DOI: 10.61782/fa.2023.0197

A Genetic Algorithm to Optimise Stage Reflectors for Self and Other Reflections

John O'Keefe[1*]
[1] O'Keefe Acoustics, Toronto, Canada

ABSTRACT

One of the fundamental challenges in stage acoustic design is striking a balance between Hearing of Self and Hearing of Other. Optimising two objectives at once is exactly what the Non-dominated Sorting Genetic Algorithm (NSGA-II) was designed for. Two Genetic Algorithm (GA) fitness functions have been developed for Hearing of Self and Hearing of Other to assist in the design of an overhead reflecting canopy above an orchestra platform. These are then used to govern the NSGA-II algorithm. Diffraction effects on the panels that make up the canopy have also been studied. Results indicate diffraction attenuations typically in the range of 3 dB at 500 Hz. Including the diffraction calculations can significantly slow the run times of the GA, depending on the desired accuracy. It is shown, however, that the accuracy of these calculations may be safely reduced during the GA's evolving optimisations, provided that more accurate calculations are performed during the final evaluation of the chosen optimal design. Exchanging the time used for diffraction calculations during the optimisation for higher GA population/generation counts leads to better canopy designs.

Keywords: *Genetic Algorithms, Stage Acoustics, Multi-objective Optimisation.*

1. INTRODUCTION

Evolutionary algorithms, such as a Genetic Algorithm (GA), fall into the class of computations known as meta-heuristics. A process that some acousticians have troubles with. A meta-heuristic is simply a routine that governs heuristic calculations, i.e. calculations that sacrifice accuracy for speed. The concern of some being that sacrificing accuracy will lead to results that might look good but do not reflect reality. This despite the fact that many successful concert halls have been designed recently using ray based computer modelling which, when compared to more accurate Boundary Element Models (BEM) or physical scale models, might be considered a less accurate, time saving heuristic. The goal of this paper will be to propose a GA for stage canopy design, examine the effects of including diffraction in the evolutionary optimisation process, then consider the question: is the heuristic compromise of excluding or reducing diffraction accuracy valid?

2. ORCHESTRA CANOPY

An image of the stage, "orchestra" and canopy used in this study is shown in Figure 1. The canopy is tilted slightly at an angle of 10° to horizonal, approximately 10 m above the stage floor. It is made up of 24 square panels, each 1600 x 1600 mm in size. There is a 100 mm gap between panels in the X-direction and a 140 mm gap in the Y-direction, the Y-axis being parallel to the longitudinal centre-line of the

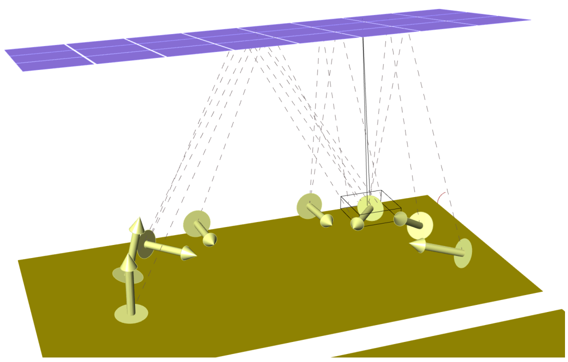

Figure 1. Reference canopy reflector and orchestra instrument directivities..

*Corresponding author: john@okeefeacoustics.com.

room. The Genetic Algorithm (GA) rotates these panels randomly about the X and Y axes. The range of the random rotations is ± 90° but rotations greater than ± 30° rarely survive into the next generation. The GA does not change the internal geometry of the square panels. They always remain flat. The 10 m height and 10° tilt of the canopy are also fixed throughout all the optimisations.

Only eight instruments (or Chairs) are used to represent the orchestra. The number being limited to increase the speed of the GA, thus allowing a larger optimisation population and a larger number of generations. The instruments and their approximate radiation directivities have been laid out to represent the so-called "American seating" plan, i.e. violins stage right, cellos and basses stage left, horns and woodwinds upstage centre. The apparently impractical layout of the Chairs shown in Figure 1 is not intended to represent 8 individuals on stage but, rather, 8 groups of individuals.

