Analytic Ray Curve Tracing
for Outdoor Sound Propagation

Qi Mo[a,*], Hengchin Yeh[a], Ming Lin[a], Dinesh Manocha[a]

[a]*Department of Computer Science, University of North Carolina at Chapel Hill, NC*

Abstract

Outdoor sound propagation, which propagates sound through inhomogeneous, moving media with complex obstacles, presents challenging scenarios for computational simulation. In this paper, we present a ray-tracing method that uses analytic ray curves as tracing primitives in order to improve the efficiency of outdoor sound propagation in fully general settings. This ray-curve tracer inherits the efficiency and flexibility of rectilinear ray tracers in handling boundary surfaces, and it overcomes the performance limitations imposed by approximating the curved propagation paths in inhomogeneous media with rectilinear rays. Adaptive media traversal, as well as acceleration structures for surfaces intersections, lead to further savings in computation. Our method's speedup over existing ray models, at least an order of magnitude for simple 2D scenarios and up to two orders of magnitude for 3D complex scenes, is demonstrated on outdoor benchmark scenes.

Keywords: Ray Tracing, Sound Propagation, Atmospheric Acoustics, Underwater Acoustics

1. Introduction

Sound propagation in outdoor environments[1, 2, 3] and room acoustics[4] are traditionally investigated separately. Room acoustics studies focus on the high-order interactions that sound waves make with surfaces within an enclosed

---

*Corresponding Author: Email, qmo@cs.unc.edu; Phone, (319)331-4931 (Qi Mo)

space, and assume a homogeneous medium. On the other hand, outdoor sound propagation (atmospheric and underwater) deals with media that are spatially varying and that move; they prominently differ from indoor models in the curved propagation trajectories that come from refraction. In the outdoor context, obstacles that constitute media boundaries are generally assumed to be sparse, and the complexity of the boundary surfaces is limited in the investigations.

Many outdoor applications can, however, benefit from simultaneous modeling of fully general media and complex boundaries. Examples include large expanses of complex terrain and sprawling urban areas, which often have their own micro-climate. As data on spatially- and temporally- varying outdoor media, such as the atmosphere and the ocean, becomes increasingly available, and geographic information on terrains and man-made structures is recorded in ever richer details, sound propagation models should ideally adapt to utilize the full scale of such data. Existing numerical methods suffer from scalability issues with large domain volume, high frequency, or complex boundaries. Geometric acoustics (GA)[5] methods are efficient at handling boundary surfaces, but with inhomogeneous media the curved trajectories become expensive to compute.

One way to speed up ray models for inhomogeneous media is to subdivide the media into spatial cells, so that within each cell the media assumes simple profile that leads to an analytic curved trajectory. The *cell method* was proposed as exactly such a model; based on a triangular subdivision of two-dimensional media, the cell method has been implemented in some acoustic simulation software (e.g. TRIMAIN[6], RAYWAVE[7]). A later work[8] adapted the idea of cell method to both light and sound propagation by decomposing three-dimensional media profiles into tetrahedral cells. This adaptation utilized mesh connectivity to trace rays across neighboring cells, and embedded boundary surfaces into the mesh so that surface intersections are also solved by the mesh-traversal process.

This paper further develops the cell method idea, and the algorithm improves upon the tetrahedral cell method[8] in three important aspects:

1. **Simple ray formulation** We select the parabolic ray curve as the ray-tracing primitive, which offers the simplest analytic form for trajectory,

intersection, and ray properties (Sec.3).

2. **Implicit cell** We use a mesh-less approach to media traversal, tracing ray curve segments of adaptive sizes based on on-the-fly sampling of the media profile. This implicit-cell approach avoids costly mesh construction, and it supports moving media as well as dynamic media (Sec.4.2).

3. **Acceleration of surface intersection** For the ray curve tracer, we adapt the hierarchical acceleration structures used in rectilinear ray tracers. Further acceleration is achieved by spatial bounding of ray curves based on their geometric properties, which offers higher culling efficiency (Sec.4.3).

Overall this analytic ray curve tracer is designed to be efficient for general media profiles and complex scenes with tens of thousands of surface primitives. Its performance is demonstrated on outdoor benchmarks (Sec.5), where it shows one to two orders of magnitude speedup over previous ray models. It is complementary to a set of numerical and geometric methods and can be extended in multiple ways (as discussed in Sec. 6).

