ZONE TO ZONE REFLECTOR OPTIMISATION USING A GENETIC ALGORITHM

J O'Keefe O'Keefe Acoustics

1 INTRODUCTION

A number of studies have recently been published by the author using Genetic Algorithms (GA) to optimise acoustic reflectors[1,2,3,4]. The challenge in those early studies being not so much the acoustic performance of the finished reflector design but the geometry of how one constructs the reflector inside the computer in the first place. The geometric methods have been streamlined along the way in an effort to increase computational speed and thus allow a GA to do what it does best: evolve solutions from the largest populations as possible, over as many generations as possible. The present study, it might be argued, takes the optical assumptions of geometric acoustics to the limit. The goal not being to calculate any of the standard acoustic parameters but, rather, to simply ask the question: can everyone in a given zone of receivers (i.e. in a seating area) see an entire zone of sources (e.g. an orchestra on stage) if they were looking at the reflector as if it were an (optical) mirror.

2 REFLECTOR CONSTRUCTION

Traditionally, the perturbation of a geometry like an acoustic reflector has been done inside a simple six sided rectilinear box, often referred to as a Bounding Box. The methods previously developed by the author, and described in references [1] and [2], allow for controlled perturbation of a reflector inside any arbitrarily chosen volume. These volumes are created from Non-uniform Rational B-Spline (Nurb) curves. Surfaces or volumes created by Nurb-curves like this are referred to as Boundary Representations or Breps. The new perturbation control volumes, first documented in ref. [1] are referred to as Bounding Breps or B-Breps. This study will present yet another refinement of the geometric methods used to build a B-Brep and perturb a reflector design inside it.*

In the example to be demonstrated here, we will be optimising a side wall balcony facia. We start with an elliptically shaped cylinder as our B-Brep. Please see Figure 1. A "spinal curve" is drawn through the centre of the B-Brep and a number of points are established on the curve. In the example shown here there are only four points (Figure 1b) but any appropriate number may be chosen. At each one of the points, a "spinal plane" is created perpendicular to the spinal curve. Then the intersections between these planes and the original B-Brep volume are calculated. This gives us four Bounding Nurb Curves (BN Curves), one for each plane (Figure 1c). Each spinal plane has its own U and V axes. Four points are established on the U axes of each of the four planes. These are the points that will be used to build the reflector surface. The points are perturbed along the V axes of their respective spinal planes, their maximum and minimum perturbations being limited by their associated BN Curve (Figure 1d).

In the previous geometric methods[1, 2, 3], the control points were free to move in the U, V and W directions. (The W direction being along the spinal curve). This proved to be a complicated process and one that was very expensive in terms of computer time. In references [4] and [5], the author found that increasing the speed of the Genetic Algorithm (GA), and thus increasing the number of optimising generations, was more important than the accuracy of calculations within each generation. For this reason, in the simplified geometric method described here, the perturbation of the control points is only performed along the V axes of each plane. This, along with the method of using pre-

determined reflection points, introduced in ref. [4], has greatly increased the speed of GA process and, eventually, the quality of the solutions.

