EUROREGIO BNAM2022
Joint Acoustics Conference
9th-11th of May, Aalborg, Denmark

# A Two Stage Embedded Genetic Algorithm to Optimise Ceiling Reflections

**John O'Keefe**

O'Keefe Acoustics, Toronto, Canada.

john@okeefeacoustics.com

**Abstract**
Recent studies have explored computer aided techniques to optimise individual NURB-based acoustical reflectors. Reconfiguring the shape of an individual reflector with the freedom that only NURB geometry allows. This study, conversely, limits itself to primitive geometries such as triangles, polygons, circles and extruded versions of each, i.e. partial spheres, etc. These primitive shapes are collected in an array as they might be in a theatre or concert hall ceiling. An embedded, two stage Genetic Algorithm (GA) is then employed to optimise their size and orientation to encourage uniform laterally reflected sound. Each of the reflectors in the array is set at a given anchor point, starting out parallel to the receiving surface. To initiate the first GA, the given reflector is oriented to direct sound towards a Target Point on the receiving surface. The first GA then optimises that (individual) reflector to cast as many reflections as possible across the receiving surface, as laterally as possible. This is repeated for each reflector in the ceiling array. Several similar arrays are then created and these form the population for the second GA. An early experiment with this 2-stage GA suggests that it might provide higher and more uniform Lateral Fractions than a human design.

**Keywords:** Room Acoustics, Genetic Algorithms, Multi-Objective Optimisation.

## 1 Introduction

The past decade has seen increasing use of Non-uniform Rational B-Spline (NURB) geometry for acoustic reflectors [1,2]. The author has developed methods to optimise NURB-based reflector geometries [3,4] using the Non-dominated Sorting Genetic Algorithm (NSGA-II) [5]. Despite the increased use of NURB-based geometries in architectural and acoustical design, however, many performing arts projects can still not afford the luxury of NURBS. Flat triangular reflectors are easier to build and install compared to the more exotic NURB geometries. This paper will try to address the practicalities of that reality.

## 2 Reflection Calculations

### 2.1 Ray Tracing

Reflections are calculated using a modified version of a simple ray tracing routine [6]. The choice of such a simple version of a ray tracer was a deliberate effort to save time on computer optimisation runs that can last several days. The routine assumes, of course, specular reflection and that the effects of diffusion and diffraction are ignored. This was thought a safe approximation as the object of this exercise is only to optimise

the arrival of first order reflections. Diffusion and diffraction are phenomena that occur after the arrival of a first order reflection. A more complete study, using a model calibrated to an existing room and employing both diffusion and diffraction, is currently being undertaken [7].

Reflections in the computer model used in the current study can be cast from either the source location or from the reflector itself. In this experiment, the reflector size was kept constant. Only the angles of inclination for the various reflectors were modified. Thus, it was computationally more efficient to cast reflections from the reflector(s), as opposed to the traditional method of ray bundles emanating from the source location. In the reflector-based method, a number of reflection points are created on the reflector, with a density typically in the range of 10. For each one of these points, an incident line is created between the source and the reflection point. A reflection line is then cast from the reflection point, based on the angle of the incident line.

For each reflection that successfully intersects the receiver surface, its position, time delay, relative amplitude and lateralness is recorded. Lateralness is calculated using a method developed by Protheroe and Day [8] that estimates the early Lateral Fraction (LF) based on a single reflection. Ref. [8] found a reasonable correlation between free-field LFs calculated with this method and measurements performed according to ISO 3382 [9]. We refer to this parameter as the Single Lateral Fraction or sLF.

$$\mathrm{sLF} = \frac{\left(p_{reflection} Sin\phi\right)^2}{p_{direct}^2 + p_{reflection}^2} \tag{1}$$

where: $p_{direct}$ is the direct sound pressure
$p_{reflection}$ is the sound pressure from a single reflection
$\phi$ is the angle between $p_{direct}$ and $p_{reflection}$

It should be noted that for extreme angles of incidence ($\phi$), the sLF will predict Lateral Fractions of 0.5 or higher, which are surely fallacious. In more reasonable ranges, say from 0.0 to 0.4, the sLF has, however, proved to be reliable.

