

METRIC BASED DEVELOPMENT OF ACOUSTIC LENSES AND WAVEGUIDES

LT MacDonald
JA Ocle-Brown
JA Hargreaves

University of Salford, Manchester, M5 4WT
GP Acoustics (UK) Ltd, Maidstone, ME15 6QP
University of Salford, Manchester, M5 4WT

1 INTRODUCTION

There are many situations where it is desirable to be able to effectively sculpt sound waves or convert them from one shape to another. For example, professional audio line array loudspeakers often include waveguides intended to receive plane waves from one or more compression drivers at circular apertures and produce near plane waves at the waveguide mouth, which is often slot-like. The acoustical design of such wave shaping waveguides is complex. Achieving high efficiency and precise directivity control, while also minimizing acoustical resonance and response irregularity, is extremely challenging. But this complex design problem has considerable application in the audio industry, and many contrasting approaches are seen in different products^{1,2}.

A recent paper by Dodd and Ocle-Brown³ describes a new class of wave-shaping waveguides. Their approach is applicable to waveguides that are thin in one dimension. This feature allows corrugations and thickness variations to be applied to the waveguide geometry, which can modify the acoustical waves propagating through the device. It also minimises acoustic resonances and encourages coherent wave propagation over a wide bandwidth. Dodd and Ocle-Brown provide an overview of the steps required to construct such devices and show three example waveguides where the required corrugations have been manually calculated; a process they described as “time-consuming and intricate”.

This paper investigates an approach for automatically calculating the required corrugations and thickness corrections based on numerical analysis of an initial smooth “prototype” waveguide. The approach is based on two numerically calculated metrics that characterize the acoustical path-length and effective area continuously throughout the prototype waveguide.

The main research questions this paper investigates are:

- How effective are the metrics at characterising problematic regions within a waveguide?
- How much correlation is there between the error shown by the metrics and the acoustic performance of a waveguide?
- How useful are the metrics as guides for modifying a waveguides geometry to make it more capable of supporting coherent wave propagation?
- How well does the pathlength equalisation technique outlined in ref. 3 work for reducing differences in the relative pathlength through a device?

Section 2 outlines core mathematical theory that the techniques and analysis is based on, including 1P waves, Laplace's equation, and the waveguide metrics. Section 3 discusses the methodology for calculating the waveguide metrics and tracing streamlines through an FEA mesh. Section 4 presents simulations of the metrics on a test domain and discusses their meaning, then section 5 explores correcting the relative pathlength by distorting the 2D mesh to form corrugations. Section 6 extends the model to 3D and studies the acoustic performance with and without pathlength and area corrections. Finally, section 7 draws conclusions on the research questions and methodology used.

2 THEORY

This section reviews the core concepts and mathematical theory on which the proposed technique is based, such as 1P waves, Laplace's solution, and the waveguide metrics.

2.1 1P Waves

An acoustic field satisfying single parameter (1P) wave propagation is one in which the pressure depends on only a single spatial coordinate, which also happens to be the distance the wave has travelled. This means that there is equal pressure magnitude across each entire wavefront. A waveguide that supports 1P wave propagation should be free of acoustic resonances and diffraction, allowing coherent wave propagation.

Geddes⁴ explains that "If the wave fronts are not 1P, then the waves, as they propagate down the horn, will not move parallel to the bounding contours but instead will be forced to either intersect the waveguide contour, reflecting off the walls as they travel; or cease to remain in contact with the walls, propagating freely into space, unconstrained and uncontrolled".

Through rigorous mathematical analysis, Putland⁵ defined the necessary conditions for 1P wave propagation. He showed that the only shapes which support exact 1P wave propagation are the cartesian, cylindrical, and spherical coordinates systems, which create planer, cylindrical, and spherical wave shapes respectively, and that no new 1P horn geometries remain to be discovered.

However, if a waveguide can be designed to support close to 1P wave propagation over a range of frequencies, it will have predictable organised wavefronts in this bandwidth. This leads to controlled directivity and good coupling capabilities as there are no cross modes or diffraction³. For these reasons, a geometry that produces more organised 1P-like wave propagation over a wide range of frequencies is assumed to have 'better' performance for the purposes of this paper.