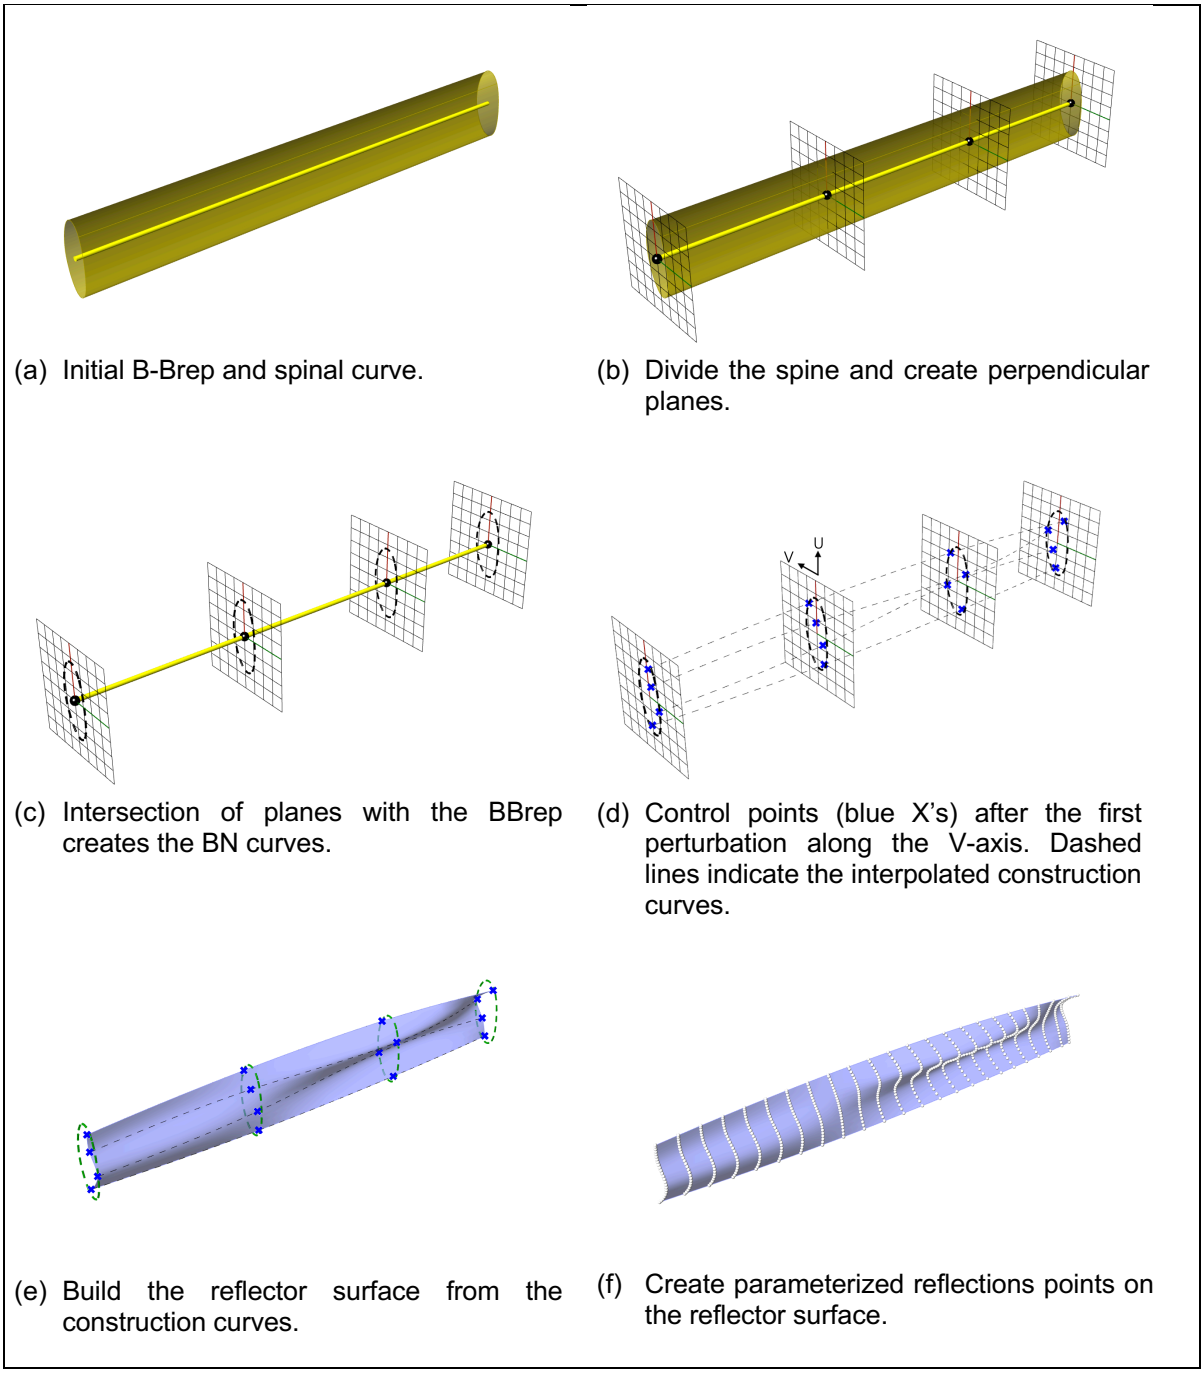

(a) Initial B-Brep and spinal curve.