In ref. [1], the author introduced a faster method for calculating first order reflections in a GA optimisation. A more complete description may be found in [2]. Rather than cast thousands of rays from a point source ray bundle, pre-determined points are established on each of the panels from which reflections are cast. Except where noted otherwise, the panels were populated with 9 reflection points each.

3. REFLECTION LEVELS & FITNESS FUNCTIONS

3.1 Reflected Sound Levels

Before defining the Fitness Functions upon which the optimisations are based, we must first introduce how the reflected sound levels were calculated. The procedure follows the attenuation of sound reflections method, as described by Rindel in Equation 1 of ref. [3]

$$\Delta L = \Delta L_{dist} + \Delta L_{abs} + \Delta L_{diffr} + \Delta L_{curve} + \Delta L_Q \quad (1)$$

Where:

$\Delta L_{dist} = 20log\dfrac{a_o}{a_1 + a_2}$	a_0, a_1 and a_2 being the direct, incident & reflected path lengths, respectively
ΔL_{abs}	0 dB
ΔL_{diffr}	as defined by Rindel in [3]
ΔL_{curve}	0 dB
ΔL_Q	Directivity of instrument

Note that for the purposes of this study, we have assumed perfectly reflecting panels, i.e. no absorption and therefore ΔL_{abs} = 0 dB. Although the GA alternates the tilt of the individual panels in the canopy, it does not change their internal geometry. They always remain flat, hence ΔL_{curve} = 0 dB. All calculations using Eqn. (1) were performed with a centre frequency of 500 Hz.

The directivity of the musical instruments (ΔL_Q) is an addition to Rindel's Equation 1. To increase computational speed, the directivity patterns are simplified elliptical versions of more realistic patterns e.g. as quantified by Pätynen, et al. [4]. The ΔL_Q directivity patterns that have been used are meant to provide an approximation of a given instrument's directivity. For example, a violin radiating mostly in the vertical direction towards the ceiling or the horns and winds radiating horizontally towards the audience.

In the following, a "Self" reflection is one that originates from an instrument played by a musician that has been successfully returned back to the same musician. "Other" reflections are those that originate from instruments played by other members of the orchestra that have be successfully cast towards and received by the musician in the "Self" location. Thus, when Self and Other reflection levels are compared, the source of the reflections will vary but the receiver is always at the Self location.

The optimisation implements the Non-dominated Sorting Genetic Algorithm (NSGA-II) algorithm developed by Deb et al.[5]. As in often done in GA optimisations, the goal will be to minimise the two Fitness Functions. NSGA-II is a Pareto based optimisation, so a goal towards minimisation suggests that the best solutions on the Pareto graphs will be the ones closest to the origin.

3.2 Self Fitness Function.

In the parlance of Genetic Algorithms (GA), the canopy above the stage is the genome and the individual reflectors that make up the canopy are the genes. A fitness function is, essentially, the question you are asking the computer. It establishes the goals to which the optimisation design is evolving towards. For Self Fitness, we want to know how much better the optimised canopy, with all its tilted panels, is when compared to a similar but completely flat canopy. The fitness function for Hearing of Self is determined by comparing the successful Self reflections cast by the overhead reflector canopy in question to the reference (flat) canopy. The reference canopy is shown in Figure 1. All of the individual panels that make up the reference canopy are in the same plane, i.e. tilted at an angle 10° to horizonal. To test the influence of diffraction and source directivity, the routine that has been developed can either include, exclude or vary the accuracy of these effects. If the reflection calculations for the genome in question include these effects, the reference reflections do so as well.

