



Raytracing in Galerkin Boundary Integral Form

Amal Emthyas¹

Acoustics Research Group

Newton Building, University of Salford, Salford, M5 4WT, United Kingdom

Jonathan A. Hargreaves²

Acoustics Research Group

Newton Building, University of Salford, Salford, M5 4WT, United Kingdom

ABSTRACT

Raytracing is an established method for computing the late-time part of room impulse responses. But it has the drawback that only very crude Monte Carlo models of boundary scattering and diffraction are possible to include without losing its attractive computational cost scaling. This happens because higher-resolution models of these processes output multiple child rays for every parent ray received, making the number of rays grow with reflection order. An emerging solution is so-called ‘Surface-Based’ Geometrical Acoustics. Here the distribution of rays arriving at a boundary is mapped onto a predefined set of spatial elements and angular interpolation functions, producing a vector of coefficients. Re-radiation of subsequent reflections is then a matrix multiplication, with the steady state solution being solvable via a Neumann series. Since rays only ever propagate one reflection order before being collected, diffraction and scattering process that cause them to multiply can be included without issue. Here we present a new formulation based on a Galerkin Boundary Element Method (BEM). A unique feature is its ability to readily change interpolation functions, so their effect on accuracy and convergence can be assessed. In this preliminary work, it is verified against an Image Source Model for a rectangular room.

1. INTRODUCTION

Prediction models are at the heart of modern acoustic engineering. They are used in a diverse range of applications, from refining the acoustic design of classrooms and concert halls to predicting how noise exposure varies through an urban environment. They also allow Auralisation to be performed for buildings and spaces before they are built or long after they are lost [1].

Geometrical Acoustics (GA) algorithms currently dominate in room acoustics applications [2]. This is an umbrella category for several algorithms, including Image Source Method (ISM) and Raytracing. These are common in their assumption that sound propagates geometrically, i.e., in straight lines in a homogeneous medium such as air in a room. Notably, GA methods do not automatically include diffraction, meaning it must be added in as an additional process where required. This contrasts with ‘low frequency’ solvers, such as the Finite Element Method (FEM) or Boundary Element Method (BEM), which directly include diffraction as part of how they solve the underlying Partial Differential Equation (PDE) that represents acoustic wave propagation within the domain under study. ISM constructs a solution from simple sources that individually satisfy the PDE, but the way it assembles them is only valid for rectangular (2D) or cuboid (3D) domains. Raytracing is related to the acoustic PDE via the eikonal approximation.

¹ a.emthyas1@edu.salford.ac.uk

² j.a.hargreaves@salford.ac.uk

ISM is effective for reflections from smooth planar boundaries. Factors that constrain the accuracy include the inability of the standard algorithm to include boundary scattering or diffraction (the latter being zero in a rectangular or cuboid domain). These must be added on as additional processes [3,4]. But it's main constraining factor is its exponential computational cost scaling with reflection order. This limits it to computing low-order early reflections only in most cases.

Raytracing is popular for late-time impulse response generation because its computational cost is linear with reflection order. However, all but the lowest-fidelity models of boundary scattering and diffraction lead to ray-splitting, and this cripples it with exponential cost scaling with reflection order. Recent work has used precomputed edge-to-edge visibility graphs to accelerate this, but this was only applied to 'direct' diffraction and the algorithm retains exponential complexity [5].

A final constraint is that the above methods bear little similarity to established methods for low frequencies (FEM and/or BEM), hindering prospects for a unified full-audible-bandwidth approach.

1.1 Surface-Based GA

An emerging category of models is what Savioja and Svensson termed "surface-based" GA [2]. These discretise the distribution of acoustic power flow versus position using a boundary mesh, just like BEM uses for low frequency problems. Indeed, it has been called "Energy-Intensity" BEM by some authors [6,7]. The oldest algorithm in this category is Acoustic Radiosity, developed by Kuttruff in the 1970s [8,9]. This does not consider the directional nature of sound arriving at the boundary and instead considers these to all scatter diffusely according to Lambert's law, which is not a realistic model of real rooms. Results tend not to be too unreasonable for late reflections, where the sound field is diffuse anyway, but it is inaccurate for early reflections, which are sparse in direction. Usefully, however, it can be used to find T60 in rooms with non-uniform absorption [8].

More sophisticated algorithms also parameterise reflection direction [6,7,10–13]. This approach appears to have been developed independently by three different groups, but they can be shown to all be equivalent once choices in dimensionality, interpolation and notation are taken into account. One of the contributions of this paper is a framework that makes this equivalence clearer. A version of the algorithm that operates on volumetric meshes has also been demonstrated [14] but, while this opens up exciting applications in inhomogeneous and structure-borne sound problems, its benefits for homogeneous room acoustics problems are less clear.

A characteristic of the surface-based GA method is that the way the solution is interpolated – mathematicians would call this the 'approximation space' – is defined in advance. Spatially this is based on elements on a mesh, for which all the standard choices of interpolation (a.k.a. 'shape' or 'basis') function are possible. Common choices are piecewise-constant (PWC), piecewise-linear or piecewise-quadratic, where the latter two can be continuous or discontinuous but PWC is always discontinuous. To capture the distribution with respect to (w.r.t.) reflection angle, a set of PWC functions on non-overlapping angle ranges can be used – ref. [13] did this in 3D – or orthogonal polynomials such as Legendre or Chebyshev [10,11]. For the latter, it is more convenient to use an edge-tangential coordinate rather than angle, since this lies in the range -1 to +1. Franzoni and co-workers used a half Fourier series in angle in 2001 and a full version in 2005, leading to problems of non-orthogonality and periodicity respectively. A cosine series in angle would have eradicated both these problems and is equivalent to Chebyshev polynomials in the tangential coordinate.