(b) Divide the spine and create perpendicular planes.

(c) Intersection of planes with the BBrep creates the BN curves.

(d) Control points (blue X's) after the first perturbation along the V-axis. Dashed lines indicate the interpolated construction curves.

(e) Build the reflector surface from the construction curves.

(f) Create parameterized reflections points on the reflector surface.

Figure 1. Construction of a reflector surface

Having created and populated the spinal planes, we now have a grid of 16 control points, four points on each of the four planes. Nurb curves are interpolated from these points, roughly parallel to the spinal curve. These are referred to as the construction curves. The curves span between spinal planes, as shown in Figures 1d and 1e. Adjacent pairs of the construction curves are then used to create surface "strips". Then, finally, these strips are joined together to create the reflector surface, as shown in Figure 1e.

3 CALCULATION PROCEDURE

3.1 Overview

The inspiration for the calculation procedure came during a concert, while sitting on the end balcony. The question being: if the facia of the side wall balcony on the opposite side of the room was a mirror, could a listener on the end balcony see all of the orchestra on stage, inside the mirror. And, perhaps more to the point, could all the listeners on the end balcony see the orchestra inside the mirror.

In conventional reflector design, one starts at a point source location then shapes or orientates the reflector to direct sound towards a receiver point or a receiver zone. In this procedure, we reverse the process, starting at a receiver point, casting rays towards the reflector (in this case the side wall balcony facia) and then determine how many of the reflected rays intercept the source zone (i.e. the orchestra on stage). This is repeated for a number of receiver points. Then, starting with a flat reflector, the Genetic Algorithm (GA) is used to create and perturb a number of surfaces to, eventually, find a single reflector that will best satisfy all the receiver points (i.e. the seating locations).

3.2 Calculating Reflections

In ref. [4], the author adopted an atypical approach for calculating the reflection of rays off a surface. The method proved much faster than the traditional ray bundle/intersection method and, in so doing, allowed the GA to work to its strengths. Namely, employing high population counts, optimised over significantly more generations.

The reflection scheme is based on pre-determined reflection points. Each 3-dimensional reflector surface is parameterised into its standard (s, t) parameters. Each parameter pair (s_i, t_i) represents a point on the reflector. An example of parameterised, pre-determined reflection points is shown in Figure 1f.

The reflection for the i^{th} reflection point is calculated as follows:

1. Evaluate the perturbed surface at the parameter pair (s_i, t_i). This will generate the i^{th} reflection point.
2. Draw a line from the receiver point to the reflection point. This is the incident path. (Remember that the listener at the receiver point is "looking" for the orchestra on stage. What would normally be the acoustic reflection path is now the optical incident path.)
3. Calculate the normal to the surface at the i^{th} reflection point.
4. Create a plane with an origin at the i^{th} reflection point and the normal vector that was determined in Step 3.
5. Create a virtual (mirror image) source point using the receiver point and the plane created in Step 4.
6. Calculate the vector between the virtual image source and the i^{th} reflection point. Use this to generate an infinite ray with its origin at the i^{th} reflection point.
7. Search for an intersection with the source zone surface. If successful, this will be the reflected path.

The advantages of this approach are two-fold. First and foremost, computer intensive intersection calculations are no longer required between a ray bundle and the reflector surface. The second

advantage to this approach is that, no matter how the reflecting surface is perturbed, the reflection points will always be in the same relative positions with respect to the surface's centre and its edges. This allows for a simple and systematic evaluation of diffraction effects using the methods developed by Rindel[6]. An advantage that has been put to use in ref. [5], studying the diffraction effects on flat surface stage reflector arrays.

