Proceedings of IASS-IWSS 2026 – September 14-18, 2026, Turin, Italy
"R-Evolution of Shapes: Sustainability, Re-Use, and New Design Paradigms"
Annual Symposium of the International Association for Shell and Spatial Structures
3rd Italian Workshop on Shell and Spatial Structures
A. Manuello Bertetto, S. Gabriele, F. Marmo, A. Micheletti (eds.)

Coupling Structural Form Finding and Ray-Traced Acoustic Optimisation for Compression-Dominant Structures

Giacomo Montiani*, Alessandro Dell'Endice [a], Achilleas Xydis [b], Tom Van Mele [a], Philippe Block [a]

*ETH Zürich, Institute of Technology in Architecture, Block Research Group
Stefano-Franscini-Platz 1, 8049 Zurich, Switzerland
montiani@arch.ethz.ch

[a] ETH Zürich, Institute of Technology in Architecture, Block Research Group
[b] ETH Zürich, Institute of Technology in Architecture, Gramazio Kohler Research

Abstract

Compression-dominant structures, such as vaults and shells, achieve material efficiency by aligning their geometry with the flow of funicular forces under prescribed loading conditions. This geometry-driven structural behaviour can be controlled through advanced form-finding methods and tools; however, the acoustic performance of these structures remains comparatively under-investigated. Vaulted spaces are acoustically challenging and are often characterised by poor acoustic quality due to their inherent concave shape, highly reflective, non-porous materials, and large spans. These conditions produce uneven sound distribution, degrading speech intelligibility. Since these acoustic deficiencies are governed by the same geometric parameters that determine structural behaviour, structural and acoustic requirements are inherently coupled. This paper introduces a computational workflow that integrates Thrust Network Analysis with a custom ray tracer based on bidirectional scattering coefficients to describe how structural form-finding influences acoustic performance and to support acoustically informed vault design. The pipeline is implemented within the COMPAS framework and applied to a barrel-vault classroom case study.

Keywords: Compression-dominant structures, Form-finding, Architectural acoustics, Acoustic simulation, Performance-based design, Ray-tracing

1. Introduction

Compression-dominant structures, such as vaults and shells, are designed so that applied loads are transferred to supports almost exclusively through compressive forces, minimising bending stresses [1]. Their capacity to span large distances with minimal material use makes them particularly relevant in contemporary sustainable design, where reducing material use and embodied CO_2 emissions is a central concern [2].

The structural behaviour of these systems has been extensively investigated, and structural form-finding enables the computational determination of optimal compressive geometries [3]. A considerable number of established tools are now available for form-finding and manufacturing optimisation [4, 5].

While these efforts have significantly advanced the understanding and design of compression-dominant forms, their acoustic behaviour has received comparatively less attention within the form-finding process itself. Although previous work has coupled structural form-finding with acoustic simulations [6], existing optimisation tools still tend to treat structural performance, manufacturing constraints, and acoustic behaviour as separate design problems rather than as a single coupled one.

The consequences of this separation are not trivial: the defining morphological features of compression-dominant structures have direct, often detrimental, effects on room acoustics. Their

Copyright © 2026 by <insert name(s) of the author(s) as listed above>
Published by the International Association for Shell and Spatial Structures (IASS) with permission.

concave geometry, combined with highly reflective interior surfaces and materials that often exhibit low scattering properties, leads to significant room acoustic problems [7]. These defects include, in particular, an uneven distribution of sound energy, which means that sound concentrates in focal points, while other areas become dead spots [8].

This paper proposes a workflow that couples structural form-finding with acoustic ray-tracing analysis to identify which form-finding parameters can be adjusted to mitigate the acoustic problems of compression-dominant structures.

1.1. Funicular form-finding for compression-dominant structures

Geometric stiffness form-finding methods determine a structure's equilibrium based solely on its geometry and applied loads, independent of material properties. This makes them particularly suitable for early design stages, where the focus is on shaping the structure rather than detailing its material composition [9].

Among these methods, Thrust Network Analysis (TNA), introduced by Block and Ochsendorf [10], extends the principles of graphic statics and thrust-line analysis to three-dimensional problems. TNA describes the equilibrium of a three-dimensional compression-only network using a pair of reciprocal diagrams: a form diagram Γ, representing the network's horizontal projection, and a force diagram Γ^*, representing its horizontal equilibrium. Manipulating either diagram allows the designer to explore a wide range of compression-only geometries while preserving structural feasibility (Figure 1).