A key benefit of these algorithms is that they maintain a fixed number of degrees of freedom (DOF) irrespective of reflection order. Reflections are advanced by a matrix multiplication; hence computational cost is linear with reflection order. Crucially, this holds even if the raytracing process that occurs within a reflection order involves a lot of ray-splitting, as is required to model scattering or diffraction accurately [2,15]. The extra rays are absorbed into the predefined approximation space and the number of degrees of freedom is kept constant and computational cost linear. Other benefits are that steady-state can be easily solved for [10,11] and decay rates could be found too [8].

The discussion above of "approximation" makes it clear that this is present in the method. Repeated discretisation of high-order reflections will not give results identical to those from tracing the rays individually. This was viewed as a limitation by Pohl and Stephenson [15], who proposed a similar idea with their 'ray recombination' algorithm, but we do not see it that way. All raytracing

algorithms acknowledge that individual rays do not describe sound intensity. Instead, it is their spatial distribution that captures this, hence receivers in classical raytracing must have volume. Nonetheless, the tracing of rays to high order makes it tempting to believe that each ray's path over many reflection orders matters but it really does not – it is just one sample within a distribution. This is especially true when boundary scattering is present. In this case reflection density increases making individual reflections no longer important [2]. Surface-based GA exploits this to simplify the algorithm and provide computational benefits, plus the approximation is controllable. This work shows that early reflections can be captured accurately if the number of DOF is sufficiently high.

Some articles have suggested that surface-based GA is viewed as computationally expensive [9], but this view appears misplaced based on the discussion above and is not backed up by the all-important computational cost scaling. What is true is that much of the computational effort is in the problem setup, so must be performed beforehand. We see such precomputation as an advantage. As Nosal *et al* [9] point out, listener position can be quickly recomputed once the boundary solution has been computed. Additionally, source position can be changed quite readily since the boundary-to-boundary interaction matrices – computation of which are the main bottleneck – do not depend on source position. Identical arguments apply to standard BEM, and the algorithm presented herein matches exactly the framework in fig. 1 of ref. [16]. The similarity of surface-based GA to BEM is discussed further in this paper's sister paper, #1005 in these proceedings [17].

1.2 Overview of this paper

The main contribution of this paper is to cast surface-based GA into a Galerkin BEM framework. This expedites comparison of different basis and test functions since it moves factors required to compensate for these out of the main code body and into a matrix operation. It is also motivated by a desire to emphasise the synergy between surface-based GA and high-frequency BEM. The algorithm and matrix structure used herein matches exactly that of the high-frequency BEM in ref. [18]. This similarity is discussed further in paper #1005 in these proceedings [17].

Source and receiver directivity are also implemented using Fourier series of arbitrary order. Stacking of their coefficients into vectors means that all acoustic wave propagation operations are represented by a matrix multiplication, be it source to boundary, boundary to boundary, boundary to receiver, or direct source to receiver. This mimics the structure in fig. 1 of ref. [16].

Results are presented for early reflections and validated against an ISM that has been modified to use the same source and receiver directivity model. Agreement is seen to be very good, though differences exist between different discretisation schemes.

2. THEORY AND FORMULATION

This section presents the background theory behind the method. There is insufficient space for a full exposition, but the key concepts and features will be presented and discussed.

2.1 The Assumed Model of Acoustic Propagation

As with all GA, it is assumed that we are working in a high-frequency regime. The first starting assumption is that acoustic propagation can be represented by particles travelling in straight lines termed rays. An individual particle would have position $\mathbf{x}$ and propagation direction $\hat{\mathbf{k}}$. The power distribution $w(\mathbf{x}, \hat{\mathbf{k}})$ of these particles is therefore a function of both position and direction.

The second starting assumption is that we will not consider phase – we only consider energy and power, albeit including its directional characteristics. This is also true for classical raytracing. It is justified by arguments of the chaotic nature of late-time sound fields and ensemble averaging [19].

The model will be formulated in two dimensions (2D) for simplicity. Within the domain Ω , position $\mathbf{x}$ can vary in 2D and the propagation direction vector $\hat{\mathbf{k}}$ can point toward any angle within a 360° range. It will therefore be convenient to define $\hat{\mathbf{k}} = [\cos \theta \quad \sin \theta]$, where θ is a global angle relative to the x-axis of the global Cartesian coordinate system. But when $\mathbf{x}$ lies on the boundary Γ the situation is slightly different. Now $\mathbf{x}$ is only a function of boundary position, a 1D quantity (the boundary of a 2D domain is a 1D line) such as length l in meters along the boundary from some




Figure 1: Geometry for source and boundary definitions: a) Incoming and outgoing directions: b) Power radiating from a point source; c) Geometry for point source to boundary mapping

reference point. The orientation of the boundary Γ at $\mathbf{x}$ is given by its normal vector $\hat{\mathbf{n}}$ and tangent vector $\hat{\mathbf{t}}$, which are perpendicular. $\hat{\mathbf{n}}$ is taken to point into the acoustic domain Ω .