### 2.2 Initial Genome.

In this study, we shall be using the language of Genetic Algorithms. An array of reflectors such as that shown in Figure 1 will be referred to as the Genome. The individual reflectors that make up the Genome are the Genes. The algorithm is executed in three stages, as described below.

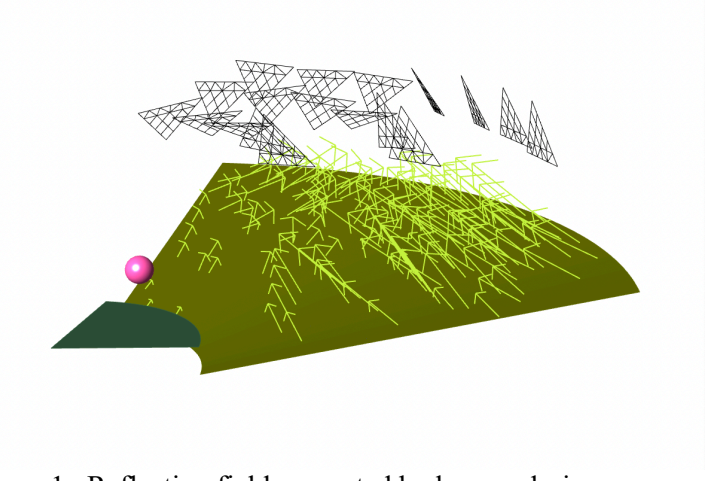

Figure 1. Reflection field generated by human design.

In the first stage of the routine, an array of Anchor Points is established in the vicinity of a receiving surface. In the example we're using here (see Figure 1) the array will form a ceiling above the seating area, but it could be any surface in the room directing reflections to any number of receiving surfaces. A reflector is built at each of the Anchor Points. The options for the six so-called "Primitive Component" shapes that the reflectors might take are shown in Figure 2. The centroid of each reflector (or Gene) is attached to its respective Anchor Point. The size of the 2-Dimensional reflectors is controlled by their edge lengths or, in the case of the circle by its radius. The 3-Dimensional shapes are simply extrusions of the polygons and the circles. Their controlling parameters are the depth of the extrusion or, in the case of the truncated pyramid (middle of bottom row, Figure 2) the depth of the extrusion and the height of the truncation. In all cases, the reflectors can be rotated about their centroids in angles of azimuth and elevation. All of these parameters: edge lengths, extrusion depth, truncation height and angles of inclination can be adjusted, or perturbed, during the Genetic Algorithm optimisation.

Once the array of reflector genes has been created at the Anchor Points, each reflector's centroid point is projected onto the receiver surface. So, in the example we're using – a ceiling of reflectors above a seating area – there will be $N_{gene}$ points projected onto the receiving surface for a ceiling with $N_{gene}$ reflector genes. These are referred to as the Target Points. Each reflector is adjusted to direct a reflection from the (single) source location, to the reflector's centroid and then to one of the (multiple) Target Points. This results in a set of $N^2_{gene}$ Reflector/Target Point combinations and the sLF value is calculated for each one. The set of $N^2_{gene}$ combinations is then sorted according to the sLF values and the combination with the highest sLF becomes the initiating genome.

### 2.3 Genome Population – 1st Stage

A population of $N_{genome}$ reflector genomes is created by perturbing the individual reflector genes created in the initial genome, as described in Section 2.2. This is the first stage of the embedded Genetic Algorithm. Each of the individual reflector genes is optimised on its own and without regard to what the other reflector genes, that will eventually make up the genome, might do. In this and the following Stage 2, optimisations are performed using the Non-dominated Sorting Genetic Algorithm (NSGA-II) [5] and the SBX cross-over [10]. The number of evolutionary optimisation generations will be referred to as $NumGen_1$ and $NumGen_2$ for Stages 1 and 2, respectively.

Although both reflector gene size and angle of inclination can be perturbed during the optimisations, only the latter was performed in this study. Two methods have been employed to change the inclination of the reflector: Target Point Perturbation and Reflector Normal Perturbation.