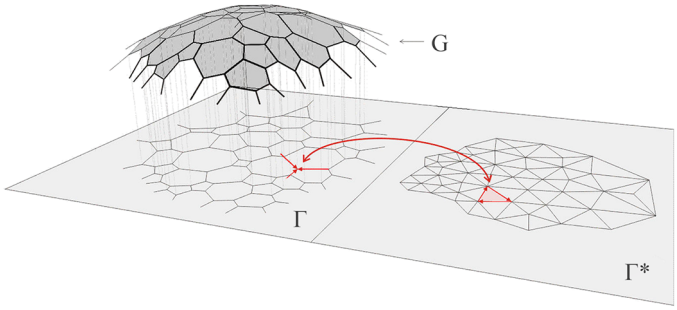

Figure 1: Relationship between compression shell (G), its planar projection (form diagram Γ) and the reciprocal diagram (force diagram Γ^*) to determine equilibrium (from Block and Ochsendorf [10]).

This flexibility is the main reason TNA is adopted as the form-finding method in the workflow presented here: it allows the vault geometry to be modified via structurally meaningful parameters while ensuring that every generated shape is stable under compression.

1.2. Geometrical acoustics and ray tracing

Geometrical acoustics methods approximate sound propagation by replacing wavefronts with rays carrying energy, an assumption valid in the mid-to-high-frequency range, where wave phenomena become less relevant [11]. Among these methods, ray tracing is one of the most widely used for architectural acoustic simulation.

In ray-tracing, a number of rays is emitted from a source. Every time a sound ray hits a surface, a portion of the ray's energy is absorbed according to the surface's absorption coefficient α, while the remainder is divided between specular and diffuse components according to the scattering coefficient (Figure 2) [12]:

$$E_{scattered} = E_{incident} \cdot (1 - \alpha) \cdot s \tag{1}$$

$$E_{specular} = E_{incident} \cdot (1 - \alpha) \cdot (1 - s) \tag{2}$$

The diffuse component is typically distributed according to Lambert's cosine law, in which the scattered energy depends on the angle of reflection relative to the surface normal [13]. The energy carried by each ray along its path is further attenuated by distance and air absorption.

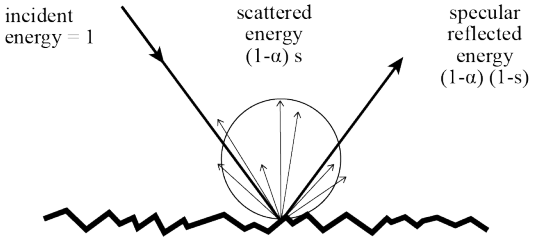

Figure 2: Sound scattering from a rough surface illustrating the separation of reflected energy into scattered and specular components (adapted from Vorländer and Mommertz [12])

A known limitation of Lambertian or random-incidence scattering models is that they collapse directional dependence and, therefore, do not capture the bidirectional behaviour of scattered fields from rough or structured surfaces. To address this, Heimes and Vorländer proposed bidirectional surface scattering coefficients, in which the scattered energy depends on both incident and reflection directions, analogous to the Bidirectional Reflectance Distribution Function (BRDF) used in optics. [14].

1.3. Related work

Several studies have explored the coupling of acoustic analysis with computational design strategies to inform the global geometry of architectural spaces. Foged et al. [15] combined genetic algorithms, parametric modelling, and acoustic simulation to shape the global geometry of performative spaces with respect to acoustic objectives. More directly relevant to the present work, Méndez Echenagucia and Block [6] proposed a method that employs a multi-objective genetic algorithm (MOGA) to couple Thrust Network Analysis (TNA) with the Rhinoceros plugin Pachyderm Acoustics for geometric acoustics simulation, optimising the acoustic behaviour of funicular shell geometries. Using the TNA model's force densities as design variables, all candidate solutions generated by the MOGA are compression-only by construction, allowing the genetic search to operate solely on acoustic objective functions.

The workflow proposed in this paper builds on the same coupling between TNA and ray-tracing simulation but differs in three respects. First, it employs a custom ray tracer that tracks the spatial distribution of speech-supporting early sound energy using the bidirectional scattering coefficient formulation, a feature not available in current acoustic simulation tools for architectural design. Second, rather than optimising aggregate acoustic indices, the workflow evaluates the spatial distribution of early sound energy on a listener-plane grid and traces back the surface regions that contribute to local concentrations, enabling targeted geometric modifications, such as ribs, in specific portions of the vault. Third, the pipeline is implemented entirely in Python within the COMPAS framework, making it software and operating system agnostic.

2. Methodology

This section describes the proposed workflow. The workflow couples Thrust Network Analysis (TNA) with a custom acoustic ray tracer in an iterative loop, in which the spatial distribution of sound energy from a given vault geometry informs the application of shape-modifying functions and the adjustment of their form-finding parameters. Each iteration is structured in four steps: (i) generation of a compression-dominant vault geometry through TNA-based form-finding, together with the definition of the room's boundaries; (ii) acoustic analysis based on ray-tracing, evaluating early sound energy distribution on a listener-plane grid; (iii) application of shape-modifying functions that, based on the analysis of ray paths, locally adjust the form diagram; (iv) acoustic simulation in Pyroomacoustics to assess the main acoustic parameters for speech intelligibility of the best iteration solutions.