It will be convenient to distinguish between incoming directions, for which $\hat{\mathbf{n}} \cdot \hat{\mathbf{k}} < 0$, and outgoing directions, for which $\hat{\mathbf{n}} \cdot \hat{\mathbf{k}} > 0$. These will be termed $\hat{\mathbf{k}}^-$ and $\hat{\mathbf{k}}^+$ respectively, and each occupies a 180° range. $\hat{\mathbf{k}}^\pm$ are defined in the local coordinate system defined by $\hat{\mathbf{n}}$ and $\hat{\mathbf{t}}$ and can be parameterised by the angle θ made with the normal vector $\hat{\mathbf{n}}$. This is shown in Figure 1a.

Another option, which is closer to what is done in refs. [10,11], is to use a boundary tangential coordinate $\chi = \sin \theta$ to represent direction. This is bounded by $|\chi| \leq 1$. $\hat{\mathbf{k}}^\pm$ are then given by:

$$\begin{aligned}\hat{\mathbf{k}}^+(\chi) &= \hat{\mathbf{t}}_n \chi + \hat{\mathbf{n}}_n \sqrt{1 - \chi^2}, \\ \hat{\mathbf{k}}^-(\chi) &= \hat{\mathbf{t}}_n \chi - \hat{\mathbf{n}}_n \sqrt{1 - \chi^2}.\end{aligned}\tag{1}$$

In this paper, $\hat{\mathbf{k}}$, θ and χ will be used interchangeably based on these definitions. Typically, θ is used when considering if propagation angles match, χ is used when evaluating angular integrals numerically, and $\hat{\mathbf{k}}$ is used to ease definitions by avoiding the need to align local coordinates.

2.2 Power incident from a point source

Consider power radiating from a point source at position $\mathbf{y}_0$, as depicted in Figure 1b. The source injects power W into the domain. The units of W is Watts/metre because the model is in 2D, and a 2D point source is equivalent to a 3D line source.

Consider now a receiver at position $\mathbf{x}$. It receives a total power density of $W/2\pi R$ due to cylindrical spreading. Here $\mathbf{R} = \mathbf{x} - \mathbf{y}_0$ and $R = |\mathbf{R}|$. The unit of this density is Watts per metre², where this metre² is area tangential to the wavefront. The intensity at $\mathbf{x}$ is this power density times the propagation direction $\hat{\mathbf{R}} = \mathbf{R}/R$, so $\mathbf{I} = \hat{\mathbf{R}} W/2\pi R$.

Now additionally consider that the source may be directional, so the power injected varies with direction θ_s (which is defined in the global sense, i.e., relative to the Cartesian axes, because $\mathbf{y}_0$ is in the domain Ω). The power injected is replaced by a directivity function $W_s(\theta_s)$. This is defined so that its integral w.r.t. θ_s from 0 to 2π equals W . It follows that $W_s(\theta_s)$, therefore, has units Watts per metre per radian. For example, an omnidirectional source has $W_s(\theta_s) = W/2\pi$.

The power density at $\mathbf{x}$ is now $W_s(\theta_s)/R$. The 2π has been removed from this statement because of the units of $W_s(\theta_s)$. Intensity is $\mathbf{I} = \hat{\mathbf{R}} W_s(\theta_s)/R$. Note that a real line source would have nearfield effects, but these have been ignored because one of the starting assumptions is that the algorithm is being derived in the high frequency regime.

The directional power distribution $w(\mathbf{x}, \hat{\mathbf{k}}) \equiv w(\mathbf{x}, \theta_r)$ is the acoustic quantity being propagated, so it must be defined for power incident from a point source. θ_r is the evaluation angle at the receiver $\mathbf{x}$ and w has units Watts per metre² per radian; here the “per radian” is w.r.t. receiving

angle θ_r . The rays emanating from this source in the high-frequency regime are perfectly directional, so are represented by a Dirac delta function in angle $\delta(\theta_s - \theta_r)$. Hence, w is given by:

$$w(\mathbf{x}, \hat{\mathbf{k}}) = w(\mathbf{x}, \theta_r) = \frac{W_s(\theta_s)}{R} \delta(\theta_s - \theta_r) = \frac{W_s(\theta_s)}{R} \delta(\cos^{-1}(\hat{\mathbf{k}} \cdot \hat{\mathbf{R}})). \quad (2)$$

The total power arriving at the receiver irrespective of direction is:

$$W_r(\mathbf{x}) = \oint w(\mathbf{x}, \hat{\mathbf{k}}) d\hat{\mathbf{k}} \equiv \int_0^{2\pi} \frac{W_s(\theta_s)}{R} \delta(\theta_s - \theta_r) d\theta_r = \frac{W_s(\theta_s)}{R}. \quad (3)$$

Here the ‘sifting’ property of the delta function has been exploited to remove the integral and set $\theta_r = \theta_s$. W_r is spatial power density with units Watts/metre², like intensity has – the ‘per radian’ has been removed by angular integration. Note that for an omnidirectional source, with $W_s(\theta_s) = W/2\pi$, this simplifies to $W/2\pi R$ as expected. A similar integral can be used to compute intensity.

2.3 Point Source to Boundary Mapping

The objective here is to map the directional power distribution $w(\mathbf{x}, \hat{\mathbf{k}})$ emanating from the source to an equivalent incoming distribution $w^-(\mathbf{x}, \hat{\mathbf{k}}^-)$ based in the boundary coordinate system.