The linear representation of Equation (1) will be referred to as:

$$E_{reflected} = 10^{\Delta L/10} \quad (2)$$

The Hearing of Self fitness is calculated as follows:

$$A_i = \sum_{j=0}^{N_{SelfRefns}} E_{reflected,genome,j} \quad (3a)$$

$$B_i = \sum_{j=0}^{N_{SelfRefns}} E_{reflected,reference,j} \quad (3b)$$

$$Self\ Fitness = \frac{1}{N_{chairs}} \sum_{i=0}^{N_{chairs}} e^{-\tau A_i / B_i} \quad (3c)$$

Where:

$N_{SeflRefns}$ is the number of Self reflections
$E_{reflected,genome,j}$ is the relative sound energy of the j^th reflection at the i^th Chair for the canopy genome in question
$E_{reflected,reference,j}$ is the relative sound energy of the j^th reflection at the i^th Chair for the flat reference canopy
A_i is the total Self reflected energy for the i^th Chair cast by the genome in question
B_i is the total Self reflected energy for the i^th Chair cast by the reference, flat reflector
N_{chairs} is the number of musicians' Chairs on stage
τ is the convergence coefficient, typically in the range of 0.15.

3.3 Other Fitness Function

The fitness function calculation for Hearing of Other is similar to that for Hearing of Self. In this case, however, the sound source locations are at every Chair on stage except the Self Chair and, in all cases, the receiver location is at the Self Chair. The convergence coefficient (τ) for Self is set to 0.15, as noted above. For Other Fitness, τ is set to 0.10.

4. TESTING THE FITNESS FUNCTIONS

4.1 Epsilon Constraint Method

One of the issues when searching for optimised designs is the simple fact that we cannot know the extent of the search space beforehand. The Epsilon Constraint method, proposed by Laumanns et al. [6] addresses this concern. An explanation of the procedure has been presented by the author in ref. [1]. Briefly stated, in a two objective optimisation search, the ε-constraint method holds one of the objectives at a fixed value (call it Fitness A) then lets the Genetic Algorithm "drive" the search to the lowest (i.e. best) possible value of the other objective (Fitness B). Remember, the Fitness Functions, defined in Section 3, are intended to be minimised. After the GA has found the lowest (free moving) Fitness B for the (fixed) Fitness A, that fixed value of Fitness A is decreased incrementally by a value of ε. The process is then repeated as often as required to find the Pareto Front of the multi-objective optimisation search.

Results of the ε-constraint searches are shown in Figures 2 and 3. The procedure and graphs are explained as follows. Figure 2 shows the search for the lowest possible fitness values for Self, using fixed values of Other. Starting close to the upper right hand corner of the graph, the Other Fitness value is fixed at 0.6. Then a population of 25 different canopy designs is optimised over 25 generations. The Self Fitness values for each member of the population are indicated by the small triangles pointing towards the left. We see that the best possible Self Fitness value (i.e. the lowest) for an Other Fitness that has been fixed at 0.6 is approximately 0.75. Then, the process just described is repeated, only this time the fixed Other Fitness is decreased by the value ε = 0.08. The procedure is repeated for the full range of the search. In this case the search was from Other Fitness = 0.6 to 0.2, as indicated by the large arrow beside the graph.

The big blue and yellow dot in both figures, shown at position (0.85, 0.45) indicates the Self/Other Fitness value for the reference (flat) reflector canopy. It can be seen that, insofar as Hearing of Self is concerned, a canopy of optimised, tilted panels can make a slight improvement over a flat canopy but, apparently, not by much.

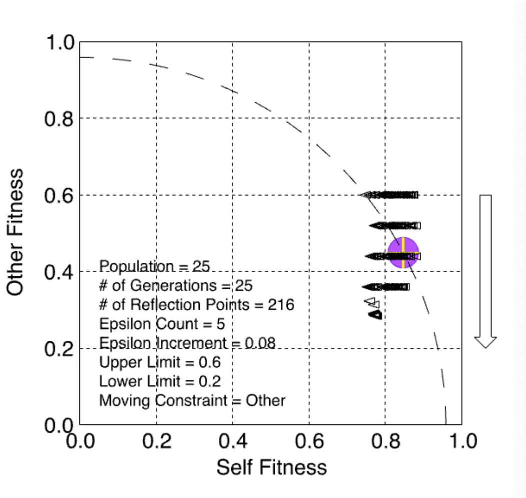

Figure 2. Epsilon Constraint search for the best Self Fitness, incrementally moving the Other Fitness from 0.6 to 0.2.