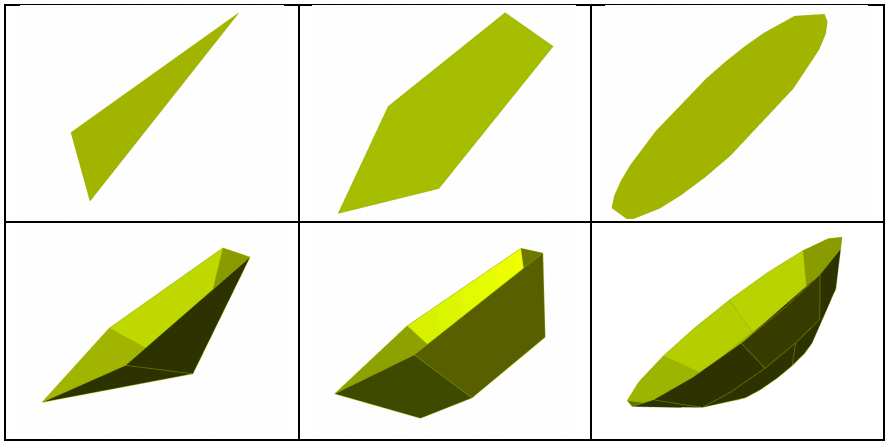

Figure 2. Primitive Component shapes available in the routine. In the top row are the 2-D options of triangles, polygons and circles. In the bottom row are 3-D extrusions of the polygons and the circles.

### 2.3.1 Target Point Perturbation

Each reflector gene has an original Target Point associated with it. It comes from the initial genome, described in Section 2.2. This point will exist on the receiver surface at the parameterised coordinates (U,V). The U,V coordinates of the Target Point are normalised with respect to the U-V span of the receiving surface. A random number generator then perturbs the normalised U,V parameters, typically over a range of U ± 0.2 and V ± 0.2. The receiver surface is then evaluated at the new U,V parameters to generate the new Target Point coordinates. A new normal for the reflector is then calculated based on the source location, the reflector's centroid and the new perturbed Target Point. The inclination of the reflector is then adjusted according to the new normal.

### 2.3.2 Reflector Normal Perturbation

Each reflector gene also has an original Reflector Normal derived, again, from the initial genome. The reflector's normal vector is normalized and, consequently, its X,Y and Z coordinates range from 0 to 1.0. A random number, typically ranging between ± 0.2, is then added to each of the vector's XYZ coordinates to create the new perturbed Reflector Normal. The inclination of the reflector is then adjusted according to the new normal.

## 2.4 Genome Population – 2nd Stage

After an individual reflector gene has been optimised through $NumGen_1$ generations (as described in Section 2.3) it is inserted into the (reflector array) genome. Then the process is repeated until we reach a population of $N_{genome}$ reflector arrays. This population is then optimised using the NSGA-II algorithm for $NumGen_2$ generations.

To summarise, and perhaps clarify, each reflector array (or Genome) will contain $N_{gene}$ reflectors (or Genes). The first stage of the routine optimises each individual reflector for $NumGen_1$ generations then contributes that reflector to the overall reflector array. A population of $N_{genome}$ reflector arrays is created this way and then the second stage of the NSGA algorithm optimises that population of arrays for $NumGen_2$ generations.

# 3 Fitness Function

Perhaps the most important component of a computer aided optimisation exercise is the Objective Function or, as it called in evolutionary genetic algorithms, the Fitness Function. A Fitness Function for the optimisation of acoustical reflectors has been developed and is described in [3]. A further explanation and extension of this proposed Fitness Function will be presented in [11]. A brief summary is presented here.

The perception of sound in a performing arts venue is a multi-dimensional experience. Any optimisation of the reflected sound in a room must provide the appropriate Reverberance, Clarity, Spatial Impression, Loudness, Warmth and Intimacy, to name a few. Evolutionary algorithms, such as NSGA-II, however, work best with only two or perhaps three optimisation dimensions. There are now some so-called "Many-Objective Evolutionary Algorithms" (MOEA) that have been developed, e.g. NSGA-III [12] but this discussion limits itself to only two objectives. These two objectives, however, address multiple psycho-acoustic dimensions.