4 FITNESS FUNCTIONS

Most of the Genetic Algorithm (GA) optimisations in this study were performed with a three objective routine, a modification of the Non-dominated Sorting Genetic Algorithm (NSGA-II)[7]. An early version of the work employed a two-objective optimisation version of NSGA-II but closer analysis of the results suggested the need for a third objective in the optimisation procedure. The fitness functions for the three objectives are presented below. Briefly summarised, the goals are to: (i) maximise the source zone area seen by the listeners on the balcony; (ii) minimise the amount of energy scattered by the reflector away from the listeners; and (iii) ensure that the listeners' view of the source zone is uniform (i.e. avoiding "hot spots" or "dead zones" on stage).

4.1 Area Fitness

As is often the case in GA optimisation, the goal will be to minimise the fitness functions. So, although we want to maximise the visible area, the actual function works towards a minimum. The area fitness function becomes:

$$Area\ Fitness\ = \frac{1}{N_{RecPts}} \sum_{i=0}^{N_{RecPts}} \left[1 - \frac{S_{visible,i}}{S_{source\ zone}}\right] \quad (1)$$

where:

N_{RecPts} = the number of receiver points in the receiver zone
$S_{visible,i}$ = the surface area of the source zone that is visible from the ith receiver point.
$S_{source\ zone}$ = the total surface area of the source zone.

4.2 Scattering Fitness

In an effort to maximize the useful reflection area of the optimized reflector, the second goal will be to minimize the scattered energy, employing the fitness function:

$$Scattering\ Fitness\ = \frac{1}{N_{RecPts}} \sum_{i=0}^{N_{RecPts}} \frac{1}{N_{RefPts}} \sum_{j=0}^{N_{RefPts}} Points_{failed} \quad (2)$$

where:

N_{RecPts} = the number of receiver points in the receiver zone
N_{RefPts} = the number of pre-determined reflection points on the reflector
$Points_{failed}$ = the number of reflection points that failed to cast reflections to the source zone from the ith receiver point

4.3 Spreading Fitness

The procedure, as described above, casts rays from a receiver point, to the reflector and, hopefully, to the source zone. The goal of the spreading fitness function is to ensure that the intersections of the reflected rays with the source zone are distributed as uniformly as possible. A similar fitness function was developed in ref. [1], which we have modified and renamed here:

$$Spreading\ Fitness\ = \frac{1}{N_{RecPts}}\sum_{i}^{N_{RecPts}}\frac{1}{N_{grid}}\sum_{j=0}^{N_{grid}} 1 - e^{\frac{-d_{NN}}{\tau_{dist}}} \quad (3)$$

where:

N_{RecPts} = the number of receiver points in the receiver zone
N_{grid} = the number of evenly distributed grid points on the source zone surface
d_{NN} = distance from the j[th] grid point to its nearest neighbour (NN) intersection point
τ_{dist} = the convergence coefficient associated with distance

5 RESULTS

The geometry for the example demonstration discussed here is shown in Figure 2. The audience layout is meant to be representative of a typical 20 m wide shoebox shaped concert hall with a single balcony. The source zone is similar to that of an orchestra on stage. The receiver zone is restricted to only one half of the balcony. The idea being that the (blue) side wall balcony facia is best used to create lateral reflections on the opposite side of the room. The receiver zone is sparsely populated with only 9 receiver points, primarily for computational speed.

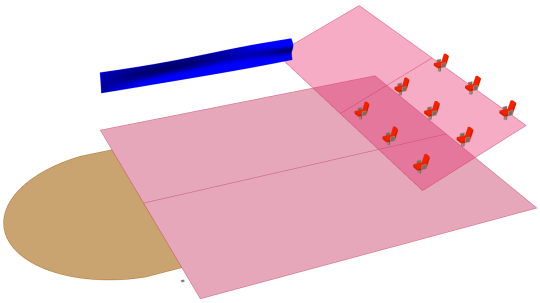

Figure 2 Perspective view of the example room geometry.

As mentioned above, the original intention of this study was to employ a standard two-objective optimisation. Figure 3 shows some of the results from one of the early 2 objective optimisations and demonstrates the need for a 3 objective optimisation search. In this case, the optimisation was driven by Equations 1 and 2, i.e. for Area and Scattering Fitness but not Spreading Fitness. The Pareto Optimisation graph in Figure 3a may be interpreted as follows.