2.2 Stretch and Flare Metrics

Laplace's equation, sometimes referred to as the heat equation, is a second order partial differential equation that is commonly used to simulate steady state heat conduction. Laplace's equation is:

$$\nabla^2 u = 0. \quad (1)$$

Laplace's equation can be used to model an equilibrium situation that is non-oscillatory. Laplace's equation is also a special case of Helmholtz' equation with wavenumber $k = 0$. The solution will depend on the geometry and the boundary conditions. The Dirichlet boundary condition, which is used in the simulations in this paper, specifies the value of u on a boundary.

We propose an extension to Putland's work that arises through applying it to solutions of Laplace's equation. This is a reasonable approach because 1P solutions of the Laplace and Helmholtz equations have closely related forms. Furthermore, we propose two metrics – 'stretch' and 'flare' – that quantify deviation of a solution u from Putland's conditions for 1P waves.

The stretch metric s is related to the relative pathlength through a domain. Non-unity in s indicates regions of a domain which will be poor at supporting 1P wave propagation as the relative pathlength through the domain is too short or too long. The 'stretch' metric s is defined mathematically as:

$$s = |\nabla \xi|, \quad (2)$$

where ξ (defined by Putland⁵) is a coordinate computed from u that measures cumulative path length. This is measured along a path that follows ∇u , so therefore runs perpendicular to planes of constant u . $s \neq 1$ means the surfaces of constant ξ are either too close together or too far apart compared to the norm. The mapping between u and ξ is found through integration:

$$\xi(u) = \int \frac{du}{|\nabla u|} \quad (3)$$

Using Putland's conditions for 1P waves, it can be shown that the flare rate term from Webster's horn equation⁶ is equal to:

$$f = \frac{1}{S} \frac{dS}{d\xi} = - \frac{\mathbf{e}_u \cdot \nabla |\nabla u|}{|\nabla u|}. \quad (4)$$

Here \mathbf{e}_u is the unit vector in the direction of ∇u . The flare rate indicates how the flare ($dS/d\xi$) changes with cross sectional area. Another identity for the flare metric can be found using logarithmic differentiation:

$$f = \frac{1}{S} \frac{dS}{d\xi} = \frac{d}{d\xi} \ln S \quad (5)$$

By rearranging the right-hand side of eq. (5) the cross-sectional area through the domain can be calculated to be:

$$a_f = e^{\int f d\xi}. \quad (6)$$

Because this involves integrating f down a ξ path, a_f can be thought of as the comparative change in cross-sectional area (relative to the start geometry) that a 1P wave *feels* as it propagates through a domain. Notably, it can differ from the physical change in cross sectional area through a domain. In this paper, a_f is therefore referred to as the 'felt area'.

3 METHODOLOGY

To calculate the waveguide metrics over an arbitrary shaped domain, the FEA package PAFEC⁷ was used to compute the solution to Laplace's equation, u . The solver allowed calculations of u – and therefore the metrics – on 2D meshes, 3D meshes, and shell meshes. The latter of these can be useful in situations where a thin 3D geometry can be well approximated by using an infinitely thin shell of 2D elements in a 3D space. Dirichlet boundary conditions with values $u = 0$ and $u = 1$ were used to represent the entrance and exit of the waveguide respectively. Thus, the Laplace solution climbs from $u = 0$ at the entrance of the waveguide to $u = 1$ at the exit. The nodal u data was then imported to MATLAB where it was used in the computation of the metrics s , f , and a_f .

To find paths on which ξ could be computed, streamlines of steepest ascent through u were traced on the FEA mesh using a Runge-Kutta 4 (RK4) method modified to work in three dimensions. The general concept of tracing streamlines is to iteratively step through a domain using the vector quantity sampled at each step to direct the algorithm towards its next sampling position, similar to how ordinary differential equations are sometimes solved numerically. On each streamline tracing step, a Newton-Raphson numerical technique⁸ was used to inverse map the global coordinates to the local element coordinates – this was chosen because it is a flexible technique that works well with complex high order elements. The nodal ∇u data was then interpolated to the RK4 sampling location using the element shape functions so the step could be traced and the next start point found.