2. Prior Work

Outdoor sound propagation has been studied extensively in underwater and atmospheric acoustics. Methods including the Fast Field Program (FFP)[9], Parabolic Equation (PE)[10, 11], and Normal Modes[1, 3] provide full-wave solutions that depend upon simplifying assumptions about the media and the obstacles in the scene. More general numerical methods, able to handle arbitrary media and obstacles, have also been proposed; the leading one is the Finite Difference Time Domain(FDTD) method, which solves the linearized Euler equation[12, 13]. The main disadvantage of the FDTD is its high computational cost, and its poor scalability with the volume of the simulation domain or with the frequency. Other numerical methods, including Finite Element Method[14, 15], Boundary Element Method[16], Pseudo Spectral Time Domain[17, 18], and Transmission Line Matrix (TLM)[19, 20, 21], face similar challenges in scalability.

Geometric acoustics (GA)[5] methods are widely used in room acoustics[4]. GA methods rely on the high-frequency approximation of wave propagation using rays, which are perpendicular to the wave fronts and can be intersected efficiently with scene surfaces. Various GA techniques have been developed, including image source method[22, 23], ray tracing[24, 25], frustum tracing[26], beam tracing[27, 28], sonal mapping[29, 30, 31], path tracing[32], and radiosity[33]. In these methods, sound speed is generally assumed to be constant, so rectilinear rays are used to trace out straight-line paths. Complex surface interactions, such as specular reflection, Snell's Law refraction, diffuse reflection [25, 32], and Bidirectional Reflectance Distribution Function (BRDF)[31, 34], are modeled.

Ray models have also been applied to inhomogeneous outdoor media by numerical integration of the differential ray equation (see survey [1, 3]). The integration step size is limited by the range within which the media can be assumed to be homogeneous, and intersection with surfaces need to be tested for each segment. With complex media the great number of small ray segments dominate the propagation computation. Higher-order integration, like the fourth-order Runge-Kutta method, have been adopted to improve efficiency, but with those methods the trajectory between each advancement of the ray integration has no explicit form, complicating intersection tests with surfaces (see details in Fig. 4). Several widely used software tools are based on ray models, including HARPO[35], BELLHOP[36], WaveQ3D[37] in underwater acoustics, HARPA[38] in atmospheric acoustics (see [39] for the collection). Although performance of these ray-based tools is orders of magnitude better than wave-based alternatives (reported recently[40, 41], a few seconds vs. thousands of seconds on 2D propagation in layered media), these methods quickly become prohibitively expensive when dealing with general media and 3D complex scenes.

3. Analytic Ray Curve

Ray models simulate wave propagation with rays that are perpendicular to the wavefront. For certain media profiles, the ray trajectories, as well as the

4
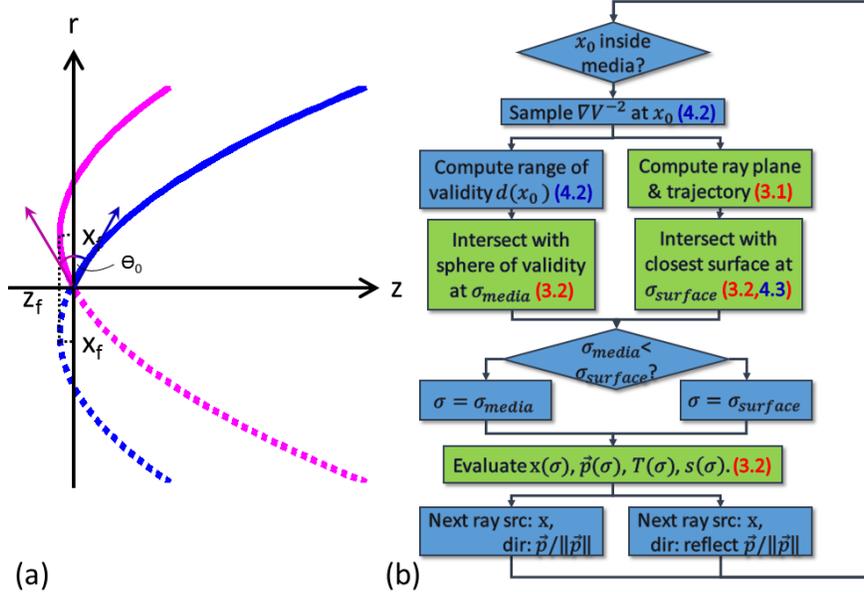
Figure 1: **(a) Parabolic ray curve.** Two rays with different initial directions (marked by arrows) are shown in the *ray plane*, with z-axis being the direction of ∇V^{-2}. (x_f, z_f), which defines the vertex of the parabola, is computed according to Eq. 6. **(b) Ray curve tracing.** The green-colored blocks have analytic solutions with the parabolic ray curve, which offer large savings in ray-tracing costs. For each key step, the relevant sections that give the analytic formulation are marked in red, and the relevant sections that describe algorithm details are marked in blue.

pressure on the ray, have analytic forms that can be evaluated at any point along the ray at constant cost. Our ray tracer utilizes one of these analytic formulations (Fig. 1(a)), chosen for its particularly simple form, and the resulting ray curve lends itself to easy intersections with boundary surfaces. As can be seen from Fig. 1(b), this particular ray curve enables analytic evaluations at a constant per-ray cost in key steps of the ray tracing (green-colored blocks). This ray formulation is a special case of the more general ray theory discussed in depth in Červený's comprehensive work[42].