2.1. TNA-based form-finding

The first step in the workflow generates the vault geometry that defines the room's ceiling using TNA. The designer specifies the vault typology by generating a form diagram and defining the corresponding boundary conditions. The procedure is implemented through the compas_tna package [16]. In the current work, the form diagram is discretised into a regular grid of $\mathbf{50} \times \mathbf{50}$ cm cells. The geometry adopted is a barrel vault, and its edges are divided into two groups according to their structural role: load-carrying edges, which transfer vertical loads to the supports, and unloaded edges, which define the horizontal layout of the network. Each group is assigned its own vertical bounds $\mathbf{h}_{\min}$ and $\mathbf{h}_{\max}$, allowing the positions of load-carrying and unloaded nodes to be constrained independently in the three-dimensional network. Once the supports and constraints are defined, the force diagram is generated, the horizontal equilibrium is solved, and the vertical equilibrium is computed under the additional constraint of a maximum overall rise $\mathbf{z}_{\max}$. The resulting shape is therefore primarily controlled by the local bounds $\mathbf{h}_{\min}$, $\mathbf{h}_{\max}$ and the global parameter $\mathbf{z}_{\max}$.

2.2. Acoustic ray-tracing analysis

The designer sets the room geometry (height and boundary conditions) to match the design scenario under evaluation. Then, they assign absorption and bidirectional scattering coefficients to every surface in the room. Finally, they add a sound source by defining its position, directivity, number of rays, and order of reflection. The directivity is controlled by an opening angle, with 360° corresponding to an omnidirectional source. The room volume is closed at the top by the vault mesh produced by the form-finding stage.

Unlike the formulation in Eqs. 1 and 2, the reflected energy from each surface is not split between a single specular and a single diffuse component. For each scattering surface, the directional reflection behaviour is described by a BRDF discretised at a 5° angular resolution, yielding up to 1296 directional bins per incident angle (Figure 3a). The energy is distributed among n reflected rays, each carrying a fraction proportional to its normalised bidirectional scattering coefficient; the bin with the highest coefficient corresponds to the dominant reflection direction and carries the largest fraction of the reflected energy (Figure 3b). The energy carried by each propagating ray is further attenuated by distance and air absorption, computed according to ISO 9613-1 [17].

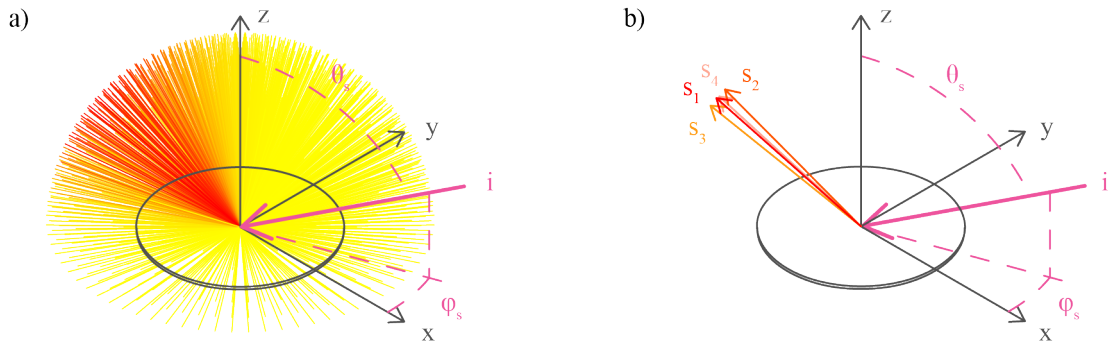

Figure 3: Visualisation of the BRDF in the custom ray-tracer. (a) Full BRDF of a flat surface at 500 Hz: 1296 directional bins per incident angle (i), with colour indicating the value of the bidirectional scattering coefficient (yellow: low; red: high). (b) Selection of the four strongest bins to be used in the ray-tracing simulation.

Sound energy is evaluated on a horizontal listener-plane grid positioned at listener height. For each cell, the ray-tracer accumulates the energy received within a configurable time window after the direct sound arrives, capturing the contribution of early reflections known to support speech intelligibility [18]. The result is presented as Sound Pressure Level (SPL) maps of the early-energy distribution across the listener plane, computed independently for each analysed frequency band (See Figure 4).