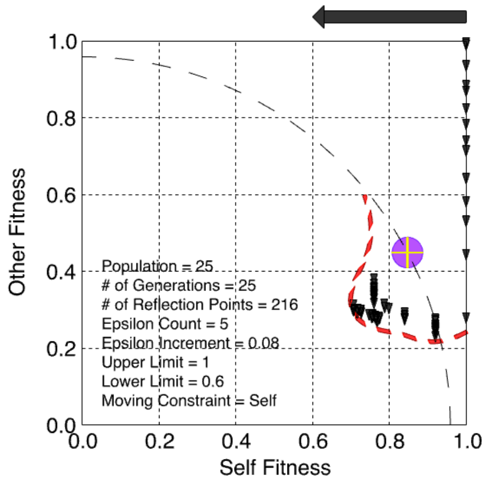

Figure 3. Epsilon Constraint search for the best Other Fitness, incrementally moving the Self Fitness from 1.0 to 0.6.

Figure 3 shows the ε-constraint search for the best possible values of Other Fitness. Here we see that it seems to be easier to improve the canopy's delivery of Other reflections than it is to improve Self reflections. Note how the triangles in Figure 3 are further below the big blue dot than the triangles are to the left of the dot in Figure 2. The thick dashed red line in Figure 3 is an approximation of where the true Pareto Front might be; a combination of the data points discovered in the Figure 2 and 3 searches. The Pareto Front is the border between the feasible and infeasible solutions for this search. For a search with a population of 25, optimised over 25 generations and with a reflection point density of 9 points per panel, the dashed red line is, in a sense, the best we can hope for.

To place this into context, consider first the two extremes for Self and Other Fitness. A value of (1.0, 1.0) means that the optimisation has returned a design with exactly the same performance as the original flat canopy. There has been no improvement. A value at the other extreme (0.0, 0.0) would mean that the Self reflection levels have been increased by 15.6 dB and the Other reflections levels by 8.5 dB. Neither of which are likely or perhaps even desirable. More realistically, moving from the flat canopy reference point (the blue dot) to the vicinity of the Pareto Front, a Self/Other value of (0.75, 0.20) implies a 3 dB increase in both Self and Other reflection levels. This means that an optimisation that returns a fitness value close to the Pareto Front has doubled or almost doubled the useful Self and Other reflections.

4.2 Comparison with Human Design

When acousticians are designing a canopy reflector, they will sometimes refer to a suggestion proposed by Meyer [7], where three of the four edges of the canopy are tilted down. Indeed, the author has employed it on several occasions. It's often useful to compare computer-aided design optimisation with "human" design. The computer does not always win! In this case, however, it did. Using the fitness functions defined in Section 3, the tilted-edge canopy, shown in Figure 4, performs essentially no better or worse than the flat canopy. It has a Self/Other Fitness value of (0.86, 0.46), compared to the flat canopy's value of (0.85, 0.45).

This is not to say, however, that the computer-aided design has created a better orchestra canopy. Although the current routine can include instrument directivity patterns, the ones used so far are rudimentary at best. In addition, the routine does not yet address hearing directivity at the receiver end of a reflection. Meyer's tilted edge concept [7] was intended to take advantage of both source and receiver directivities. In this study, the GA will find apparently superior canopy designs but they are only superior within the definition of the Fitness Functions established and the modelling that has been developed so far.

5. DIFFRACTION EFFECTS

The amount of diffraction attenuation imparted on a reflection will depend on where it lands on the reflector [3]. The concept of pre-determined reflection points on the reflector, proposed in ref. [1] presents the opportunity for a systematic study of diffraction effects.

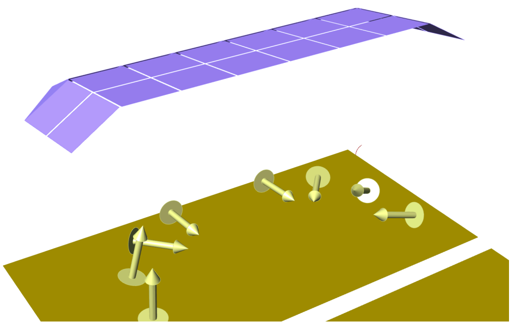

Figure 4 The tilted canopy, based on Meyer's suggestions [7].