The streamlines allow integration with respect to ξ , as is required to compute a_f (see eq. (6)). The integration was performed starting from the entrance using the cumulative trapezoid rule. Data was then mapped between the nodes and streamline points using a global shape function matrix.

To evaluate the acoustic performance of the 3D domain, the Helmholtz wave equation was simulated in PAFEC using FEA. These simulations used did not include any nonlinearities or viscous losses,




which could influence the performance of a real waveguide, particularly at high SPL or if the geometry is extremely thin.

4 UNDERSTANDING THE METRICS

The metrics in section 2 measure how much a Laplace solution computed for a specified geometry deviates from Putland’s 1P conditions. They therefore indicate how effective regions of a domain are at supporting 1P wave propagation. This section demonstrates the waveguide metrics on a test case geometry and shows how streamlines of the vector field ∇u are used to calculate them.

To make this section clearer to the reader, different colour scales have been used to plot the three waveguide metrics. Table 1 shows which colours represent which quantities.

Table 1: Waveguide metric colour guide

	u – solution to Laplace's equation
	s – stretch
	a_f – felt area

4.1 Test geometry

For this paper a single test case geometry is examined. This is termed the “sine-curve” and it was chosen to help illustrate the metrics and waveguide correction technique. However, the same process could be applied to many different shapes and designs. The test geometry was constructed by adding 8cm of width to a single period (50cm) of a sinusoidal curve with 5cm amplitude.




Figure 1: Solution to Laplace’s equation u on sine-curve geometry with 50 isolines of constant u overlaid

Figure 1 shows the solution to Laplace’s equation on the test shape with 50 isolines of constant u overlaid. The geometry has no expansion in area from the entrance to exit, so no change in a_f is expected. The shape can be thought of as two connected bends which curve in opposite directions. Each bend causes a change in the relative pathlength through the domain, but at opposing sides, which makes it a good test case to demonstrate the stretch metric s .

The entrance is on the boundary near $x = 0$ and the exit is on the boundary near $x = 0.5$. At these $u = 0$ and $u = 1$ respectively. The distance between the isolines is related to ∇u and illustrates how there is a steeper (higher) ∇u close to the inside edge of each curve.




Figure 2: Curvature κ of the central line through the geometry

Figure 2 shows the curvature of the central line through the geometry. The curvature is at a maximum at the peak and trough of the sinusoid shape ($x = 0.125, 0.375$ (m)), and goes to zero in the middle, start, and end. It shall be seen that the curvature of a channel of constant width has an interesting effect on the waveguide metrics.

4.2 Stretch

Stretch s measures the gradient of u . In a Helmholtz problem, this measures relatively how fast a part of the wave has to travel in order to keep up with the rest of the wavefront it is part of. Waves at the outside of a bend would somehow have to travel faster than those at the inside of a bend if they are all to emerge together. But for the Laplace problem its meaning is rather arbitrary – the main thing that defines it is the length of the duct. Because of this, relative stretch compared to a reference path is instead considered. Once a reference path through the domain is chosen, the relationship of arc length ξ to the solution to Laplace’s equation u can be used to analyse the rest of the geometry. The rest of the domain is compared to the reference ξ path, so depending on the ξ path chosen different values of s will be obtained, which may be greater or less than one depending on whether the path is shorter or longer (giving steeper or less steep gradients in u respectively).




Figure 3: Streamlines of the vector field ∇u traced through the sine-curve geometry

Streamlines that trace the steepest ascent of ∇u were traced through the domain, starting on the entrance boundary where $u = 0$. Because the streamlines follow ∇u they are ξ paths following Putland’s conditions⁵. Figure 3 shows 10 streamlines of ∇u traced through the sine-curve geometry.