The big blue and yellow dot is the reference point. Diffusion or scattering studies often use a sphere or cylinder as a point of reference. The blue and yellow dot indicates the fitness function value for a semi-cylindrical facia reflector the same size and in the same location as our optimised reflectors. The semi-cylinder has an Area/Scattering Fitness value of (0.33, 0.93). This suggests that with a semi-cylindrical facia reflector, listeners on the opposite (end) balcony will not be able to see 33% of the orchestra on stage if they tried to find it as a mirror reflection off the side wall balcony facia. The scattering fitness value of 0.93 suggests that the semi-cylindrical reflector scatters 93% of the visibility rays away from the intended receiver points on the balcony. Or, put another way, only 7% of the reflector surface is directing sound towards the nine listeners in receiver zone.

The solid and open circles in Figure 3a indicate how the solutions have been sorted into the first and second Pareto Fronts by the Non-dominated Sorting Genetic Algorithm (NSGA-II)[7]. The solid circles are for Pareto Front 1, the open circles for Pareto Front 2. The goal of the optimisation was to minimise the Area and Scattering Fitness Functions and all of the solutions on both Pareto fronts are, indeed, closer to the graph's origin than the reference point for the semi-cylinder.

Figure 3b shows the results from one of the better solutions in Figure 3a. The one indicated by the arrow. It has an Area/Scattering Fitness of (0.20, 0.71). The view of the stage, in Figure 3b, is for the end seat in the 2nd row, the one that has been circled in yellow. The grey spheres on the right side of Figure 3b indicate where the visibility rays from the balcony front have intersected the source zone. The white surface indicates the part of the source zone that is visible from the balcony seat, as calculated by the intersection points. Their irregular distribution leads to an erroneous over estimation of the visible source zone area. Hence the need for a 3 objective optimisation search.

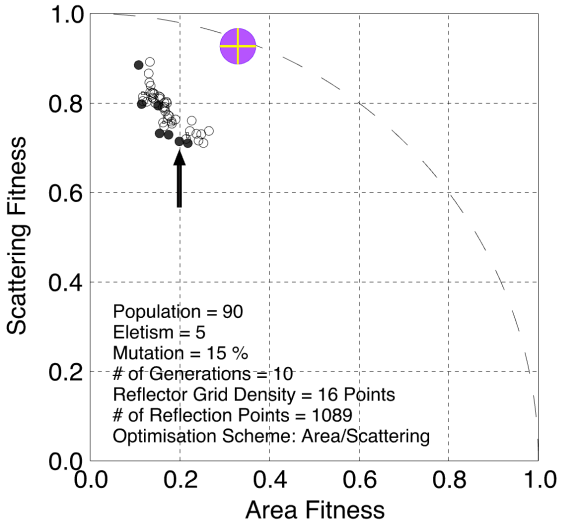

Figure 3a. Pareto graph for a 2 objective optimisation search. The blue and yellow dot indicates the performance of the reference semi-cylindrical reflector. Solid circles are for the 1st Pareto Front, open circles for the 2nd.

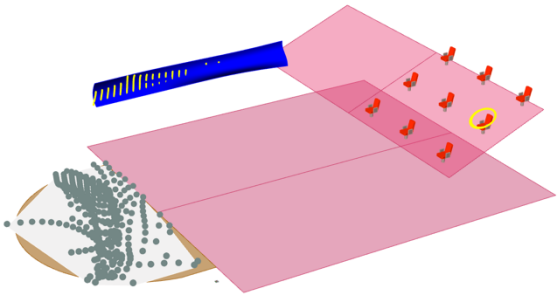

Figure 3b. Image of the view of the stage from the seat circled on the balcony, using the reflector given by the solution indicated by the arrow in Figure 3a. The yellow dots on the blue reflector indicate the reflection points. The grey spheres indicate intersections with source zone.

A 3 objective optimisation was therefore performed on a population of 315 over 50 generations. The resulting Pareto Graphs are shown in Figures 4a and 4b. Quite impressively, all of the solutions have been sorted onto the first Pareto Front and appear to have reached their best possible fitness values. This would not have been possible without such a high population count or the high number of optimising generations. And neither of those would have been possible without the computer time saving methods described above, notably the concept of pre-determined reflection points.