3.1. Ray formulation

Within a medium with spatially-varying sound speed, we denote $V(\mathbf{x})$ as the sound speed at location $\mathbf{x}$. The *arrival time* or *travel time* field T is defined

as the time that a sound wave takes to travel from its source to a field location, and the spatial derivative of T: $\vec{p} = \nabla T$ is called the *slowness vector*. Within such a medium, the ray trajectories can be solved from the general *Hamiltonian* form of the Eikonal equation, which in Cartesian coordinates reads:

$$\mathcal{H}(x_i, p_i) = n^{-1}\{(p_i p_i)^{n/2} - V^{-n}\}, i = 1, 2, 3, \tag{1}$$

with n as a real number corresponding to different forms of the equation. Eq. 1 can be solved in terms of *characteristics*: 3-D trajectories $x_i = x_i(u)$ parameterized by u, along which 1 is satisfied, as follows:

$$\frac{\mathrm{d}x_i}{\mathrm{d}u} = (p_k p_k)^{n/2-1}p_i, \qquad \frac{\mathrm{d}p_i}{\mathrm{d}u} = \frac{1}{n}\frac{\partial}{\partial x_i}(V^{-n}), \tag{2}$$
$$\frac{\mathrm{d}T}{\mathrm{d}u} = (p_k p_k)^{n/2} = V^{-n}.$$

When the gradient of V^{-n} is constant, 2 can generally be solved analytically. When we assume that the gradient of V^{-2} is constant (i.e. $n = 2$: $V(\mathbf{x})^{-2} = A_0 + \vec{\mathbf{A}} \cdot \mathbf{x}$), we obtain the simplest analytic form for ray trajectory and travel time, in terms of the parameter $u = \sigma$:

$$x_i(\sigma) = x_{i0} + p_{i0}(\sigma - \sigma_0) + \frac{1}{4}A_i(\sigma - \sigma_0)^2, \tag{3}$$

$$p_i(\sigma) = p_{i0} + \frac{1}{2}A_i(\sigma - \sigma_0), \tag{4}$$

$$T(\sigma) = T(\sigma_0) + V_0^{-2}(\sigma - \sigma_0) + \frac{1}{2}A_i p_{i0}(\sigma - \sigma_0)^2 + \frac{1}{12}A_i A_i(\sigma - \sigma_0)^3. \tag{5}$$

σ is related to the travel time T and to arclength s by $\mathrm{d}\sigma = V^2\mathrm{d}T = V\mathrm{d}s$.

3.2. Ray properties

We take advantage of two key properties of the ray curve given in Sec. 3.1:

1. **Intersection** The ray is a parabolic plane curve and can be intersected analytically with planar surfaces to solve for the intersection point.

2. **Evaluation** For an arbitrary parameter σ along the ray, the corresponding position, tangent direction, and the ray properties contributing to pressure

6

computation, including the slowness vector, arrival time, and arclength, can be evaluated analytically at constant cost.

It is evident from Eq. 3 that the analytic ray lies in a plane with the normal of $\vec{A} \times \vec{p_0}$. We denote the direction of media gradient $\vec{A}$ as z-axis and $(\vec{A} \times \vec{p_0}) \times \vec{A}$ as r-axis; the r-z plane is then called the *ray plane*. When the origin of the coordinate system is placed at ray origin $\mathbf{x_0}$, the ray trajectory is:

$$z(r) = \frac{\alpha}{4\xi_0^2}(r - r_f)^2 - z_f, \qquad r_f = -\frac{2\xi_0\sqrt{-\xi_0{}^2 + n_0{}^2}}{\alpha}, z_f = \frac{n_0{}^2 - \xi_0{}^2}{\alpha}, \quad (6)$$

where $\alpha = \|\vec{A}\|$ is the magnitude of media gradient, $\xi_0 = V(0)^{-1}cos\theta_0$, $V(0)$ is the sound speed at ray origin, and θ_0 is the angle between the ray direction at origin and the r axis (see Fig. 1). The parabolic plane curve given by Eq. 6 has closed-form solutions for its intersections with planar surfaces.