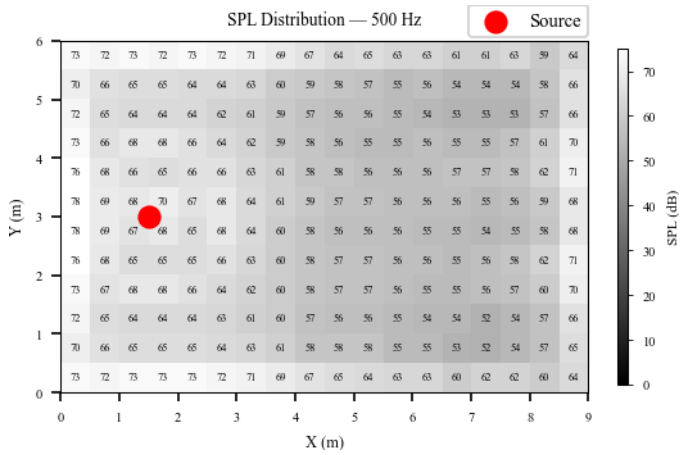

Figure 4: Example SPL map of the early-energy distribution across the listener plane at 500 Hz, showing per-cell SPL values [dB]. The red marker indicates the source position. Cells closer to the room boundaries (within 1.5 m) are excluded from the descriptor computation.

For each candidate geometry, two descriptors are computed per frequency band across the N listener-plane cells. The mean SPL $\bar{L}$ is first computed as:

$$\bar{L} = \frac{1}{N} \sum_{i=1}^{N} L_i \tag{3}$$

The first descriptor is the standard deviation σ of the SPL across the listener plane:

$$\sigma = \sqrt{\frac{1}{N} \sum_{i=1}^{N} (L_i - \bar{L})^2} \tag{4}$$

with lower values indicating a more homogeneous early-energy distribution. The second descriptor is the coverage ratio $C_{x\,\mathrm{dB}}$, defined as the percentage of listener-plane cells whose SPL falls within $\pm x$ dB of the field mean:

$$C_{x\,\mathrm{dB}} = \frac{1}{N} \left| \{ i : |L_i - \bar{L}| \leq x \} \right| \times 100 \tag{5}$$

where x is a tolerance threshold set by the designer. Higher values indicate broader, more uniform coverage across the audience area. Both descriptors are computed independently for each frequency band and averaged across bands to obtain an overall value.

2.3. Shape-modifying functions

Once the structural and acoustic boundary conditions are set, the workflow proceeds in two iterative phases. In the first phase, multiple form-finding runs are performed by varying the vertical bounds $\mathbf{h}_{\min}$ and $\mathbf{h}_{\max}$ assigned to the edge groups of the form diagram. The variation ranges are specified by the designer, producing a set of candidate vault geometries. For each candidate, the descriptors σ and $C_{x\,\mathrm{dB}}$ are evaluated from the early-energy distribution; the geometry with the best performance is retained as the starting point for further refinement.

In the second phase, shape-modifying functions are applied to redistribute the early energy more uniformly across the listener plane. The procedure first identifies the listener-plane cells corresponding to energy hot spots, traces back the first-order reflections contributing to those concentrations, and projects their reflection points onto the vault mesh. The regions of the mesh where these projections cluster

identify the portions of the vault responsible for the focusing effect, and their corresponding edges in the form diagram are selected for modification. All modifiers operate by locally adjusting the form diagram and rerunning the form-finding stage; the global parameter $\mathbf{z}_{\max}$ is fixed at this stage to preserve the overall vault rise established in the first phase. The modifiers can be applied individually or sequentially in an automated loop.

2.4. Acoustic parameters evaluation

At the end of each iteration, the best-performing geometries are also tested using an independent simulation with Pyroomacoustics [19], a Python library that combines ray-tracing with the image source method. The room model used in Pyroomacoustics matches the geometry, material assignments, and source configuration of the custom ray-tracing simulation. This benchmark allows the room impulse response (RIR) and standard acoustic parameters for speech — Clarity (C50), Reverberation Time (T60) and Speech Transmission Index (STI) — to be evaluated for each iteration. The comparison between the two simulators is qualitative rather than quantitative: the custom ray-tracer is limited to three reflections and focuses on the early-energy distribution, whereas Pyroomacoustics extends the simulation to capture the full impulse response.

3. Case study application

The workflow is applied to the design of a barrel vault covering a hypothetical classroom with a plan of 9 × 6 m and a height of 4 m. The absorption coefficients assigned to each interior surface are listed in Table 1; Figure 5 shows the surfaces they are assigned to. In the custom ray-tracer, each surface is assigned the bidirectional scattering coefficients of a geometrically flat surface to isolate the effect of the global vault shape on the early-energy distribution; in Pyroomacoustics, random-incidence scattering coefficients are used instead.