The source is radiating rays in all directions, and the boundary can accept rays from any direction within the domain, but only ones that align in both receiving position and angle are geometrically valid. This creates an equivalence between θ_s and θ_x (see Figure 1c), but this is not immediately straightforward to state because they are defined relative to different reference directions. But it can be expressed easily by requiring that $\cos^{-1}(\hat{\mathbf{k}}^- \cdot \hat{\mathbf{R}}) = 0$. This is satisfied by $\hat{\mathbf{k}}^- = \hat{\mathbf{R}}$, hence a boundary distribution including an angular term $\delta(\cos^{-1}(\hat{\mathbf{k}}^- \cdot \hat{\mathbf{R}}))$ will sift out a ray direction (blue arrow) matching the arrival direction from the source (orange arrow).

The other aspect that must be accounted for is whether the incident ray arrives from the correct side of the boundary and whether the ray path is occluded. Such so-called “visibility” and “validity” checks [2] are encapsulated into a function $V(\mathbf{x}, \mathbf{y}_0)$, which takes the value 0 or 1. However, this is always 1 in a convex domain, such as the rectangular room considered herein, so the detail of this is considered here no further. The incoming boundary power distribution is therefore:

$$w^-(\mathbf{x}, \hat{\mathbf{k}}^-) = \frac{W_s(\theta_s)}{R} \delta(\cos^{-1}(\hat{\mathbf{k}}^- \cdot \hat{\mathbf{R}})) V(\mathbf{x}, \mathbf{y}_0) \quad (4)$$

2.4 Re-radiation from the Boundary

The objective here is to determine how the energy distribution $w^+(\mathbf{y}, \hat{\mathbf{k}}^+)$ leaving the boundary at $\mathbf{y}$ propagates through the domain Ω and then maps onto a distribution $w(\mathbf{x}, \hat{\mathbf{k}})$ at a point receiver $\mathbf{x}$ or an incoming distribution $w^-(\mathbf{x}, \hat{\mathbf{k}}^-)$ on another boundary Γ_x . Figure 2 shows the geometry. The ray (shown in orange) connects $\mathbf{y}$ and $\mathbf{x}$ and the direction of propagation aligns with it, hence $\hat{\mathbf{k}}^+ = \hat{\mathbf{k}} = \hat{\mathbf{k}}^- = \hat{\mathbf{R}}$. But how this is parameterised as an angle varies:

- For at $\mathbf{y} \in \Gamma_y$ the angle θ_y is relative to the normal vector $\hat{\mathbf{n}}_y$ at $\mathbf{y}$;
- For a point receiver $\mathbf{x} \in \Omega$ the angle θ_r is relative to the global Cartesian x-axis;
- For $\mathbf{x} \in \Gamma_x$ the angle θ_x is relative to the inverted normal vector $-\hat{\mathbf{n}}_x$ at $\mathbf{x}$.

In all three cases, the units of w are Watts per metre squared per radian, where the “per metre squared” is taken in the plane perpendicular to the direction of propagation. The formula will be developed for a domain receiver $w(\mathbf{x}, \hat{\mathbf{k}})$ but this also applies to a boundary point $w^-(\mathbf{x}, \hat{\mathbf{k}}^-)$.




Figure 2: Geometry for evaluating boundary to boundary interactions

All points on Γ_y can radiate to $\mathbf{x}$ and the power density received will be the summation of their contributions. Because of this, there is an integral over boundary position. The boundary can radiate in any angle too, hence we also have an integral over θ_y . $w(\mathbf{x}, \hat{\mathbf{k}})$ therefore is given by:

$$w(\mathbf{x}, \hat{\mathbf{k}}) = \int_{\Gamma} \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \frac{w^+(\mathbf{y}, \hat{\mathbf{k}}^+)}{R} \delta(\cos^{-1}(\hat{\mathbf{k}} \cdot \hat{\mathbf{R}})) \delta(\cos^{-1}(\hat{\mathbf{k}}^+ \cdot \hat{\mathbf{R}})) V(\mathbf{x}, \mathbf{y}) \cos \theta_y d\theta_y d\Gamma_y \quad (5)$$

In this, $\mathbf{x}$ and $\hat{\mathbf{k}}$ are arguments of w . This means it cannot be assumed that $\hat{\mathbf{k}}$ is aligned with $\hat{\mathbf{R}}$, but the statement needs to ‘sift’ out values of $\hat{\mathbf{k}}$ that are. The $\delta(\cos^{-1}(\hat{\mathbf{k}} \cdot \hat{\mathbf{R}}))$ term achieves this. Similarly, $\hat{\mathbf{k}}^+$ is a function of the integration variable θ_y , but the geometric model of acoustic propagation means that only radiation from the boundary in the direction $\hat{\mathbf{k}}^+ = \hat{\mathbf{R}}$ contributes. The term $\delta(\cos^{-1}(\hat{\mathbf{k}}^+ \cdot \hat{\mathbf{R}}))$ sifts this out. The $1/R$ arises due to spherical spreading from the boundary sources. It converts the ‘per radian’ of $w^+(\mathbf{y}, \hat{\mathbf{k}}^+)$ at $\mathbf{y}$ into a ‘per metre’ tangential to the wavefront at $\mathbf{x}$. One of the ‘per metre’ of $w^+(\mathbf{y}, \hat{\mathbf{k}}^+)$ at $\mathbf{y}$ will be integrated out by the spatial integral over Γ_y . But this ‘per metre’ of $w^+(\mathbf{y}, \hat{\mathbf{k}}^+)$ is in the plane perpendicular to the direction of propagation, whereas the integral is w.r.t. boundary length. The $\cos \theta_y$ term compensates for this. The second delta function can be exploited to analytically evaluate the integral w.r.t θ_y via the sifting property:

$$w(\mathbf{x}, \hat{\mathbf{k}}) = \int_{\Gamma} \frac{w^+(\mathbf{y}, \hat{\mathbf{R}})}{R} \delta(\cos^{-1}(\hat{\mathbf{k}} \cdot \hat{\mathbf{R}})) V(\mathbf{x}, \mathbf{y}) \hat{\mathbf{n}}_y \cdot \hat{\mathbf{R}} d\Gamma_y. \quad (6)$$

This is the directional energy distribution arriving at a point receiver in the domain or at a point on a boundary. $\cos \theta_y = \hat{\mathbf{n}}_y \cdot \hat{\mathbf{R}}$ has been substituted into it for ease of numerical evaluation.

Eq. (5) is equivalent to the Frobenius-Perron operator that the Nottingham group [10,11] use as their starting point. The only significant difference is that they integrate (and later sift) in momentum (equivalent to position and χ together), whereas we have integrated w.r.t θ_y , instead of χ_y , and sift in θ_y and θ_x . This choice means they omit the $\cos \theta_y$ and $1/R$, but a later coordinate transform reintroduces these and a $\cos \theta_x$ term. This makes our formulation equivalent, as is the one from Duke.

2.5 Discretisation

We need to discretise $w(\mathbf{x}, \hat{\mathbf{k}})$ on the boundary. Other papers have usually discretised only the incoming or the outgoing sound energy, and then included somehow the fact that w^- is reflected to give w^+ . We will discretise w^- and w^+ separately, and then account for boundary scattering via multiplication with a ‘scattering’ matrix.

The boundary Γ is meshed using a closed loop of N boundary elements. Position on this is indexed by a unitless position parameter $0 \leq \mu \leq N$ where N is the total number of mesh elements. The boundary quantities w^- and w^+ are discretised by a weighted sum of basis functions:

$$w^\pm(\mathbf{x}(\mu), \hat{\mathbf{k}}^\pm(\chi)) = \sum_{n=0}^N \sum_{q=0}^Q \mathbf{w}^\pm[nQ + q] b_n(\mu) a_q(\chi). \quad (7)$$

Here $\mathbf{w}^-$ and $\mathbf{w}^+$ are vectors of coefficients for the incoming and outgoing waves w^- and w^+ respectively. The two sets of indices n and q have been collapsed into a single index $nQ + q$. $b_n(\mu)$ is the spatial basis function. This is chosen to be piecewise-constant (PWC), so 1 on element Γ_n and zero otherwise. $a_q(\chi)$ is the angular interpolation function. It might be an orthogonal polynomial, such as Legendre or Chebyshev, or PWC discretisation of $-1 \leq \chi \leq 1$ into Q regions.

The source directivity $W_s(\theta_s)$ will be represented by a Fourier series in angle θ_s :

$$W_s(\theta_s) = \sum_{q=0}^{Q_s} a_q \cos q\theta_s + \sum_{q=1}^{Q_s} b_q \sin q\theta_s. \quad (8)$$

Here Q_s is the source order and the set of coefficients a_q and b_q are parameters supplied by the user. For examples, An omnidirectional source has $Q_s = 0$ and $a_0 = W/2\pi$, where W is the total source power in Watts / metre. A Fourier series in angle θ_r will also be used to represent directional energy arriving at receivers. This is akin to 2D Ambisonics. $w(\mathbf{x}, \hat{\mathbf{k}})$ will be represented by:

$$w(\mathbf{x}, \hat{\mathbf{k}}(\theta_r)) \approx \sum_{p=0}^{P_r} a_p \cos p\theta_r + \sum_{p=1}^{P_r} b_p \sin p\theta_r \quad (9)$$

Here P_r is the receiver order. The coefficients a_p and b_p are output data computed by the algorithm by finding w at $\mathbf{x}$ and then mapping it onto the representation in eq. (9). This mirrors the framework used in ref. [16] for auralisation of BEM.

2.6 Galerkin BEM

There is insufficient space here to fully explain the details of the Galerkin BEM formulation. It involves a second pair of integrals over position on the boundary and receiving angle that must be evaluated, though another of these can be sifted out leaving just two boundary integrals to evaluate. These integrals include another pair of basis functions, sometimes called the ‘testing functions’ [18]. These allow choice in how error in the solution is minimised. Choosing them to match the basis functions in eq. (7) forms a ‘Bubnov–Galerkin method’, which minimises mean-squared error in the same approximation space as the solution. Choosing them to be different forms a ‘Petrov–Galerkin method’, which minimises mean-squared error in some other space. Rowbottom and Chappell recently opted to do the latter [20], using delta functions as basis in angle (so a finite set of angles instead of a distribution) and PWC testing functions. This produces an effect very similar to the ‘ray direction quantisation’ discussed by Pohl and Stephenson [15]. We opt to do the former for reasons of simplicity and symmetry, and synergy with ref. [18].