For an arbitrary point on the ray curve with coordinates (r_p, z_p) in the ray plane, the corresponding parameter σ_p can be solved from Eq. 3 as $\sigma_p = r_p/\xi_0 + \sigma_0$, where σ_0 is the σ at ray origin. Given σ_p, we can calculate the position $\vec{x}(\sigma_p)$, the slowness vector $\vec{p}(\sigma_p)$, and travel time $T(\sigma_p)$ from Eq. 3,4,5, respectively. The tangent direction $\vec{t}$ coincides with the direction of the slowness vector $\vec{p}$ and can be evaluated by $\vec{t}(\sigma_p) = \vec{p}(\sigma_p)/\|\vec{p}(\sigma_p)\|$. The arclength s of the parabolic ray curve can be evaluated at (r_p, z_p) by

$$s(r_p) = \left.\frac{\sqrt{(ar_p + b)^2 + 1}(ar_p + b) + \text{arcsinh}(ar_p + b)}{2a}\right|_0^r, \qquad (7)$$

$$\text{where } a = \frac{1}{2}\frac{\alpha}{\xi_0^2}, \quad \text{and } b = \frac{\sqrt{-\xi_0 + V(0)^{-2}}}{\xi_0'}.$$

The arclength s determines the attenuation along the ray path, and it is one of the key variables that we measure for accuracy of the ray tracing (see Sec. 5.1).

Similar to rectilinear rays, the parabolic ray curve comes with efficient operations for finding intersection and for updating the key properties at intersection points, making it amenable to be used as a tracing primitive. The ray curve serves as a building block for a ray tracer that handles general media; it does

so by computing the propagation paths, which consist of consecutive segments of the aforementioned ray curves. The flow chart in Fig. 1(b) illustrates the steps involved in tracing one such segment, with the computational savings highlighted by green-colored blocks. In the next section, we give the details of our ray-tracing algorithm using the parabolic ray formulation.

4. Ray Tracing Algorithm

4.1. Tracing Analytic Segments

Our ray tracer computes the propagation paths in a segment-by-segment fashion, with each segment being a parabolic ray curve (formulated in Sec. 3). There are two criteria for terminating one segment and starting the next one:

1. The media no longer satisfy the assumption of constant ∇V^{-2},
2. The boundary surfaces in the scene are encountered.

For the former criterion, the change in the media gradient ∇V^{-2}, we propose an adaptive segment size that is computed from on-the-fly sampling of the media; this is discussed in detail in Sec. 4.2. For the latter criterion, in which the segment encounters a surface, we use acceleration data structures and methods adapted from rectilinear ray tracer; these are discussed in detail in Sec. 4.3. Both aspects of the ray-tracing algorithm contribute to its efficiency and offer performance improvements over previous methods (see Fig. 2 and Fig. 3).

Based on these criteria for ending segments, we find out how far a particular segment will go in two ways. For the surface-encounter criterion, we compute the closest intersection point; for the media gradient criterion, we compute the point at which the ray curve reaches the edge of the *range of validity*. The closest point between the two cases is taken as the end point of the current segment, and σ, $\mathbf{x}(\sigma)$, $\vec{t}(\sigma)$, $T(\sigma)$, $s(\sigma)$ are evaluated at this point. The next segment starts at $\mathbf{x}(\sigma)$, and its trajectory is computed from the local media gradient $\nabla V^{-2}(\mathbf{x}(\sigma))$. σ, $T(\sigma)$, $s(\sigma)$ are continued over the next segment; the ray direction $\vec{t}(\sigma)$ is also continued, except in the case of the surface encounter.

8
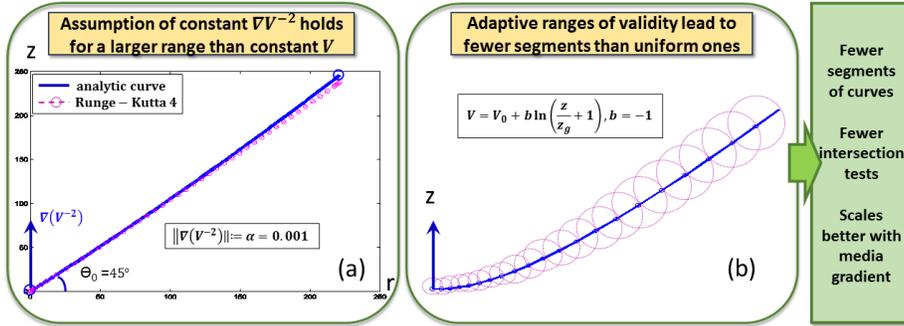
Figure 2: **Efficient media traversal** (Sec. 4.2). **(a)** With constant $\|\nabla V^{-2}\|$, one segment of the analytic curve (blue line) suffices, as compared to the many steps (magenta circles) taken by numerical ray integration. The trajectory traced out by RK-4 visibly diverges from the parabolic curve; this can be improved by further reducing step sizes; **(b)** With a logarithmic sound speed profile, our ray tracer proceeds in segments of adaptive-sized curves, evident from the spheres of validity drawn in magenta.

In the case of the surface encounter, we compute a reflection direction $\vec{t}'$ as the initial direction of the next segment, based on the surface-reflection model.