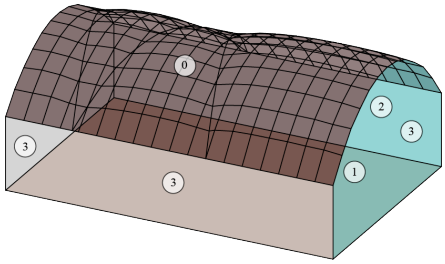

Figure 5: Interior surfaces of the room model and their assigned materials, numbered according to Table 1.

Table 1: Absorption coefficients of the model's surface materials at 250 Hz and 500 Hz.

Material	α **250 Hz**	α **500 Hz**
0. Rough concrete	0.03	0.03
1. Plywood floor	0.22	0.17
2. Glass glazing	0.25	0.18
3. Plaster wall	0.14	0.09

For both simulations, the sound source is placed 1.5 m above floor level along the room's horizontal symmetry axis, 1.5 m from the room's centre. Receivers are positioned at the centres of the 50 × 50 cm listener-plane grid cells at least 1.5 m from all room boundaries. The early-energy window is set to 50 ms after the arrival of the direct sound, and the coverage threshold to $x = 3$ dB. The analysis covers octave bands at 250 and 500 Hz, whose wavelengths (1.4 m and 0.7 m) are comparable to the scale of the geometric modifications applied to the vault.

The case study proceeds in two stages. The first stage (Baseline) explores different vault rises, comparing the resulting early-sound energy distributions through the descriptors σ and $C_{3\,\mathrm{dB}}$ introduced in Section 2.2. The selection prioritises $C_{3\,\mathrm{dB}}$, with σ used as a secondary criterion when coverage values are close; in cases of further ambiguity, the speech-related acoustic parameters obtained from Pyroomacoustics (C_{50}, T_{60}, and STI) are used to disambiguate among candidates. In the second stage, two shape-modifying configurations, ribs and diagonal ribs, are applied independently to the selected Baseline geometry using the functions introduced in Section 2.3., targeting the regions of the vault contributing most to first-order reflections converging on the listener plane.

3.1. Baseline iteration results

This iteration analyses barrel vaults whose rise (measured from the top of the room's vertical walls) ranges from 0.60 to 2.10 m in 0.30 m increments. Since h1.80 produces the most favourable early-energy distribution, and h2.10 shows a reduction in average $C_{3\,\mathrm{dB}}$ despite its marginally lower σ, two additional configurations at a finer 0.10 m step (h1.90 and h2.00) are evaluated to verify whether further refinement of the rise in that range yields additional improvement.

Among the configurations explored, those with a higher rise yield vault sections closer to a parabolic profile, result in a more uniform distribution of early sound energy across the listener plane and indicate reduced acoustic focusing (Table 2, Table 3). The improvement is most consistent in the coverage ratio $C_{3\,\mathrm{dB}}$, while σ varies less markedly across configurations. From a structural perspective, sections closer to a parabolic arch follow the funicular line under uniform vertical loading more closely, resulting in a more efficient distribution of compressive forces along the vault.

Since h1.80 and h2.00 both achieve the highest average $C_{3\,\mathrm{dB}}$ (36.1%), they are jointly selected as the Baseline candidates for subsequent refinement. Between the two, h2.00 additionally exhibits a near-optimal σ (5.21 dB, within 0.02 dB of the lowest value recorded at h2.10), while h1.80 performs slightly better at 500 Hz in terms of coverage (43.1% vs 40.3%). Both configurations are therefore carried forward to the shape-modifying stage.

Table 2: Standard Deviation σ [dB]

Rise	250 Hz	500 Hz	Avg
h0.60	5.57	5.09	5.33
h0.90	5.60	4.87	5.23
h1.20	5.82	4.84	5.33
h1.50	5.93	4.88	5.40
h1.80	**5.72**	**4.79**	**5.26**
h1.90	5.58	4.90	5.24
h2.00	**5.53**	**4.90**	**5.21**
h2.10	5.50	4.88	5.19

Table 3: Coverage $C_{3\,\mathrm{dB}}$ [%]

Rise	250 Hz	500 Hz	Avg
h0.60	29.2	29.2	29.2
h0.90	27.8	26.4	27.1
h1.20	27.8	34.7	31.2
h1.50	27.8	41.7	34.7
h1.80	**29.2**	**43.1**	**36.1**
h1.90	31.9	37.5	34.7
h2.00	**31.9**	**40.3**	**36.1**
h2.10	25.0	38.9	31.9

3.2. Ribs configuration

This configuration evaluates two distinct rib variants, generated by a shape-modifying function that adjusts the $\mathbf{h}_{\min}$ and $\mathbf{h}_{\max}$ values of selected arches at 1.5 m intervals in the region of the vault exhibiting the strongest first-order reflections, as identified by the ray-tracer analysis. This locally lowers the vault surface, producing ribs of varying height $\mathbf{h}_{\mathrm{r}}$. The first variant applies these adjustments along the arch direction only (Figure 6a). The second extends the modification to the edges perpendicular to the arches as well, introducing a slight curvature in the transverse direction (Figure 6b).