By definition, a ξ path follows ∇u so is always perpendicular to surfaces of constant u . This means that a rigid edge boundary is always a ξ path since it is parallel to ∇u . A boundary edge therefore makes a straightforward choice of reference path. The arc length ξ along the reference path is found by integration, then this value was mapped to all nodes in the domain with matching u (since Putland defined path-length $\xi(u)$ to be a function of u) and the stretch s metric was calculated as $|\nabla \xi|$.




Figure 4: stretch metric s on sine-curve geometry

In Figure 4 the bottom boundary was used as the reference ξ path to calculate s . Therefore, on this boundary $s = 1$, indicating no error in relative pathlength. For the first half of the curve the reference edge is on the inside edge of the curve, so the longer opposing boundary has $s > 1$, indicating how a wavefront there would need to travel faster to ‘keep up’. Then on the second curve the channel bends in the other direction so $s < 1$, indicating that the relative pathlengths are shorter than the reference edge. s also highlights how the areas with biggest path length difference are around the middle of each bend where the curvature is highest (see Figure 2). This makes sense because it is also where the channel turns the tightest corner.

It is far easier to add pathlength to a waveguide than it is to take it away. Ref. 3 gives an example of how to do this. In Figure 4 the flat domain could be distorted somehow to add pathlength when $s < 1$, but a problem arises when the domain is too long, i.e. when $s > 1$, since no mechanism exists to take pathlength away. In other words, the stretch metric is most useful when it shows how much shorter part of the domain is compared to the longest part at that position, as the domain can then be modified to add pathlength to these regions to compensate. A method of normalisation is therefore required that uses the longest part of the domain at each point as the reference ξ path.

To achieve the required normalisation, the stretch metric was calculated multiple times using each of the traced rays as the reference ξ path, and then taking the minimum calculated stretch value at each node. This yielded a minimum-normalised stretch metric s_n . This method of normalisation means that an abstract reference ξ path is constructed, which comprises sections of whichever streamlines have the longest relative pathlength (or biggest $d\xi/du$) in each region of the domain.

Figure 5 shows normalised stretch s_n on the sine-curve geometry. A maximum value of one is found on the outside edge of each curved section. This is where the reference path has fallen since it has the longest relative path length. This is a useful result because pathlength can be added to the domain in areas where $s < 1$. It is interesting to note how $s_n \rightarrow 1$ in the middle at the intersection between the two curved sections.




Figure 5: Minimum-normalised stretch s_n on sine curve geometry

Since the normalisation acts as a scaling, the ratio of the maximum and minimum values of stretch on a surface normal to the ∇u (a wavefront) will always be the same regardless of it. For example, comparing s on Figure 4 and Figure 5, the ratio of s between the top to bottom boundaries at $x = 0.13$ is around 1.78 in both cases.

4.3 Felt area a_f – the area metric

The relative change in cross sectional area or felt area a_f was calculated from flare f using Eq. (6). The integration was performed down the streamlines in the same direction as ∇u , therefore the entrance surface (with Dirichlet boundary condition of $u = 0$) always has $a_f = 1$ and a_f indicates how the relative cross-sectional area changes through the geometry from entrance to exit.




Figure 6: felt area a_f on sine-curve geometry

Figure 6 shows a_f on the sine curve geometry. Interestingly, the change in curvature through the shape causes a_f to indicate a decrease in area close to the inside edge and an increase close to the outside edge of each bend. The domain has a constant width, so the physical cross-sectional area does not change - it suggests that for a wave travelling through a duct of constant width, the change in curvature causes an apparent change in area which varies across the width of the channel.

5 EQUALISING THE RELATIVE PATHLENGTH

In a thin channel, one method of equalising the relative pathlength difference shown by s_n is to distort the geometry in areas where the domain is too short ($s_n < 1$). This concept of distorting a thin domain to equalise the relative pathlength is outlined in ref. 3. This section automates that approach and uses Vol. 44. Pt. 3. 2022

corrugations that are a function of s_n to equalize the relative pathlength; the process should scale arc length by a factor of $1/s_n$. Sinusoidally shaped corrugations were used, but other profiles are possible.