The ray traversal continues in this segment-by-segment manner until one of three things happens: it reaches outside the media volume; it interacts with surfaces a pre-determined number of times; or the pressure amplitude along the ray drops below a pre-set threshold. To compute the pressure field, the ray tracer finds the set of propagation paths starting from the sound source, the ray segments making up those paths can be queried for any field location.

4.2. Adaptive Media Traversal

Within a smoothly varying medium, the gradient ∇V^{-2} can be assumed to be constant within a certain range of validity. This assumption is valid for naturally occurring media under stable conditions, such as the atmosphere or the ocean. Furthermore, our ray tracer uses the heuristic that the range of validity for a locally constant ∇V^{-2} should be proportional to $\|\nabla V^{-2}\|$, so that the constant gradient remains valid for a relatively small spatial range in areas of great media variation. Such adaptive sizing facilitates accurate traversal of the media without the cost incurred by a uniformly small segment size.

Based on the aforementioned heuristic, we compute the range of validity $d(\mathbf{x})$

9
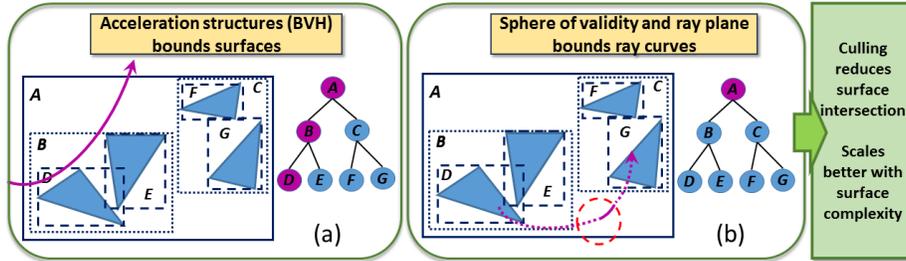
Figure 3: **Efficient surface intersection** (Sec. 4.3). **(a)** A parabolic ray (magenta arrow) is tested against surfaces and their bounding volumes, shown on the left, and the corresponding BVH is shown on the right. Only the magenta nodes are traversed. **(b)** A ray segment is bound by the ray plane and sphere of validity, which provides further culling against scene surfaces. Here only the root node of the BVH is traversed.

by $\delta = \frac{1}{8}\|\nabla V^{-2}\|d^3(\mathbf{x})$, controlled by a global parameter δ. The ray curve given the local $\nabla V^{-2}(\mathbf{x})$ is computed from Eq. 6 and continued until it intersects the *sphere of validity*, centered at the ray origin $\mathbf{x}$ and with a radius $d(\mathbf{x})$. The parabolic curve has a closed-form intersection with this sphere, and the exit point $\mathbf{x}'$ from the sphere is used to sample the media for the next segment.

Given an input media profile $V(\mathbf{x})$, on-the-fly sampling of $\nabla V^{-2}(\mathbf{x})$ can be computed either analytically (if the input profile is available as an analytic function), or by differentiation techniques, such as central differences. Moving media can be approximated by the standard *effective sound speed* approach (adding to or subtracting the wind velocity $w(\mathbf{x})$ from $V(\mathbf{x})$ for upwind or downwind propagation respectively). Particularly for our ray tracer, the on-the-fly media-gradient sampling enables us to handle the vector wind field by adding the wind speed projected onto the ray direction $V(\mathbf{x}) + \vec{w}(\mathbf{x}) \cdot \vec{t}(\mathbf{x})$.

4.3. Handling Boundary Surfaces

For rectilinear ray tracers, the most expensive part of the computation is computing ray intersections with complex surfaces within the scene; many acceleration structures and techniques have been developed to speed up this process. We adopt the hierarchical acceleration structures from rectilinear ray tracers that bound the scene surfaces. The parabolic ray curve as our tracing primitives can be used directly with these structures to cull non-intersecting surfaces.
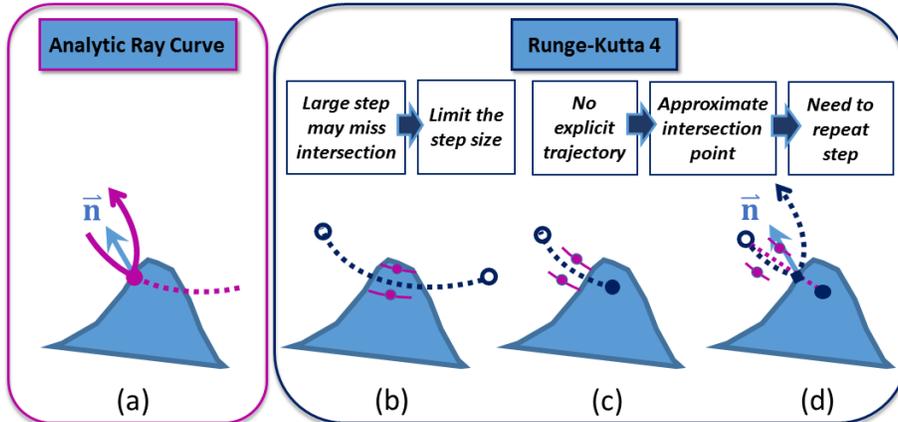
Figure 4: **Ray curve vs. RK4 on surface intersection: (a)** with the ray curve (in magenta), intersection with obstacles (blue block) is efficient and robust; **(b,c,d)** with higher-order numerical methods (blue dotted lines with the 2 mid-points in RK-4 drawn in magenta), challenges include: **(b)** step size is limited, **(c)** intersection point can only be approximated, and **(d)** after the intersection point is approximated (assuming linear trajectory here as shown by the magenta line), step needs to be repeated adjusting for the intersection distance. The poor performance of RK4 on scenes with complex surfaces (Table 3) reflect these challenges.