A feature of Galerkin BEM is that it does not compute the coefficients of eq. (7) directly. Instead, it computes a vector of ‘projections’, from which the coefficients can be found by solving a matrix equation. The matrix involved in this equation is an identity matrix of the approximation / testing space, and is commonly called the ‘mass matrix’, due to its similarity with the mass matrix in FEM. For an orthogonal discretisation and testing scheme, such as arises from pairs of PWC or Legendre functions, it is diagonal so has NQ non-zeroes. For PWC discretisation in space and angle its entries are $l_m \times 2/Q$, where l_m is the length of the m^{th} element. For PWC discretisation in space and Legendre in angle, its entries are $l_m \times 2/(2p + 1)$. Chebyshev functions are not orthogonal with the chosen weighting in χ , making their mass matrix block diagonal with NQ^2 non-zeroes. All mass matrices were therefore seen to be highly sparse and readily invertible. The ones encountered during this study all had condition numbers of 40 or below.

2.7 Algorithm Structure

The process for solving the discretised version of the scheme above is as follows:

- 1) First the directional power from the source, represented by $\mathbf{w}_s$, is mapped onto the boundary as an incoming power distribution $w^-(\mathbf{x}, \hat{\mathbf{k}}^-)$, represented by $\mathbf{w}_0^-$
- 2) This is reflected in the boundary, including scattering and absorption as appropriate, to give an outgoing power distribution $w^+(\mathbf{x}, \hat{\mathbf{k}}^+)$, represented by $\mathbf{w}_0^+$
- 3) The outgoing power distribution $w^+(\mathbf{x}, \hat{\mathbf{k}}^+)$ is propagated across the domain to the boundary, where it is mapped onto an incoming power distribution $w^-(\mathbf{x}, \hat{\mathbf{k}}^-)$, represented by $\mathbf{w}_1^-$
- 4) Either steps 2 and 3 are repeated up the required number of reflection orders, and $\mathbf{w}_i^\pm$ are summed with respect to reflection order i to obtain the total boundary power distributions, or a Neumann series is used to solve for total steady state power distribution in one step [10].
- 5) Power distributions at receivers are computed, represented by $\mathbf{w}_r$. This comprises:
 - a) Computation of power direct from sources (represented by their $\mathbf{w}^s$)
 - b) Computation of power from the boundary (represented by $\sum_i \mathbf{w}_i^+$)

Step 5 can also be applied to individual reflection orders $\mathbf{w}_i^+$, allowing the power density received from those to be visualised if that is of interest (e.g. for diagnostic purposes). This is done below.

The process above can be summarised by the following pair of matrix equations:

$$\mathbf{w}_i^- = \begin{cases} \mathbf{M}^{-1} \mathbf{T}_{sb} \mathbf{w}_s & \text{for } i = 0 \\ \mathbf{M}^{-1} \mathbf{T}_{bb} \mathbf{w}_{i-1}^+ & \text{for } i > 0 \end{cases}, \quad \mathbf{w}_i^+ = \mathbf{M}^{-1} \mathbf{R} \mathbf{w}_i^- \quad (10)$$

Here $\mathbf{T}_{sb}$ is a matrix describing source-to-boundary propagation and $\mathbf{T}_{bb}$ describes boundary-to-boundary propagation. $\mathbf{M}$ is the mass matrix. $\mathbf{R}$ is a reflectance matrix that describes boundary absorption and scattering. Boundaries are rigid in the examples herein, so $\mathbf{R} = \mathbf{M}$.

Total power distribution at a receiver is then given by:

$$\mathbf{w}_r = \mathbf{T}_{sr} \mathbf{w}_s + \mathbf{T}_{br} \sum_i \mathbf{w}_i^+ \quad (11)$$

Here $\mathbf{T}_{sr}$ is a matrix describing direct source-to-receiver propagation and $\mathbf{T}_{br}$ describes boundary-to-receiver propagation. This form directly matches that used for BEM in eq. 32 and 41 of ref. [18].

3. RESULTS

A rectangular room was simulated. It had dimensions 4m x 3m and was centred on the origin of the coordinate system. The source was located at [-1, -0.5] and the receiver at [1, 1]. The setup is shown in Figure 3b. A 3rd order source directivity was defined with $a_0 = a_1 = a_2 = 1$, $a_3 = 1/2$, and all $b_q = 0$. This gives $W_s(\theta_s)$ (via eq. (8)) as shown in Figure 3a. The receiver is order 5. Direct propagation to the receiver gives the power distribution shown in Figure 3c. Power received directly from the source was 0.762 W/m and main lobe points in the direction of the source, as expected.




Figure 3: a) Source directivity polar plot, b) Room geometry used with source and receiver locations, c) Received power directivity for direct sound

For the initial results, a maximum element size of 0.25m was used, giving $N = 56$ elements, and $Q = 17$ was used to discretise angle. This gave the system 952 degrees of freedom. Boundary integrals were evaluated by Gauss-Legendre quadrature and the $1/R$ singularity in eq. (6), which affects adjacent but non-coplanar elements, was regularised using the Sato transform [21].

Validation was against a cuboid room ISM code modified from the standard approach to include source and receiver directivity. This uses the same matrix mapping validated in Figure 3c but with image source directivity ‘flipping’ to respect reflection of sources in walls. A rigid boundary condition was assumed for simplicity. This, however, forces validation to be performed over finite reflection orders since neither algorithm converges to a finite steady state when energy is trapped.