## 3.1 dMean Fitness

The first objective is referred to as the dMean Fitness and its goal is to distribute the reflections as evenly as possible on the receiver surface. To calculate the dMean Fitness, a grid of points is created on top of the receiving surface, the nodes of which are separated by a distance (dMean), derived from the density of received reflections on any given iteration of the optimisation.

$$d_{mean} = \sqrt{\frac{Area\ of\ Receiving\ Surfaces}{\#\ of\ Receiving\ Points}} \quad (2)$$

The dMean Fitness Function then minimises the distance between any given reflection intersection point on the receiver surface and its Nearest Neighbour (NN) grid point.

The goal of this part of the Fitness Function should not be confused with a desire to produce a more "diffuse" sound. Rather, the intention is to prevent the optimisation from concentrating the reflections in one corner or along one edge of the receiving surface.

### 3.2 sLF Fitness

The second part of the Fitness Function encourages the reflections to arrive from the side before 80 ms. Again, this should not be confused with a desire only to improve the Lateral Fraction and, hence, Spatial Impression. Reflections arriving before 80 ms will also improve Clarity (C80) and early Strength (G80). The sLF Fitness function is designed to minimise the difference between a target Lateral Fraction ($sLF_{goal}$) and the sLF calculated for each reflection intersection of the receiver surface. sLF values are calculated according to Equation (1). For this study, $sLF_{goal}$ was set at 0.50. This apparently unrealistic number was chosen to avoid sLF Fitness levels too close to zero. In effect, it increases the dynamic range of the sLF Fitness Function as the routine converges towards some of the best solutions.

## 4 Results – Man vs. Machine

In one of the experiments to test the 2-stage optimisation algorithm described here, a competition was run between a human design of a ceiling reflector array (the author's) and the computer's optimisation efforts. The reflection field resulting from the human design is shown in Figure 1. The yellow arrows indicate the reflections that have intersected the receiver surface. The base of the arrow indicates where the reflection arrived, the direction indicates where it came from and the length of arrow indicates how lateral it is with respect to the source location (i.e. its sLF value).

The comparison between the human and the computer is seen in Figure 3. This is a Pareto Analysis plot of the two part Fitness Function: dMean Fitness vs. sLF Fitness. In the Fitness Function regime that has been adopted here, the goal of the Pareto Analysis is to minimise both the dMean and sLF Fitnesses, which means the better solutions are the ones closer to the origin.

The star at position (0.27, 0.28) indicates the human's performance. The sLF values for the human's design suggest a space averaged Lateral Fraction (LF) of 0.31. The square on top of the star indicates the performance of the computer's initial genome, i.e. before the optimisation has started (see Section 2.2). It appears to have a dMean Fitness and sLF Fitness similar to the human's. But the space averaged sLF value (which is not the same as sLF Fitness) is suggesting a LF of 0.24. The Just Noticeable Difference (JND) for the Lateral Fraction is 0.05 [13]. This suggests that the human's performance is better than the computer's first attempt by slightly more than a single JND. However, after optimisation ($N_{genome}$ = 25, $NumGen_1$ = 25, $NumGen_2$ = 25) the 2-stage genetic algorithm has matched the human's space averaged sLF of 0.31 but, perhaps more importantly, it has provided a better distribution of the sLF values. That is to say, the computer's solutions are closer to the origin on the (horizontal) sLF axis which means that more of the received reflections have values closer to $sLF_{goal}$.

To explain, the sLF Fitness axis quantifies the uniformity of the sLF values on the receiver surface, with respect to the target value of $sLF_{goal}$. At its maximum, we might have sLF values very close to $sLF_{goal}$ in only one corner of the receiver surface and very low values everywhere else. At its minimum, i.e. as the values approach the origin, we will have a more uniform distribution of sLF values, i.e. more of them will be closer to $sLF_{goal}$.