The corrugations have two parameters: amplitude A and repeat length L . The relationship between these and added arc length is not linear. The left-hand side of Figure 7 show how sinusoidal curves with a smaller period require less amplitude for the same extension. These curves all have an arc length of 7.64, but differing choice of L ($2\pi/3, \pi, 2\pi$) has meant differing amplitudes A achieve this.

The right-hand side shows the relationship between extension factor ($1/s_n$), corrugation amplitude (A) and period (L). This non-linear function is used to find the required amplitude A from s_n . This allows tailored corrugations to be added throughout a domain which have the correct amplitude – and therefore arclength – to equalise pathlength through the geometry.




Figure 7: left: corrugation curves all with an arc length of 7.64, right: non-linear relationship between extension ($1/s_n$) and corrugation amplitude A

Figure 8 shows a shell mesh of the test geometry where sinusoidal corrugations have been used to correct relative pathlength. The corrugation amplitude A was computed from s_n (shown in Figure 5) using the relationship in Figure 7. This causes the largest corrugations to be close to the inside boundary of each bend where s_n is lowest, and the channel curvature is highest.

u was used as the coordinate system with which to build the periodic corrugations and move through the domain from entrance to exit. This ensured that the corrugations were always perpendicular to the edges of the domain and parallel with surfaces of constant u . Using u as the coordinate system in which to build the periodic corrugations means that the spatial period (wavelength) of each corrugation changes depending on s_n . This effect was accounted for when calculating the s_n to amplitude relationship.

In the example in Figure 8, a small stretch offset $s_o = 0.01$ was subtracted from s_n before building the corrugations. This adds a small amount of corrugation everywhere and helps avoid having heavily distorted elements close to the outside edges of each bend, where the corrugation amplitude would otherwise fade to zero.




Figure 8: sine-curve geometry with corrugations to equalise the relative pathlength

To quantify whether the corrugations have helped equalise pathlength, the new geometry was re-analysed. But instead of using stretch normalised to the longest ξ path, the average normalised stretch s_a was used; this is the ratio of the pathlength to the mean pathlength through the domain. It highlights regions where too much or too little pathlength has been added by the corrugations.




Figure 9: average normalised stretch s_a on the corrugation-corrected sine curve geometry

Figure 9 shows s_a on the corrugated shell mesh. s_a is much closer to unity on the corrugated shell mesh than on the original 2D flat domain in Figure 4 (note different colour limits), which indicates that the differences in the relative pathlength through the geometry have been reduced and the design will therefore be better at supporting 1P wave propagation. The most intense spots in s_a are at the peaks and troughs of the largest corrugations. There are also periodic patterns in s_a because sine wave shape corrugations add pathlength in a cyclic manner.




Figure 10: felt area a_f on the corrugation-corrected sine curve geometry

Figure 10 shows the effect of pathlength equalisation on the felt area metric a_f . Interestingly, the range of a_f is significantly less than in the non-corrugated geometry (Figure 6), which shows that correcting the relative pathlength through a domain also reduces the change in area indicated by a_f . It is worth noting that because the elements are infinitely thin on a shell mesh, the metrics in Figure 9 and Figure 10 only analyse the domain parallel to the surface and not perpendicular to it.

6 THREE-DIMENSIONAL GEOMETRY

Unlike an infinitely thin 2D planer or shell domain, a 3D domain has a third dimension (its thickness) that can be used to adjust the cross-sectional area and correct changes indicated by the felt area metric a_f . This is referred to as area correction. This chapter investigates the acoustic performance of the geometry with and without pathlength and area correction.

Converting a shell mesh into a 3D mesh was achieved by computing vectors normal to the surface at each node and then lofting equally in the positive and negative direction to extrude into a 3D domain. This means the central surface is the same as the original shell mesh. Figure 11 shows the corrugated sine-curve geometry extruded into 3D with a base thickness t_b of 5mm.