For scene-surfaces culling, we built a Bounding Volume Hierarchy(BVH)[43], which groups boundary surfaces and bounds them spatially. A BVH is a hierarchical tree structure whose internal nodes are bounding volumes(BVs); for the shape of the nodes, we choose axis-aligned bounding boxes(AABBs). The parabolic ray curve can be intersected with an AABB much faster than it can with all the AABB's enclosed surfaces, and when the ray curve and the AABB are not intersecting, all those surfaces can be culled from further computations. To find the closest intersection with the scene surfaces, the BVH is traversed by each ray curve in a top-down fashion, achieving on average a logarithmic intersection time with regard to the number of surfaces in the scene. The BVH can also be efficiently re-fitted or rebuilt to accommodate dynamic scenes.

Besides spatial culling of the scene surfaces with a BVH, we also perform spatial bounding and culling of ray curves. As we traverse the media (as described in Sec. 4.2), we compute a bounding sphere (the sphere of validity) for each segment as a by-product. The culling efficiency is improved when inter-

secting such a segment with the scene BVH by first intersecting the bounding sphere of the segment with the BVs. And since the parabolic ray is a plane curve, we use the ray plane to reduce the dimension of the intersection test.

At the intersection point with a boundary surface, the interaction can be modeled in a variety of ways: by specular reflection, by Snell's law refraction, or by BRDF-based reflection, depending upon the surface properties. We can then compute the direction of the secondary ray segment accordingly. In addition to modifying the ray direction, a surface impedance model can be used to modify the amplitude and the phase of the pressure along the ray.

5. Results

5.1. Performance Comparison of Ray Models

Ray models are the foundation of many software tools in underwater acoustic simulation[6, 7, 36, 37], and in atmospheric acoustics[38]. The majority of the available tools assume media profiles that vary in at most two dimensions, with the exception of HARPO and HARPA, which computes three-dimensional Hamiltonian ray tracing on clusters. The tools are written in languages with different performance characteristics (e.g. MATLAB, Fortran). We reviewed the relevant data in Sec. 2; in this section we compare our method's performance to that of a numerical ray integration implemented in our code base. Platform-neutral results such as number of ray steps are reported in place of raw running time, in order to highlight the algorithmic difference in comparison.

Existing ray models perform direct numerical integration of the ray equations, using either the Euler method or a higher-order method, such as the fourth order Runga-Kutta (RK4). In these methods, the size of the integration timestep is an important parameter that controls the trade-off between performance and accuracy. Similarly, our ray tracer's balance between performance and accuracy depends on the global parameter δ (Sec. 4.2), which controls step size. As in the numerical-integration methods, smaller step sizes produce more accurate results, at the cost of longer computation time. To highlight this trade-

12
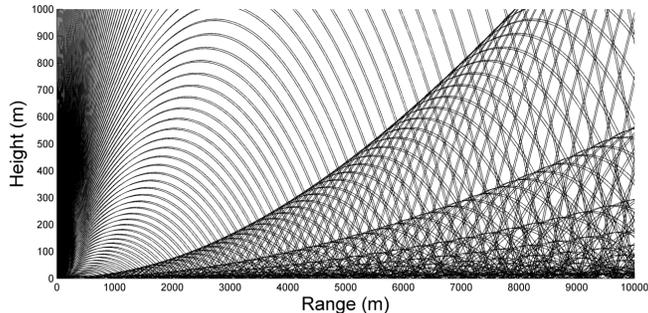
Figure 5: **Downward refracting atmosphere:** computed by the ray curve tracer.

off, we conduct two types of comparisons: same-quality comparisons, for which ray models are tuned to achieve the same level of error, then compared in terms of performance; and same-speed comparisons, in which ray models are run at the same level of performance, then compared in terms of resulting errors.