3.3. Diagonal ribs configuration

This configuration also evaluates two variants. It parametrically generates diagonal ribs within the region of the vault responsible for the strongest first-order reflections, placing the diagonals sequentially with each element mirroring the previous one to produce a continuous chevron pattern across the target area. As in the previous configuration, the first variant acts along the arch direction only (Figure 6c), while the second extends the modification to the edges perpendicular to the arches as well, introducing a slight curvature in the transverse direction (Figure 6d).

3.4. Optimisation results

Table 4 and Table 5 show the results of the shape-modifying configurations. Among the tested configurations, the diagonal rib variants applied to the h2.00 baseline (configurations c and d) produce the most consistent improvement, both reaching an average $C_{3\,\mathrm{dB}}$ of 41.7%, compared to 36.1% for the

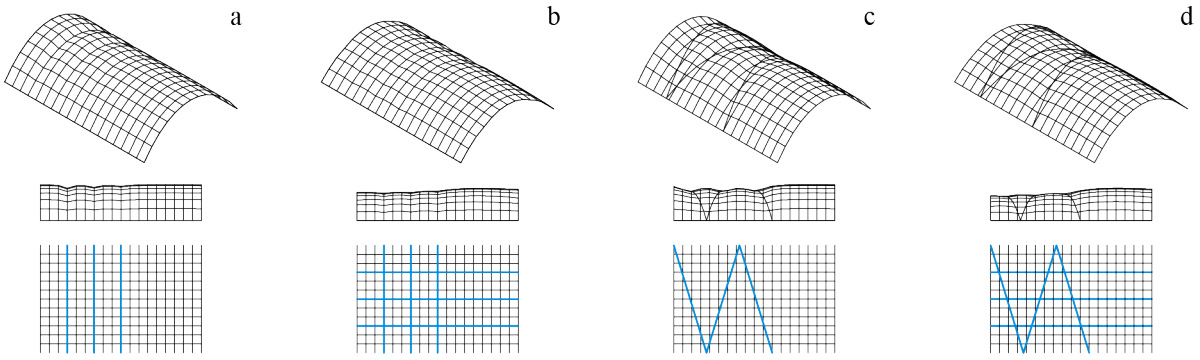

Figure 6: Shape-modifying configurations applied to the Baseline geometry. Ribs configuration: (a) along the arch direction only; (b) with added transverse curvature. Diagonal ribs configuration: (c) along the arch direction only; (d) with added transverse curvature.

unmodified h2.00 baseline, a gain of 5.6 percentage points. The transverse rib variant (b) applied to h2.00 also performs well, achieving 39.6% with the lowest σ recorded across all tested configurations (5.15 dB).

It should be noted that the ranking is based exclusively on the early-energy descriptors σ and $C_{3\,\mathrm{dB}}$, which capture only the first 50 ms of the sound field. The Pyroomacoustics benchmark (Figure 7) provides a complementary evaluation through the full impulse response. All configurations remain within the *Fair* to *Good* STI range (0.50–0.54), with the diagonal rib variants (c, d) achieving marginally higher intelligibility than the baseline. The clarity index C_{50} improves across the rib configurations at both frequencies, corroborating the early-energy results. As expected, the reverberation time RT_{60} remains above 1.0 s across all configurations, indicating that the geometric modifications alone are insufficient to meet the speech reverberation target; additional absorptive treatment would be required to address this. The two evaluations are therefore complementary: the early-energy descriptors identify configurations that redistribute sound more uniformly, while the Pyroomacoustics parameters confirm their effect on overall speech intelligibility.

Table 4: Standard Deviation σ [dB]

Configuration	**250 Hz**	**500 Hz**	**Avg**
h1.80	5.72	4.79	5.26
h2.00	5.53	4.90	5.21
a. h2.00; h_r 0.15 to 0.25	5.51	4.72	5.11
b. h1.80; h_r 0.05 to 0.15	5.74	4.77	5.26
b. h2.00; h_r 0.05 to 0.15	5.51	4.79	5.15
c. h2.00; h_r 0.10 to 0.30	5.67	4.82	5.25
c. h2.00; h_r 0.15 to 0.35	5.73	4.85	5.29
d. h1.80; h_r 0.05 to 0.30	5.81	4.78	5.29
d. h2.00; h_r 0.05 to 0.35	5.83	4.79	5.31