Rindel gives two approximate "cut-on" frequencies above which geometric acoustic principles may be assumed. He refers to them as f_g and f_{g1}. Below these frequencies, diffraction effects should be taken into account. The f_g cut-on frequency is based on the assumption that the reflection is incident at the centre of the reflector. If one knows the exact location where a reflection lands on the reflector (as we do with the pre-determined reflection points) the more accurate of Rindel's two cut-on frequencies is f_{g1}. Figure 5 shows the distribution of both cut-on frequencies taken from a trial run of the routine. Observing the f_{g1} distribution, the graph suggests that we should consider diffraction effects at 1250 Hz or lower. All of the diffraction calculations in this study have therefore been calculated at a centre frequency of 500 Hz, i.e. where the diffraction effects should be significant.

As explained in the Introduction, the very nature of a heuristic calculation, such as a Genetic Algorithm (GA), is to exchange calculation accuracy for calculation speed. The processing speed for the GA that has been developed for this study is dependent on the accuracy of the diffraction calculation performed during the optimisations. The accuracy being related to how many pre-determined reflection points are established on the panels.

What might be considered a "quick" optimisation run can be done on a population of 10 over 10 generations with 9 reflection points per panel. This currently takes a little less than an hour of computing time, which we will round up to a full hour and use as a benchmark. Table 1 shows a summary of comparative running times as the diffraction calculation accuracy is increased.

Table 1. Calculation Speeds

# of Reflection Points	Time (hours)
9	1.0
16	1.6
36	3.8
64	6.2
121	12.5

It can be seen from Table 1 that increasing the diffraction calculation accuracy is costly in computer time. The growth rate is approximately quadratic. The question then arises: how how much accuracy do we need and does the accuracy influence the final result?

Figures 6 and 7 show a statistical analysis of the diffraction attenuation effects calculated during the runs that led to Table 1. The data is from a single canopy sample, taken from the population of 10. The attenuations were calculated at a centre frequencies of 500 Hz. Figure 6 shows the distribution of the diffraction attenuations for Self measurements, calculated as outlined by Rindel in ref. [3] and described here in Section 3.1. Figure 7 shows the same for Other reflections. Starting with Figure 7, it is apparent that approximately 50% of the Other reflections experience diffraction attenuations in the range of 3 dB, regardless of the density of the reflections points on the panels. At the (3.0 - 3.5 dB) distribution bin, notice how, on either side of the bin, there is no significant difference

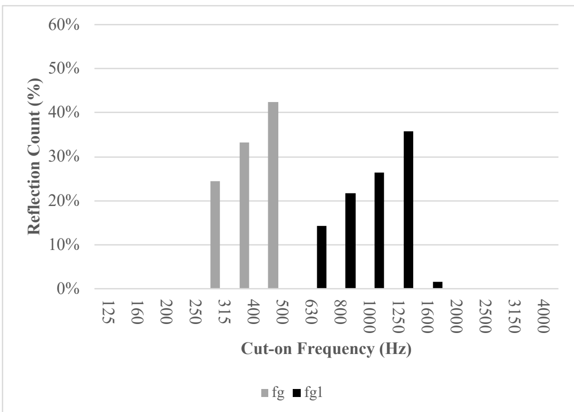

Figure 5 Distribution of cut-on frequencies for geometric acoustics assumptions.