We measure the accuracy of our ray tracer by the two variables: hit location of the ray at a particular range and the arclength of the entire trajectory. These variables characterize the shape and size of the computed propagation path, which determine the location to which the path contribute and the phase and attenuation along the path for the pressure field computation. To quantify the difference between two sets of results, rt and $\tilde{rt}$, tracing the same set of initial ray directions, we compute the relative error: $E_{rel} = \frac{\|rt - \tilde{rt}\|}{\|rt\|}$.

We select two media profiles as test cases: a downward-refracting atmosphere with a sound-speed gradient of 0.1, and the Munk profile from ocean acoustics. Both profiles have ∇V^{-2} and ∇V in analytic form, which decouple the measurements of ray-tracing accuracy from the accuracy of gradient estimation. We also limit our test to 2D cases with only flat ground as the obstacle, so as to fairly compare our ray tracer's performance to that of the majority of existing ray-based software. The performance improvements for 3D scenes with complex obstacles is expected to be greater, as demonstrated in Section 5.2.

13

The first media profile we tested is a linear sound speed profile that leads to a downward-refracting atmosphere: $c(z) = c(0) + 0.1 m/s$, where $c(z)$ is the effective sound speed at the height z, and $c(0)$ is the effective sound speed on the ground. We compute the ray-tracing results for the range $0 - 10 km$ and the height $1 - 1 km$. The ray-tracing results are gathered and compared at the $10 km$ mark, a point at which the rays have been reflected up to 25 times. An example of the computed ray trajectories can be seen in Fig. 5, which shows a fan of 201 rays between the elevation angle -90˚and 90˚.

Rel. Error	1e-2	5e-3	1e-3	5e-4	2e-4	
Ray Curves	219	892	2,810	3,972	8,881	(segments/path)
RK4 Steps	467	2,330	4,881	24,588	49,158	(steps/path)

Table 1: **Same-quality comparison**. A fan of rays is traced under a downward-refracting atmosphere to a range of 10km. At each level of relative error in the ray-tracing results, the average number of ray-curve segments per propagation path is compared with the average number of Runga-Kutta 4 steps. The analytic ray curve is able to achieve same level of accuracy with a much smaller tracing cost.

# Ray steps/ RK4 segments	1×10^2	5×10^2	1×10^3	5×10^3	1×10^4
RT rel. error	1.31e-2	2.76e-3	8.04e-4	1.98e-4	4.33e-5
RK4 rel. error	4.35e-2	9.44e-3	4.92e-3	2.14e-3	9.12e-4

Table 2: **Same-speed comparison**. A fan of rays is traced under a downward-refracting atmosphere to a range of 10km. With a comparable number of ray curve segments and Runga-Kutta 4 steps, our ray tracer (RT) is able to achieve lower levels of relative error in ray-tracing results (both hit points and path length) than RK4.

The same-quality and the same-speed comparisons between our ray tracer and the RK4 ray integrator are listed in Table 1 and 2. To achieve the same error level, our ray tracer runs an order of magnitude faster than the numerical integrator. And with the same running time, the numerical integrator produces results with higher levels of error. For this comparison only the computation of ray trajectories are included, the pressure computation is excluded, and the intersection computation is minimum with only a flat ground. The higher number

of ray steps used by the numerical integrator indicate that, if it had to deal with complex intersection tests and additional pressure-related computations at each step, the performance improvements of our ray tracer would be even higher.

We also test the scalability of the ray models with increasing magnitudes of media gradient in Fig. 6 left. The media gradient is increased gradually from 0.04 to 0.4 with 0.04 increments, and the relative error is kept at 2×10^{-4} by tuning the step size for RK 4 and the parameter δ for our ray curve tracer. Our ray tracer is shown to scale better with increased media variations than RK 4.

5.1.2. Munk profile

The second media profile we tested is the Munk profile commonly used in ocean acoustics: $c(z) = 15001.0 + 0.00737[x - 1 + exp(-x)]$, where $c(z)$ is the effective sound speed at the depth z, $z \leq 5000m$. We compute the ray-tracing results for the range $0 - 100km$ and the depth $0 - 5km$. An example of the computed ray trajectories can be seen in Fig. 8, which shows a fan of 21 rays between the elevation angle -13°and 13°. The relative errors in the results from the numerical integrator with decreasing step sizes are plotted in Fig. 6; compare this to the same cost-accuracy analysis results for our analytic ray tracer. Again, for this comparison only the computation of ray trajectories and arclength are included, and the pressure computation is excluded.