Figure 11: 3D corrugated sine-curve geometry

When extruding into 3D, the base thickness t_b could also be adjusted spatially by using a_f from the 2D model (Figure 6) to get a corrected thickness t_c .

$$t_c = \frac{t_b}{a_f} \quad (7)$$

The effect on the acoustic performance with and without this area correction is explored in Figure 13.

To understand the effectiveness of the pathlength and area correction techniques, FEA simulations of the Helmholtz wave equation were performed on the 3D geometry. The domain was excited by a constant velocity over the entrance surface and was terminated with a $\rho_0 c$ boundary condition on the exit surface.

Figure 12 shows the pressure distribution at 20kHz, with and without pathlength and area correction. On the original flat geometry, the bends cause the pressure to intensify close to the outside edges, and the pressure distribution at the exit ($x = 0.5$) is visibly uneven. In contrast the model with pathlength and area correction has much more aligned wavefronts throughout the geometry. However, there are slight fluctuations in the pressure visible across some of the wavefronts.

The maximum range of the SPL over all the nodes on the exit surface of the geometry was used to characterise the acoustic performance over a wide range of frequencies. If the range of the SPL is zero it implies that the pressure wave at the mouth is completely uniform and coherent and therefore no cross-modes have been excited.

Figure 13 shows the range of SPL across the mouth with different correction techniques. Frequencies lower than 2kHz have not been included as the larger wavelengths mean the uncorrected design already has close to plane wave propagation.

The two models with pathlength correction perform significantly better than the original flat geometry over most frequencies. The geometry with both pathlength and area correction performs best, with a maximum range in SPL of around 4dB at 5kHz. These results suggest that the best technique is to correct both the pathlength and the area (s_n and a_f) together.




Figure 12: acoustic pressure distribution at 20kHz, top: original flat geometry, bottom: geometry with pathlength and area correction

