4 Comparison with Previous SRC Results

One exercise used to become familiar with the operation of the SRC algorithm, as well as to verify the implementation was to attempt to reproduce selected results in Berger and Silverman [1]. The Berger and Silverman problem consisted of a line array of microphones, and one target (desired) source, along with a line of undesired sources parallel to the array. The target source was in the center of the unwanted sources.

Berger and Silverman called their cost function the extended power spectral distribution (PSDX). The power spectral distribution (PSD) measures the power output from the array over a range of frequencies from a given sound source. The PSDX is a min-max formulation that seeks to minimize the maximum noise power detected from the unwanted spatial region relative to the power output from the target source.

Berger and Silverman implemented a closed form expression for the PSDX. In our replication of their experiment the PSDX was implemented using our array simulator. The array output due to a source in the unwanted region was calculated at many discrete frequencies and averaged to approximate the PSD for that source. This was repeated for many sources throughout the unwanted region. From this a ratio was formed between the nominal output of the array from the target source and the maximum output from any unwanted source.

The optimal array spacings found using the simulated PSDX are compared to the results published by Berger and Silverman in Table 1. These optimizations are for 5, 7, and 9-element arrays that are oriented to have one element directly in front of the target source, and the remaining elements symmetrically arranged on either side of the center element. Because of the symmetry the number of optimization variables is 2,3, and 4, respectively.

	Berger Spacings(m)	Berger PSDX(dB)	Derived Spacings(m)	Derived PSDX(db)
5	0.076, 0.119	-8.15	0.076, 0.119	-8.12 dB
7	0.150, 0.102, 0.058	-10.46	0.149, 0.103, 0.058	-10.45 dB
9	0.171, 0.144, 0.063, 0.059	-12.19	0.175, 0.147, 0.064, 0.059	-12.21 dB

Table 1: PSDX values derived using the array simulator software compared with published values.

Berger and Silverman did not report the spatial resolution they used to sample the unwanted region. A spatial resolution of $5cm$ was found by trial and error to most closely approximate the Berger and Silverman results. Berger and Silverman’s closed form expression included all frequencies between 500 Hz and 6000 Hz. Experiments showed that a frequency resolution of 10 Hz is adequate to approximate the closed form PSDX.

After these PSDX parameters were determined there was still an unexplained discrepancy between the simulated PSDX and the Berger and Silverman published results. Without more detail on the specific implementation used by Berger and Silverman it is difficult to determine the cause of the remaining discrepancy. This discrepancy is most notable in arrays with large numbers of microphones (i.e. greater than eight microphones).

5 Loudspeaker Array Optimization

The next phase of testing the applicability of the SRC optimization technique to loudspeaker array design was to choose a known array and attempt to find the array parameters by optimizing to match the known array performance.

The array chosen to use as an optimization target was a three-element uniform endfire array. An endfire array has all its elements in a line, and time delays are used to direct the beam parallel to the array axis.

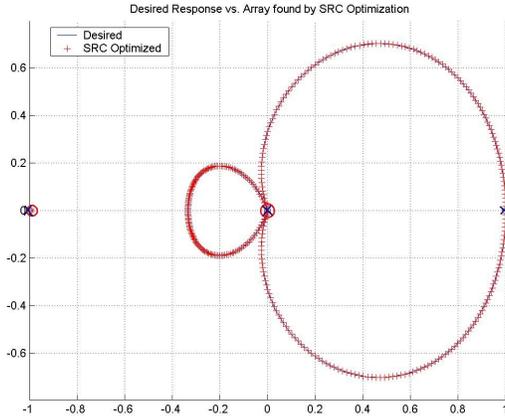


Figure 1: The theoretical performance of a three element uniform endfire array compared with the SRC-derived array optimized to match the theoretical performance. The locations of the theoretical and SRC-derived array elements are also shown superimposed.

The delay between each element is related to the design frequency of the array. The element at one end of the array will receive the signal first, and then each element along the array in order will receive the signal. More detail on various types of arrays can be found in reference [2].