3.1 Validation via Power Distribution on the boundary

To validate the codebase and provide insight into the boundary power distribution $w(\mu, \chi)$ and choice of basis functions, plots of it are presented below. Here $w(\mu, \chi)$ is shorthand for the left-hand side of eq. (7) and no discrimination is necessary between incoming and outgoing waves because the boundary is planar and rigid, meaning no absorption or scattering occur.

Figure 4b gives an example of such a plot for PWC basis in both position and angle and Figure 4a illustrates its meaning. Figure 4b pertains to the lower boundary depicted in Figure 4a. This has 16 elements, hence the horizontal coordinate μ in Figure 4b runs from 0 to 16. χ parameterises reflection angle θ via $\chi = \sin\theta$ and θ is depicted in Figure 4a. This has the range $-\pi/2 \leq \theta \leq \pi/2$ so $-1 \leq \chi \leq 1$. In ISM it is common to think only of the ray direct from the source to the receiver. This is shown towards the right of Figure 4a. It’s position and angle in boundary reflection space (μ, χ) are indicated by the yellow cross on Figure 4b. But the source doesn’t radiate only to the receiver – it radiates to all angles, hitting all positions on the boundary. This leads to cases where θ may be negative (green cross) or zero (blue cross) – these are also overlaid Figure 4b. Evaluating arrival angle for all positions on the edge produces the red line, onto which all three crosses fall. This is the loci of the source arrival direction. It can be seen that the PWC discretisation scheme in position and angle causes a 2D ‘faceting’ of the distribution into small rectangles in (μ, χ) space. The fading out around $\chi \approx 0$ is due to the choice of source directivity in Figure 3a.

This clearly shows approximation. The fine loci of the ISM method has been coarsely pixelized and its energy smudged across these. But this is not inappropriate for late-time responses, where the sound field becomes diffuse and individual reflections can no longer be distinguished. Here it illustrates that smaller element and higher Q cause less approximation since the ‘pixels’ are smaller.




Figure 4: Reflection of a source from a boundary: a) layout of a source, receiver and connecting rays in room; b) projection onto boundary (μ, χ) space with PWC basis funcs.

Figure 5 shows the same plot as Figure 4b but for all four edges of the room and for PWC and Legendre angular basis functions. Using PWC functions makes the solution on the boundary discontinuous, with data quantised into blocks. But Legendre polynomials are continuous over the whole range $-1 \leq \chi \leq 1$, yielding a series of smooth stripes. These exhibit some side-lobing in χ , as is typical of orthogonal polynomials, but the peak in the distribution is clearly where it is expected to be. Peak values are different due to the differing ‘smudging’ in the two schemes.




Figure 5: Mapping of reflection angles onto the four edges of the rectangular room using (a-d) Piecewise-Constant functions, and (e-h) Legendre polynomials, as angular basis functions

3.2 Validation via Power Directivity at Receiver

Figure 6 shows the receiver directivity patterns for first and second order reflections. Figure 6a only involves the source-to-boundary and boundary-to-receiver propagation processes $\mathbf{T}_{sb}$ and $\mathbf{T}_{br}$, but Figure 6b involves the boundary-to-boundary propagation process $\mathbf{T}_{bb}$ too, hence all three are validated. All three discretisation schemes match the ISM response well, Legendre and Chebyshev especially so. For these the error in total power (zeroth order directivity) was 0.1% for first-order reflections and 4.4% for second order. For PWC it was higher at 4.5% and 18.8% respectively.

That the Legendre and Chebyshev schemes gave identical results (to reported precision) was unexpected since ref. [11] reported that Chebyshev basis performed less well. To validate that this was not a coding error, the structure of the Chebyshev mass matrix was checked. Its version of the plots in Figure 5 (not shown) also showed slightly different sidelobes, proving it is different.




Figure 6: Received power directivity for: a) first order reflection, b) second order reflection

3.3 Convergence Study

In this section, element size and Q were varied to examine how the various schemes converge to the ISM result. The error metrics were normalised error versus ISM in total receiver power, as reported above, and a similar normalised RMS error over all the receiver coefficients. Thus, error is reported for receiver order 0 and receiver order 5.

Figure 7 shows results with Q fixed at 20 and element size varying from 0.1m to 1m. The Legendre and Chebyshev schemes again exhibit identical error trends. Error for PWC is again higher. Element size appears to have a fairly weak effect on accuracy, though convergence is visible in the RMS (directional to order 5) results for the Legendre and Chebyshev schemes. The PWC scheme shows little evidence of convergence in any case.




Figure 7: Normalised error (left) and RMS error (right) for varied element size and fixed $Q = 20$

Figure 8 kept the maximum element size constant at 0.2m and varied Q between 1 and 20. The Legendre and Chebyshev schemes show a clear convergence trend for the RMS (directional) results, including at second order. A weaker trend is also seen for total power. The PWC scheme, however, again shows little evidence on convergence in any case.




Figure 8: Normalised error (left) and RMS error (right) for varied Q and fixed elem. size (0.2m)

One reason for poor convergence of the PWC scheme might be its discontinuous nature. This affects its convergence in standard BEM. But accuracy could also be compromised by application of Gauss-Legendre integration to evaluate matrix entries for these discontinuous functions. The boundary basis functions are smooth over each element, but the angular kernel is not, which could compromise integration accuracy. Refs. [9,19] addressed this in 3D using spherical triangles.