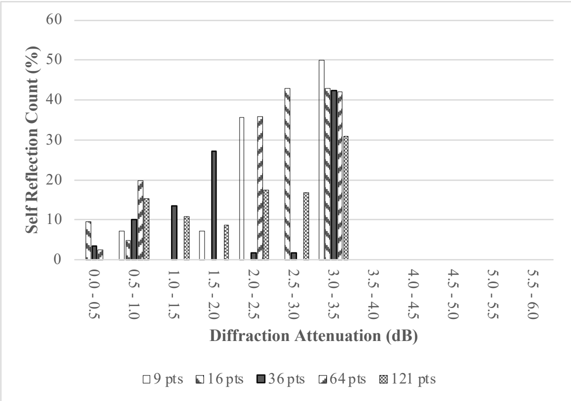

Figure 6. Statistical distribution of diffraction attenuations experienced by "Self" reflections, depending on the number of reflection points on each reflecting panel.

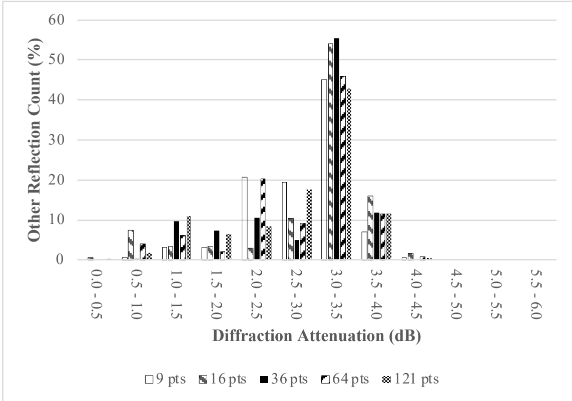

Figure 7. Same as Figure 6 but for the "Other" reflections.

between the 9 reflection point bar and the 121 point bar. This suggests two things. First, that including diffraction in the Genetic Algorithm (GA) evolution calculations might be important – it can have an effect of up to 3 dB. Perhaps more importantly though, we see that the effect is immediate. It is significant at the lowest reflection point resolution (9 points) and does not change much as the resolution is increased to as high as 121 points per panel. Referring to Table 1, this suggests that, insofar as diffraction attenuation accuracy is concerned, a 12.5 hour calculation (for 121 points) may be no better than a 1.0 hour calculation (9 points).

The big question however is this: is the improved accuracy, provided by diffraction calculations, which take longer, more useful for the evolution towards an improved canopy design? Or, would it be better run a faster GA on a larger population over more generations of optimization? In effect, letting the GA do what it does best – randomly evolve towards a better solution.

The following experimental comparison was performed. Optimization runs were performed under the same conditions used for the ε-constraint search described in Section 4. That is, on a population of 25, optimising the design over 25 generations. Like the ε-constraint search, this run included the effects of diffractions, employing the minimum calculation accuracy; a reflection point density of 9 points per panel. These optimisation runs typically took about 7:30 hours of computer time.

A second set of (comparison) optimisations was performed, similar to the first but with a higher reflection point density (36 points per panel) and, thus, higher diffraction accuracy during the optimisation. These runs took considerably longer; in the range of 40 hours. All 50 canopy designs evolved from the two competing optimisation runs were evaluated according to the fitness functions defined in Section 3. In both cases, however, the fitness function calculations were performed using the reflection point density employed in the ε-constraint search, i.e. 9 points per panel. Which is to say that, although the two samples were evolved with different diffraction accuracies, they were evaluated with the same diffraction accuracy.

The results are presented in two ways. Figure 8 compares the fitnesses in a Pareto graph similar to the one in Figure 3. We see that the design solutions found by the faster search, the ones with only 9 reflection points (solid dots), are closer to the Pareto Front and are, therefore, slightly better than the slower 36 point search solutions (open circles).

Figure 9 compares the same data but in a less abstract fashion. Rather than fitness functions, we compare the improvement in Self and Other reflection levels that the optimised reflectors have provided. The change in level being the difference in reflected levels cast by the optimised canopy compared to the reference flat canopy. Again, we see that the more accurate diffraction