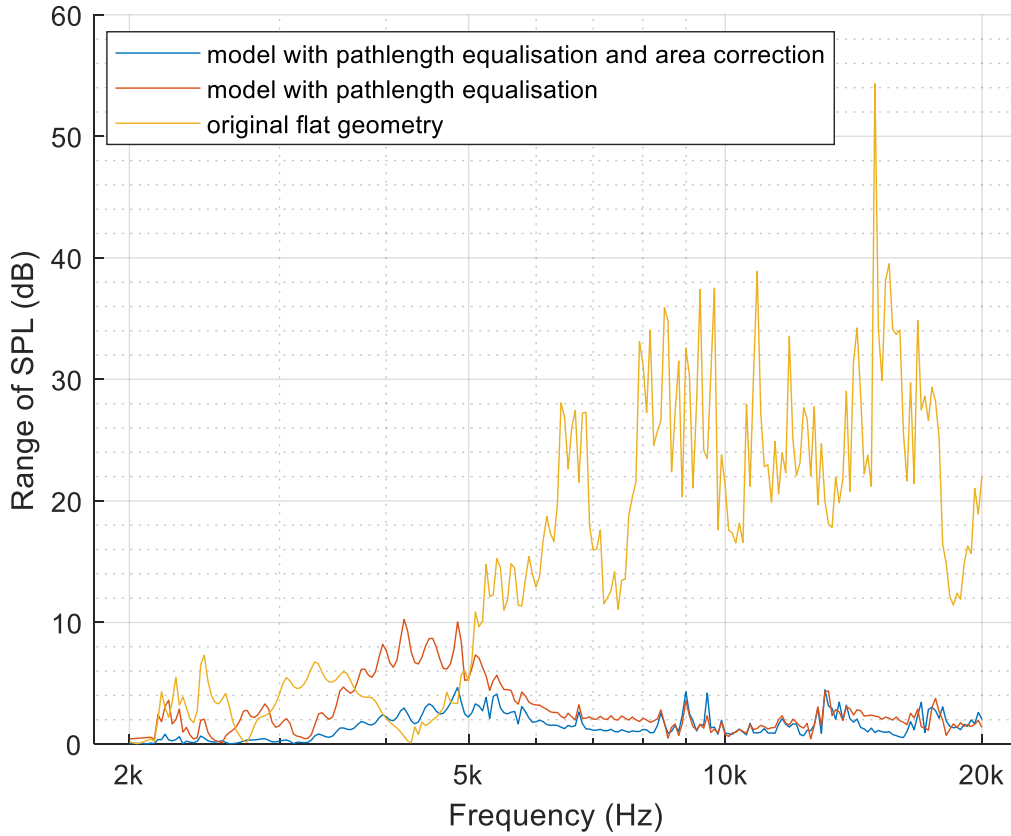


Figure 13: range of SPL across exit surface with and without pathlength and area correction
Vol. 44, Pt. 3, 2022

7 DISCUSSION AND CONCLUSIONS

This paper has tested the hypothesis that metrics based on Laplace's equation can be used to modify waveguides to better support 1P acoustic waves. Results show a geometry with judiciously added corrugations to have significantly less pathlength discrepancy than the original planar geometry. Acoustic simulations back this up, showing wavefronts that are much more aligned and a significantly smaller range of SPL over the mouth, indicating the geometry better supports 1P wave propagation. Hence, this hypothesis appears to hold for the test case simulated, suggesting the corrugation-based pathlength correction technique³ can be a useful tool for refining designs of acoustic waveguides.

Although the performance of the waveguide is significantly improved, it is not perfect. Fundamentally, the limitation on the performance of the modified waveguides is related to the findings by Putland⁵, in which he describes mathematically how there are only three coordinate systems which will support perfect single parameter waves. This is important as the prototype geometric surfaces studied in this project are not of these shapes, so will never support 1P wave propagation perfectly for all frequencies no matter how well the correction works (except for the unlikely scenario where the correction procedure creates one of Putland's isolated cases). However, through reduction of non-unity in s_a and a_f , the designs become closer to supporting 1P wave propagation over a limited bandwidth.

The extend of this bandwidth is limited by the accuracy of the pathlength correction (and therefore reduction in non-unity of s_n), and the thickness of the thin corrugated domain, which should be less than one wavelength to avoid reflections of the propagating waves from the boundaries

7.1 Further Research

While the test case geometry used in this paper is useful for illustrating and investigating the metrics and correction technique, the methods could be used to design waveguides or acoustic lenses for many different situations. It would be interesting to tackle some more complex real-world acoustic wavefront manipulation problems, such as the circle to slot adapter used in line array sources.

8 CREDIT AUTHOR STATEMENT

Lewis MacDonald: Methodology, Software, Investigation, Visualisation, Writing - Original Draft.
Jack Ocle-Brown: Conceptualisation (metrics), Resources, Supervision, Writing – Review & Editing.
Jonathan Hargreaves: Methodology, Supervision, Writing – Review & Editing.

9 REFERENCES

1. C. Heil, 'Sound wave guide', US5163167A, Nov. 10, 1992
2. J. Spillmann and S. Riemersma, 'Acoustic Waveguide', US9571923B2, Feb. 14, 2017
3. M. Dodd and J. Ocle-Brown, 'Wave-shaping using novel single-parameter waveguides', presented at the Audio Engineering Society Convention 153, Oct. 2022
4. E. Geddes, Acoustic Waveguide Theory. J. Audio Engineering Soc. Vol. 37, No 7/8 p554. (1989).
5. G.R. Putland, Every One-Parameter Acoustic Field Obeys Webster's Horn Equation. J. Audio Engineering Soc, Vol. 41, No. 6. (1993).
6. A. Webster, Acoustical Impedance and the Theory of Horns and the phonograph, Proc. Natl. Acad. Sci. USA, 5, pp. 275–282; reprinted in J. Audio Engineering Soc., 25 (1977), pp. 24–28, (1920).
7. PACSYS Ltd., PAFEC-FE. PACSYS Ltd., 2008.
8. G. Silva, R. Le Riche, J. Molimard, and A. Vautrin. Exact and efficient interpolation using finite elements shape functions. (2007). HAL <https://hal.archives-ouvertes.fr/hal-00122640v2>