Figure 5 presents a more in depth analysis of one of the better solutions taken from this optimisation. The solution has an Area/Scatter/Spreading Fitness value of (0.08, 0.75, 0.82). It is being compared to the reference point semi-cylindrical reflector on a seat by seat basis. The open bars are for the semi-cylinder, the solid bars for the optimised reflector. In every seat, the optimised reflector provides

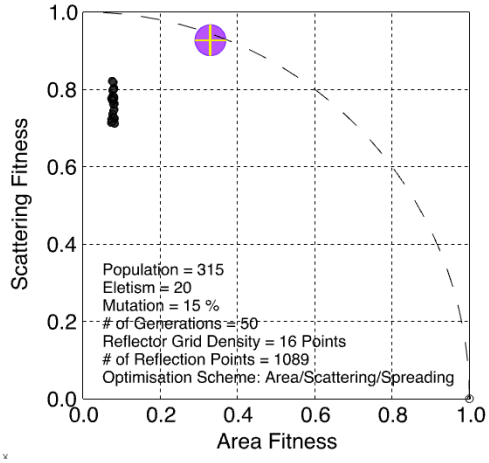

Figure 4a. Area/Scatter fitness Pareto analysis.

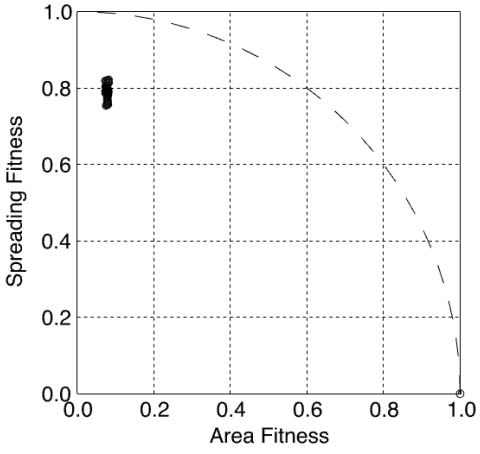

Figure 4b Area/Spreading fitness Pareto analysis.

a significantly better view of the source zone than the benchmark semi-cylinder. In the range of 90 to 95%.

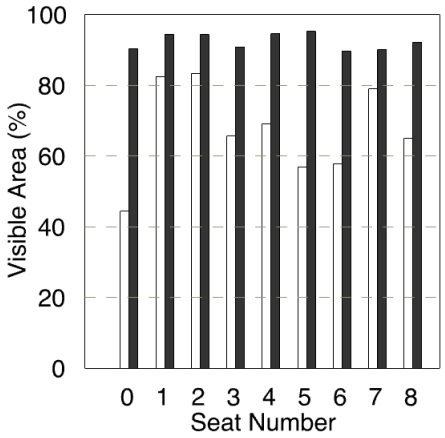

Figure 5. Percentage of visible areas from the 9 balcony seats. Solid bars for the optimized reflector, open bars for the reference semi-cylindrical reflector.

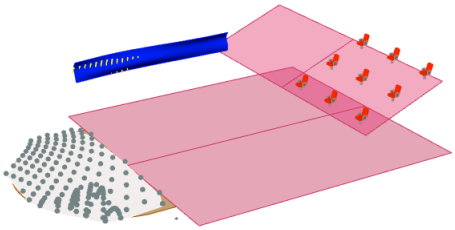

Figure 6. View of the stage from the same seat as depicted in Figure 3b. Including the 3rd optimization objective provides a more uniform view of the stage.

Figure 6 shows the visible area of the stage from the same seat as depicted in Figure 3b, this time using the reflector analysed in Figure 5. Including the 3rd objective (Spreading Fitness) has produced a much more uniform view of the stage and, hence, a more accurate estimation of the visible surface area.