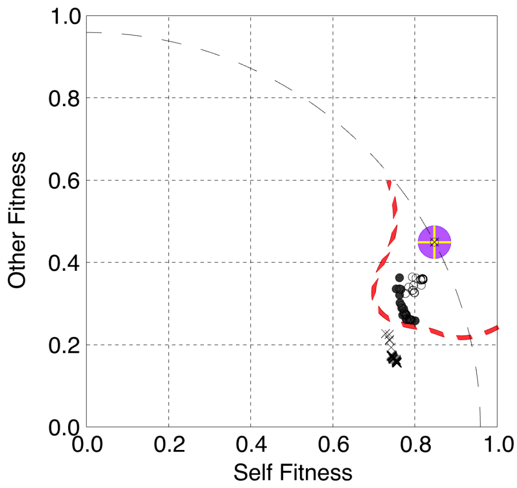

Figure 8. Pareto analysis of 9 reflection point optimisation (solid dots) and 36 point optimisation (open circles). The thick red dashed line is the Pareto Front.

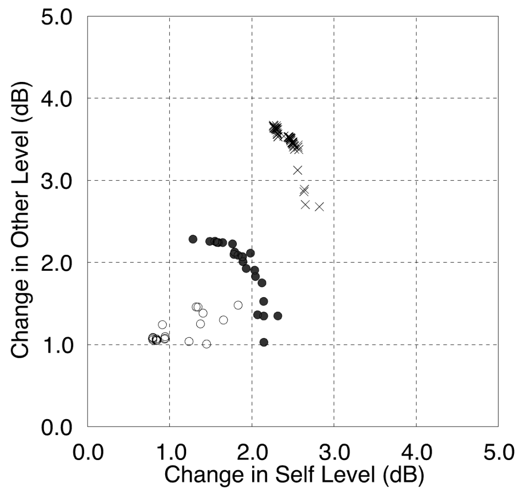

Figure 9. Change in Level analysis of 9 reflection point optimisation (solid dots) and 36 point optimisation (open circles). The X's indicate the change in level for the 61/61 run.

calculations applied during the evolution of the canopy designs did not render better results. The 25 canopy designs evolved during the lower resolution, 9 point search (solid dots) have, in all cases, done a better job at increasing both Self and Other reflections levels, when compared to the 36 reflection point search (open circles).

From this and other experiments like this, we can conclude that reducing the accuracy to which one calculates diffraction *during* the evolution of designs in a Genetic Algorithm (GA) does not have a detrimental effect on the quality of the final results of that evolution. If anything, the final results might actually be better using the lower resolution reflection point density. This is not to say however that the higher accuracy diffraction calculations should be abandoned altogether. Scale model studies [8] on a canopy similar to the one used in this experiment have proven the validity of Rindel's method [3]. It would be wise to employ it. The higher accuracy calculations can be easily performed *after* the GA optimisation has finished. In the cases demonstrated here, that would be on the final 25 designs of the GA's population. This takes only a matter of seconds of computer calculation time, compared with the hours it takes if diffraction is included during the evolutionary process.

Figures 8 and 9 also show data from a final experiment, one in which the time used for more accurate diffraction calculations during the GA's evolution was exchanged for time spent on a higher population/generation count. In this case, a population of 61, optimised over 61 generations. The results are indicated with the X's. As might be expected, the results are even better than those found in the Pareto Front search (the thick dashed red line). This is because that original Pareto Front search was limited to a population of 25 over 25 generations. The extra time that would have been spent on diffraction calculations has allowed the GA to evolve designs that further increase the Self and Other reflected sound levels, in the range of 2.5 dB and 3.5 dB respectively, as seen in Figure 9.

6. DISCUSSION

6.1 The Question of Balance

Before concluding, we should point out that the routine described here does not optimise for the Balance between Self and Other. Rather, the fitness functions described in Section 3 encourage the GA to improve both Self and Other reflected energy levels at the same time, whilst maintaining a balance between the two. Self/Other Balance remained fairly uniform throughout the experiments in this study, in the range of -6.0 dB. Indeed, it would be difficult to design for an optimum Self/Other Balance because it is very difficult to state what the goal might be.