4. CONCLUSIONS

A surface-based geometrical acoustics raytracing algorithm has been presented in Galerkin boundary integral form. The inclusion of the ‘mass’ matrix was seen to simplify comparison between different interpolation schemes. Moreover, the framework and notation emphasise synergy with a high-frequency BEM, which we see as being a route towards a unified algorithm for simulation of the full audible frequency band. The scheme was validated against the image source method for early reflections in a rectangular room with rigid boundary conditions.

Future work is to add support for absorbing and scattering boundaries and study convergence w.r.t. reflection order. Other basis functions, e.g., continuous piecewise linear, could be studied too. Later, it would be valuable to add corner diffraction to the scheme and investigate T60 decay rates.

5. REFERENCES

- [1] M. Vorländer, *Auralization : Fundamentals of Acoustics, Modelling, Simulation, Algorithms and Acoustic Virtual Reality*, 2nd ed., Springer, Cham, Switzerland, 2020.
- [2] L. Savioja, U.P. Svensson, Overview of geometrical room acoustic modeling techniques, *J. Acoust. Soc. Am.* 138 (2015) 708–730.
- [3] P.T. Calamia, B.E. Markham, U.P. Svensson, Diffraction Culling for Virtual-Acoustic Simulations, *Acta Acust. United With Acust.* 94 (2008) 907–920.
- [4] A. Erraji, J. Stienen, M. Vorländer, The image edge model, *Acta Acust.* 5 (2021) 17.
- [5] C. Schissler, G. Mückl, P. Calamia, Fast diffraction pathfinding for dynamic sound propagation, *ACM Trans. Graph.* 40 (2021) 1–13.
- [6] L.P. Franzoni, D.B. Bliss, J.W. Rouse, An acoustic boundary element method based on energy and intensity variables for prediction of high-frequency broadband sound fields, *J. Acoust. Soc. Am.* 110 (2001) 3071.
- [7] J.W. Rouse, L.P. Franzoni, Improvement of a High-Frequency Broadband Energy-Intensity Boundary Element Method to Include High Resolution Specular Reflection, *J. Comput. Acoust.* 13 (2005) 99–125.
- [8] H. Kuttruff, A simple iteration scheme for the computation of decay constants in enclosures with diffusely reflecting boundaries, *J. Acoust. Soc. Am.* 98 (1995) 288–293.
- [9] E.-M. Nosal, M. Hodgson, I. Ashdown, Improved algorithms and methods for room sound-field prediction by acoustical radiosity in arbitrary polyhedral rooms, *J. Acoust. Soc. Am.* 116 (2004) 970–980.
- [10] D.J. Chappell, G. Tanner, S. Giani, Boundary element dynamical energy analysis: A versatile method for solving two or three dimensional wave problems in the high frequency limit, *J. Comput. Phys.* 231 (2012) 6181–6191.
- [11] D.J. Chappell, G. Tanner, Solving the stationary Liouville equation via a boundary element method, *J. Comput. Phys.* 234 (2013) 487–498.
- [12] S. Siltanen, T. Lokki, S. Kiminki, L. Savioja, The room acoustic rendering equation, *J. Acoust. Soc. Am.* 122 (2007) 1624–1635.
- [13] S. Siltanen, T. Lokki, L. Savioja, Room Acoustics Modeling with Acoustic Radiance Transfer, *Proc. Int. Symp. Room Acoust.* 2010. (2010) 1–6.
- [14] J. Bajars, D.J. Chappell, N. Søndergaard, G. Tanner, Transport of phase space densities through tetrahedral meshes using discrete flow mapping, *J. Comp. Phys.* 328 (2017) 95–108.
- [15] A. Pohl, U.M. Stephenson, A Combination of the Sound Particle Simulation Method and the Radiosity Method, *Build. Acoust.* 18 (2011) 97–122.
- [16] J.A. Hargreaves, L.R. Rendell, Y.W. Lam, A framework for auralization of boundary element method simulations including source and receiver directivity, *J. Acoust. Soc. Am.* 145 (2019) 2625–2637.
- [17] J.A. Hargreaves, A comparison between the high-frequency Boundary Element Method and Surface-Based Geometrical Acoustics, in: *Proc. Internoise, Glasgow*, 2022.
- [18] J.A. Hargreaves, Y.W. Lam, The Wave-Matching Boundary Integral Equation — An energy approach to Galerkin BEM for acoustic wave propagation problems, *Wave Motion.* 87 (2018) 4–36.
- [19] J. Bajars, D. Chappell, N. Søndergaard, G. Tanner, Computing High-Frequency Wave Energy Distributions in two and three dimensions using Discrete Flow Mapping, in: *Proc. 22nd Int. Congr. Sound Vib., Florence*, 2015.
- [20] J. Rowbottom, D.J. Chappell, On hybrid convolution quadrature approaches for modeling time-domain wave problems with broadband frequency content, *Int. J. Numer. Methods Eng.* 122 (2021) 7581–7608.
- [21] M. Sato, S. Yoshiyoka, K. Tsukui, R. Yuuki, Accurate numerical integration of singular kernels in the two- dimensional boundary element method, in: C. Brebbia (Ed.), *Bound. Elem. X*, Vol. 1, Springer, Berlin, 1988: pp. 279–296.