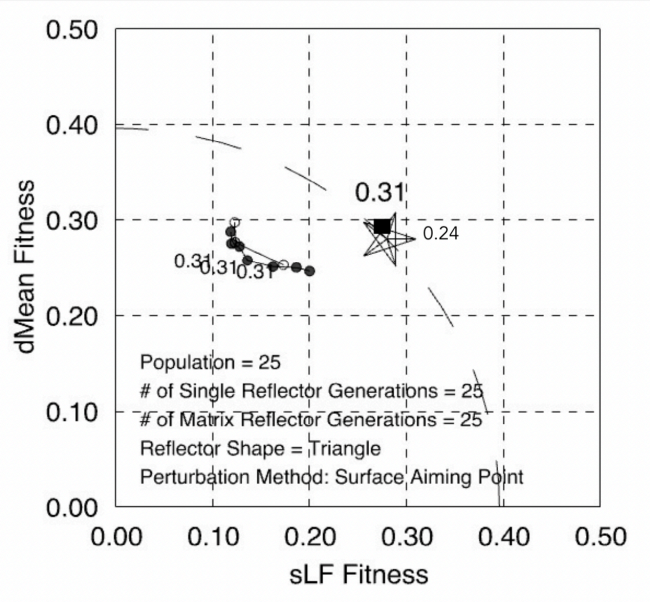


Figure 3. Pareto analysis of the "human" reflector design, compared to the computer's optimisation.

On this run, the computer has only delivered a slightly better spatial distribution of the reflections on the receiving surface. That is, compared to the human's design (the star in Figure 3), some solutions are slightly closer to origin on the (vertical) dMean Fitness axis but others are not. This is an atypical comparison between the human and the computer. In most cases, the human's distribution is worse than what is shown in Figure 3. One might say that the comparison was biased in the human's favour. In either case, the results suggest that the computer can improve on the human's performance, even when the human is at his or her best.

## 5 Conclusion

A two-stage embedded genetic algorithm to optimise acoustic reflector arrays has been presented. The first stage of the algorithm optimises the individual reflectors before they are allowed into a composite reflector array. The second stage then optimises the composite array itself. The routine limits itself to traditional (i.e. non-NURB) geometries such as triangles, pyramids and partial spheres. In this, the early stages of the algorithm's development, a trial run has shown that it can produce acoustical solutions equal to or slightly better than the human design.

## References

[1] Peter Exton, Harold Marshall, The Room Acoustic Design of the Guangzhou Opera House, *Proc. Inst. of Acoustics,* (33) 2:117-124, 2011.

[2] Thomas Scelo, J. Valentine, Harold Marshall, Chris Day, Implementing the Acoustical Concept tor The Philharmonie De Paris, Grande Salle, *Proc. of Inst. of Acoustics*, 37(3), Pt. 3: 118-127, 2015.

[3] John O'Keefe, Geometric Algorithms for Machine Based Optimisation of Acoustic Reflectors, *Proc. of International Conference on Immersive and 3D Audio,* 2021.

[4] John O'Keefe, Applications of Machine Learning Bounding-Breps for Optimised Acoustical Reflectors, *Proc. of Euronoise*, 2021

[5] K. Deb, A. Pratap, S. Agarwal, T. Meyarivan, A fast and elitist multi-objective genetic algorithm: NSGA-II *IEEE Transactions on Evolutionary Computation.* 6 (2), p.182A, 2002.

[6] Krokstad, S Strom, S. Sorsdal, Calculating the acoustical room response by the use of a ray tracing technique, *J. Sound Vib.,* 8 (l), 118-125, 1968.

[7] John O'Keefe, Henrik Moller, Applying Computer Aided Optimization to a Problematic Fan-shaped Auditorium, *Proc. ICA-2022*(to be published)

[8] D. Protheroe, C. Day, Validation of lateral fraction results in room acoustic measurements, *InterNoise Melbourne*, 2014.

[9] ISO 3382-1, 2009. Acoustics - Measurement of room acoustic parameters - Part 1: Performance spaces, *International Organization for Standardization*, Geneva.

[10] K. Deb, R.B. Agrawal, Simulated Binary Crossover for Continuous Search Space, *Complex Systems* Vol. 9, pp.115-148. 1995.

[11] John O'Keefe, A Genetic Algorithm Fitness Function for Acoustic Reflectors in Performing Arts Venues, *Proc. ICA-2022*(to be published).

[12] K. Deb, An Evolutionary Many-Objective Optimization Algorithm Using Reference-point Based Non-dominated Sorting Approach, Part I: Solving Problems with Box Constraints, *Proc. of IEEE*, 2013.

[13] T.J. Cox, W.J. Davies, and Y.W. Lam, The sensitivity of listeners to early sound field changes in auditoria, *Acustica* (79), pp. 27-41, 1993.