Table 5: Coverage $C_{3\,\mathrm{dB}}$ [%]

Configuration	**250 Hz**	**500 Hz**	**Avg**
h1.80	29.2	43.1	36.1
h2.00	31.9	40.3	36.1
a. h2.00; h_r 0.15 to 0.25	27.8	45.8	36.8
b. h1.80; h_r 0.05 to 0.15	27.8	47.2	37.5
b. h2.00; h_r 0.05 to 0.15	33.3	45.8	39.6
c. h2.00; h_r 0.10 to 0.30	27.8	55.6	41.7
c. h2.00; h_r 0.15 to 0.35	29.2	52.8	41.0
d. h1.80; h_r 0.05 to 0.30	30.6	51.4	41.0
d. h2.00; h_r 0.05 to 0.35	27.8	55.6	41.7

4. Conclusions and Discussion

The proposed workflow demonstrates a viable approach for the early-stage design of compression-dominant structures, in which acoustic performance is integrated directly into the form-finding process. By coupling Thrust Network Analysis with a custom acoustic ray-tracer employing bidirectional scattering coefficients for more accurate representation of sound reflections, and by combining the spatial distribution of early sound energy with backward tracing of the contributing surface regions, the workflow makes the relationship between specific geometric features of the vault and the resulting sound-energy distribution explicit and interpretable.

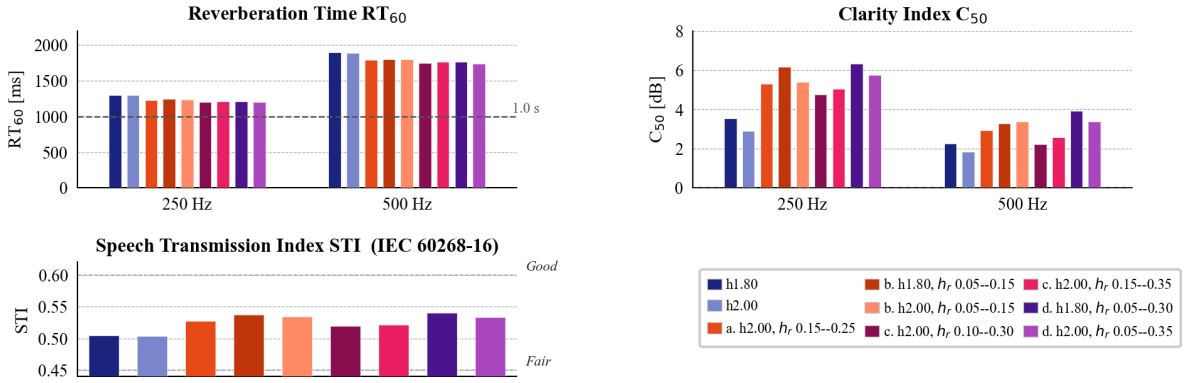

Figure 7: Pyroomacoustics simulation results — RT_{60}, C_{50}, and STI (IEC 60268-16) — for selected configurations at 250 and 500 Hz. Baseline (h = 1.80 and 2.00 m); a: ribs along the arch direction; b: ribs with transverse curvature; c: diagonal ribs along the arch direction; d: diagonal ribs with transverse curvature.

The application to the barrel-vault classroom shows that adjusting the height bounds of the transverse arches provides effective control over the distribution of early sound energy at the listener plane. Among the tested rise heights, configurations closer to a parabolic profile produce a more uniform energy distribution; the configurations at $h = 1.80$ m and $h = 2.00$ m both achieve the highest average $C_{3\,\mathrm{dB}}$ (36.1%) and are jointly selected as baseline candidates, offering a favourable trade-off between acoustic uniformity and structural efficiency. The subsequent application of localised diagonal ribs (configurations c and d) further improves the coverage ratio by up to 5.6 percentage, while the Pyroomacoustics benchmark confirms a corresponding improvement in C_{50} and STI relative to the baseline.

Current limitations open future research directions. The workflow presented in this paper relies on geometric acoustics, which is reliable in the mid-to-high frequency range but does not capture wave phenomena that become significant at lower frequencies. A natural next step is to integrate a wave-based solver, such as FEM-based Python libraries, directly into the workflow, allowing a unified treatment of the full audible frequency range. Then, the geometric-acoustics approach can be validated against this solver to establish the lower bound on its frequency range of validity. A further extension is the development of a BRDF library covering different surface pattern typologies, such as coffers, ribs, and stereotomic textures, whose dimensions interact primarily with shorter wavelengths. Integrating such a library would allow the workflow to address scattering behaviour at higher frequencies, where surface microstructure becomes acoustically significant, complementing the global geometric modifications already covered by the current implementation. Further developments will also extend the methodology to other compression-dominant structure typologies.