6 CUT-ON FREQUENCIES

As mentioned at the outset, the optimisations described here have centred around the optical visibility of the source zone. The acoustical implications should, of course, be taken into account. With that in mind, a brief study of the reflecting panels' cut-on frequencies was undertaken. Rindel has outlined a method for approximating the diffraction attenuation of rectangular reflectors similar to the ones considered here, i.e. reflectors with a large ratio between length and height[8]. Two cut-on frequencies are calculated, one for the short dimension (f_{g1}) and one for the long dimension (f_{g2}). Above these frequencies, the reflector will cast uniform reflections, at lower frequencies diffraction effects will attenuate the reflections at a rate of 3 dB per doubling of frequency between f_{g1} and f_{g2} and 6 dB per doubling below f_{g2} .

A cursory study of the cut-on frequencies was carried out as follows. A single source-receiver combination was chosen, with the source and receiver points at the centroids of the source and receiver zones respectively. Reflections were cast off of the original flat panel at 1089 pre-determined reflection points. For the long dimension (11.4 m), the f_{g2} cut-on frequencies are between 15 and 18 Hz and, hence, are not a concern. For the short dimension (1.3 m) the cut-on frequencies range from 1545 to 7554 Hz. A contour map of these frequencies is shown in Figure 7. As indicated by the yellow dots in Figure 3b and Figure 6, most of the successful reflections are cast from points on the reflector

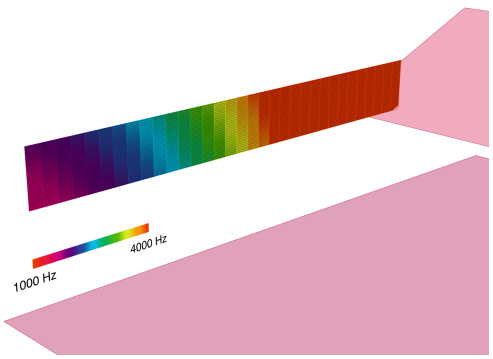

Figure 7 Contour map of f_{g1} cut-on frequencies superimposed on the balcony facia.

that are closest to the stage. The contour map in Figure 7 suggests that the cut-on frequencies in this area are in the range of 2000 to 2500 Hz. An octave below these frequencies (i.e. in the 1000 Hz octave band) reflection levels will be attenuated by 3 dB. In the 500 Hz octave band they will be attenuated by 6 dB, and so on into the lower frequencies.

7 CONCLUSIONS

A new procedure has been presented to optimise a reflector casting reflections from one zone to another zone. If one accepts the inherent limitations of geometric acoustics, i.e. the optical equivalency, the new procedure does, indeed, improve how well a group of listeners might "see" another group of sound sources as a mirror reflection inside the reflector. Using a 3 objective Genetic Algorithm (NSGA-II), the receiver zone's view of the source zone has been significantly optimised. A standard semi-cylindrical reflector provided a view of 67% of the source zone. An optimised reflector of the same size and in the same location provided a view of 92% of the source zone.

8 REFERENCES

1. O'Keefe, J. Geometric Algorithms for Machine Based Optimisation of Acoustic Reflectors, Proc I3DA, Bologna, 2021.
2. O'Keefe J., Applications of Machine Learning Bounding-Breps for Optimised Acoustical Reflectors, Proc. of Euronoise, 2021
3. O'Keefe, J., A Two Stage Embedded Genetic Algorithm to Optimise Ceiling Reflections. Proc. of Euroregio BNAM 2022 Joint Acoustics Conference; 9th-11th May 2022; Aalborg, Denmark.
4. O'Keefe, J., Genetic Algorithm Fitness Functions for acoustic reflectors in performing arts venues, Proc. 24th ICA (2022), Seoul, South Korea.
5. Rindel JH. Attenuation of Sound Reflections due to Diffraction, Proc. of NAM Conference 20-22 August 1986.
6. O'Keefe, J. A Genetic Algorithm to Optimise Stage Reflectors for Self and Other Reflections, Proc. 10th Forum Acusticum, Torino (2023).
7. Deb, K. Pratap, A. Agarwal, S. Meyarivan, T. A fast and elitist multi-objective genetic algorithm: NSGA-II IEEE Transactions on Evolutionary Computation. 6 (2), 2002, p.182
8. Rindel, J.H., Acoustic Design of Reflectors, Proc. of IOA, Vol. 14 (2), 1992.