In 1985, Naylor [9] found that, in anechoic conditions, the optimum balance between a performer's instrument and others' could be anywhere between -23 dBA and +7 dBA depending on whether the music is in unison, double or triple counterpoint. Later, in 2009, Dammerud found that 81% of players agreed with the statement "Acoustics for performers depends on the correct balance between hearing yourself and hearing other players" [10]. Despite this, not much progress has been made on the subjective percept of Balance since Naylor's initial findings. Leaving Wenmaekers et al. to state in 2017: "Optimal ranges for (stage) acoustic parameters are not confirmed or non-existent." [11]

6.2 Future Work

The most obvious improvement to the algorithm described here might be to include more realistic directivity patterns for the musical instruments on stage. This may be more important than the diffraction effects which, it appears, may be safely omitted during the optimisation evolutions. Including the directionality of hearing would also be beneficial.

Further development of the fitness functions presented in Section 3 might also be considered. Currently they are designed to increase either Self or Other reflections to an unspecified level, implicitly assuming that more is better. This, of course, may not always be the case. Specified levels in the fitness functions might address this concern.

7. CONCLUSIONS

To answer the question set out at the beginning of this study: can the accuracy of diffraction effects be safely reduced during the evolutionary optimisation design of a stage canopy, the answer is conditional. If the Genetic Algorithm (GA) run is reasonably fast, say on the order of seconds or minutes the answer might be no, it's probably better to include accurate diffraction calculations. If however the GA calculation is on the order of hours or days it would be better to exclude the high resolution diffraction calculations during the evolutionary optimisation process but then include them when the final optimised design is evaluated. Given a limited amount of computer time, it's better spent on higher GA population and generation counts than it is on diffraction calculation accuracy. To conclude, the heuristic trade-off sacrificing accuracy for speed in a GA is a safe one and is, indeed, worthwhile.

8. ACKNOWLEDGEMENTS

The author would like to thank Henrik Möller, Eckhard Kahle and Jean-Dominique Polack for their helpful comments and suggestions on an earlier version of this work.

9. REFERENCES

[1] J. O'Keefe, "Genetic Algorithm Fitness Functions for acoustic reflectors in performing arts venues" in *Proc. of the 24th International Congress on Acoustics,*(Seoul, Korea) 2022

[2] J. O'Keefe, "Zone to zone reflector optimisation using a genetic algorithm" in *Proc. of Institute of Acoustics;* (Athens, Greece) 2023

[3] J.H. Rindel, "Acoustic Design of Reflectors in Auditoria" in *Proc. of the Institute of Acoustics,* (Birmingham, UK), Vol. 14, Part 2, 1992.

[4] J, Pätynen, T. Lokki, "Directivities of Symphony Orchestra Instruments", *Acta Acustica United with Acustica*, Vol. 96, pp.138 – 167, 2010

[5] K. Deb, A. Pratap, S. Agarwal, T. Meyarivan, "A fast and elitist multi-objective genetic algorithm: NSGA-II" *IEEE Transactions on Evolutionary Computation*, Vol.6, no. 2, p.182, 2002,

[6] M. Laumanns, L. Thiele, E. Zitzler "An Adaptive Scheme to Generate the Pareto Front Based on the Epsilon-Constraint Method, Practical Approaches to Multi-Objective Optimization", in *Dagstuhl Seminar Proceedings*, 2005.

[7] J. Meyer, "Acoustics and the Performance of Music, 5th Edition", *Springer*, Bransweig, Germany, p. 230

[8] A. Szela, M. Lewin, J. Rubacha, "Verification of mathematical formulae applied to overhead stage canopy design", *Applied Acoustics*, Vol. 110, pp. 61-71, 2016.

[9] G. Naylor, "Level Difference and Mutual Hearing Amongst Musicians", *Proc. of the 12th International Congress on Acoustics*, Paper K5-7, Toronto, Canada, 1985,

[10] J. Dammerud, "Stage Acoustics for Symphony Orchestras In Concert Halls", *Ph.D. Thesis*, University of Bath (UK), 2009

[11] R. Wenmaekers, C. Hak, M. Hornikx, A. Kohlrausch, "Methods to measure stage acoustic parameters: Overview and future research", *J. Acoustical Soc. of America,* Vol. 141, no. 5, p. 3498, Boston 2017