The response of a three-element uniform endfire line array [5] with element spacing equal to $1/4$ the wavelength of the design frequency has a characteristic hypercardioid-like pattern. The SRC optimization algorithm was set up to optimize the position and time delay values of a three-element array. Two relative positions and delay values completely describe a three-element array, so one position and one delay value are fixed to match those of the ideal array. The ideal values for position and delay are shown in Table 2 and the response of the ideal and optimized arrays is shown in Figure 1. This optimization was done using a min-max cost function, although using a least squares cost function produces a similar result. Clearly, the optimization algorithm correctly converges on the correct solution in simple cases such as this.

Ideal Position(m)	Ideal Delay(s)	SRC Position(m)	SRC Delay(s)
$(-1,0)$	0	$(-0.989, 0.001)$	0
$(0,0)$	0.0029	$(0,0)$	0.0031
$(1,0)$	0.0059	$(1.018,0.004)$	0.0059

Table 2: Ideal and SRC-derived positions and delay for a 3-element uniform endfire array. Values in italics were fixed.

Ultimately, this optimization technique will hopefully be useful to design loudspeaker arrays with performance requirements that are not easily met by using the well-known analytically-derived array configurations. To demonstrate this concept, Figure 2 shows an arbitrary desired array response, and the corresponding performance of the array found by the SRC optimization. The resulting array does not appear to be a trivial modification of a well-known array configuration.

The quality of the optimization results is strongly influenced by the design of the cost function. For these simple examples, both min-max and least squares approaches gave reasonable results. It is not clear from these examples if a min-max or least squares approach is always preferable.

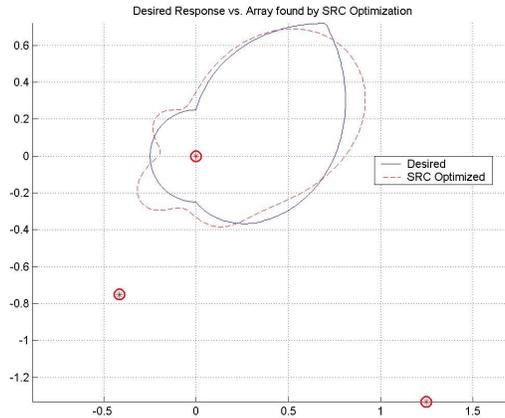


Figure 2: An arbitrary desired response and the optimum three-element array found using the SRC with a minimum mean squared error cost function.

6 Conclusions and Future Work

The SRC method of global optimization holds potential as a method for designing arrays of acoustic transducers. The earlier work of Berger and Silverman on microphone arrays has been partially verified and translated to loudspeaker arrays. It has also been shown that the SRC technique can find the parameters for well-known types of arrays when the appropriate target performance is used to setup the optimization software.

Two key challenges are involved in applying the SRC algorithm to the design of acoustic arrays: determining an appropriate cost function, and ensuring that the cost function associated with a particular problem meets the criteria in Section 3.

To guarantee an optimal solution it is particularly important that the problem has a sufficiently strong global optimum. However, in practical applications, it may be good enough to simply find a solution that is close to the global optimum, even if the global optimum cannot be known with certainty. Designing cost functions for further problems is an area for more study. It is possible that different cost functions will be better suited to different types of problems and different acoustical situations.

Continuing work with the SRC algorithm and array simulation software will focus on designing array configurations that out-perform analytically-derived arrays.

References

- [1] Mordechai F. Berger and Harvey F. Silverman, “Microphone Array Optimization by Stochastic Region Contraction”, *IEEE Trans. Signal Proc.*, vol. 39, no. 11, 2377-2386, November 1991.
- [2] Christopher J. Whitt and Cecilia Moloney, “Simulation of Loudspeaker Arrays”, *Proceedings of IEEE NECEC 2001*, St. John’s, November 2001.
- [3] J. F. Tang and Q. Zheng, “Automatic design of optical thin-film systems - merit function and numerical optimization method”, *J. Opt. Soc. Am.*, vol. 72, no. 11, 1522-1528, November 1982.
- [4] Victor Manuel Alvarado, “Talker Localization and Optimal Placement of Microphones for a Linear Microphone Array Using Stochastic Region Contraction”, *Ph.D Thesis, Brown University*, May 1990.
- [5] Jefferson A. Harrell, “Constant-Beamwidth One-Octave Bandwidth End-Fire Array of Loudspeakers”, *J. Audio Eng. Soc.*, vol. 43, no. 7/8, p581-591, July/August 1995.