Acknowledgements

AI assistance disclosure: This paper was developed with the assistance of an AI language model for tasks including structuring and editing text, and refining clarity and coherence. All critical analyses, methodological designs, and scientific content were developed independently by the authors.

References

[1] J. Heyman, *The Stone Skeleton: Structural Engineering of Masonry Architecture*. (1st paperback ed.). New York, NY: Cambridge University Press, 1997.

[2] S. Bhooshan, A. Dell'Endice, F. Ranaudo, T. Van Mele and P. Block, "Unreinforced concrete masonry for circular construction", *Architectural Intelligence*, vol. 3, no. 1, 2024, p. 7. doi:

10.1007/s44223-023-00043-y.

[3] S. Adriaenssens, P. Block, D. Veenendaal and C. Williams Eds., *Shell Structures for Architecture: Form Finding and Optimization*. London; New York: Routledge/Taylor & Francis Group, 2014.

[4] Block Research Group, *COMPAS FormFinder for Rhino*. Software, ETH Zurich, 2024–. Available: https://github.com/BlockResearchGroup/compas-FoFin.

[5] Block Research Group, *RhinoVAULT: Funicular Form-Finding for Rhino*. Software, ETH Zurich. Available: https://blockresearchgroup.gitbook.io/rhinovault.

[6] T. Méndez Echenagucia and P. Block, "Acoustic Optimization of Funicular Shells", in *Proceedings of the International Association for Shell and Spatial Structures Symposium 2015*. Amsterdam, The Netherlands, 2015, pp. 1–10. Available: https://brg.ethz.ch/publications/98.

[7] T. Wulfrank, Y. Jurkiewicz and E. Kahle, "Design-Focused Acoustic Analysis of Curved Geometries Using a Differential Raytracing Technique", *Building Acoustics*, vol. 21, no. 1, 2014, pp. 87–95. doi: 10.1260/1351-010X.21.1.87.

[8] M. L. S. Vercammen, *Sound Concentration Caused by Curved Surfaces*. PhD thesis, Technische Universiteit Eindhoven, 2012. doi: 10.6100/IR732483.

[9] T. Mihajlovska, A. Trombeva-Gavriloska, M. Cvetkovska and L. Dimevska-Sofronievska, "Graphic statics in the digital age: A critical review of current methods and trends", *Gradjevinski materijali i konstrukcije*, vol. 68, no. 3, 2025, pp. 163–173. doi: 10.5937/GRMK2500009M.

[10] P. Block and J. Ochsendorf, "Thrust network analysis: A new methodology for three-dimensional equilibrium", *Journal of the International Association for Shell and Spatial Structures*, vol. 48, no. 3, 2007, pp. 167–173.

[11] H. Kuttruff, *Room Acoustics*. (6th ed.). Boca Raton: CRC Press, 2016. doi: 10.1201/9781315372150.

[12] M. Vorländer and E. Mommertz, "Definition and measurement of random-incidence scattering coefficients", *Applied Acoustics*, vol. 60, no. 2, 2000, pp. 187–199. doi: 10.1016/S0003-682X(99)00056-0.

[13] T. Cox and P. D'Antonio, *Acoustic Absorbers and Diffusers: Theory, Design and Application*. (3rd ed.). Boca Raton: CRC Press, 2016. doi: 10.1201/9781315369211.

[14] A. Heimes and M. Vorländer, "Bidirectional surface scattering coefficients", *Acta Acustica*, vol. 9, 2025, p. 41. doi: 10.1051/aacus/2025026.

[15] I. W. Foged, A. Pasold and M. B. Jensen, "Evolution of an Instrumental Architecture", in *Proceedings of the 32nd eCAADe Conference*, Vol. 2, Newcastle upon Tyne, UK, 2014, pp. 365–372.

[16] T. Van Mele and Block Research Group, *compas_tna: TNA Package for the COMPAS Framework*. Software, ETH Zurich, 2017–. Available: https://github.com/BlockResearchGroup/compas_tna.

[17] International Organization for Standardization, *ISO 9613-1:1993 Acoustics — Attenuation of Sound during Propagation Outdoors — Part 1: Calculation of the Absorption of Sound by the Atmosphere*. Geneva: ISO, 1993.

[18] J. S. Bradley, H. Sato and M. Picard, "On the importance of early reflections for speech in rooms", *Journal of the Acoustical Society of America*, vol. 113, no. 6, 2003, pp. 3233–3244. doi: 10.1121/1.1570439.

[19] R. Scheibler, E. Bezzam and I. Dokmanić, "Pyroomacoustics: A Python package for audio room simulation and array processing algorithms", in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Calgary, Canada, 2018, pp. 351–355. doi: 10.1109/ICASSP.2018.8461310.