For this profile we also show the adaptive segment size adopted by our ray tracer by plotting the segment size with depth within the profile (Fig. 7). We trace rays spanning the launch angles between -90°and 90°to record the step sizes adopted, and we contrast this with a plot of sound speed of the Munk profile. The Munk profile's local coherence of media properties is taken advantage of by our adaptive ray tracer, to focus computation and improve efficiency. Even with the analytic ray formulation, a uniform step size (instead of an adaptive one) will lead to significant performance degradation.
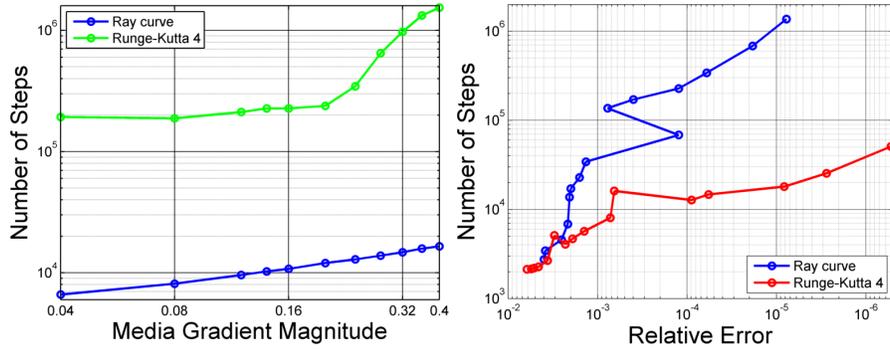
Figure 6: **Left: downward-refracting atmosphere.** As the magnitude of media gradient increases, both RK 4 and ray curve tracer adopt smaller steps to keep the error below 2×10^{-4}; RK 4 shows both higher cost and lower scalability. **Right: Munk profile** The cost-accuracy trade-off of Runga-Kutta 4 and ray curve tracer. No adaptive step/segment size is used for this comparison.
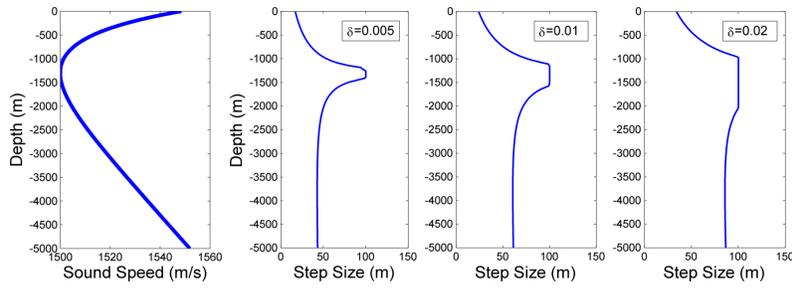
Figure 7: **Adaptive segment size.** The range of validity for each segment of ray curve is controlled by the parameter δ and adapts to local media gradient. **(a)**Munk profile sound speed. **(b,c,d)** Segment sizes plotted with depth for the Munk profile, for $\delta = 0.002, 0.001, 0.0005$, respectively. Without adaptive sizes, the required number of ray segments will be 3 to 10 times higher to achieve comparable accuracy.
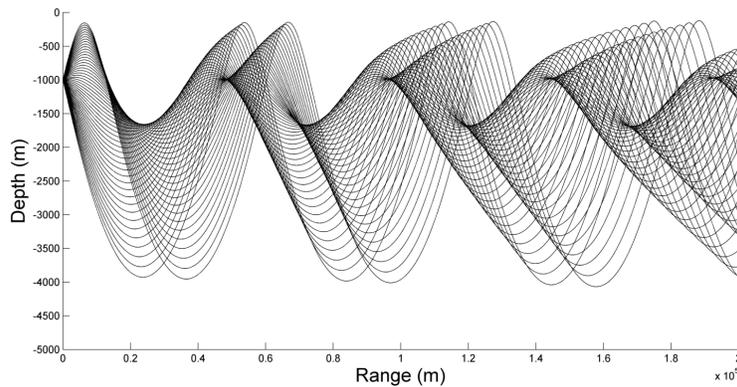
Figure 8: **Munk profile** characteristic ray plot computed by the ray curve tracer.
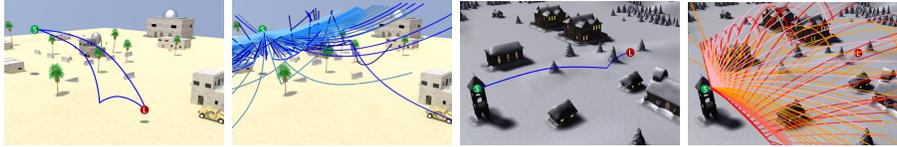
Figure 9: **Outdoor benchmarks** Desert and Christmas village, upward and downward refraction, respectively. For purposes of illustration, the media gradient is exaggerated and the ray path number is kept very small. See Table 3 for scene stats and actual performance numbers when tracing 10k rays up to 3 orders of reflections.

5.2. Outdoor Applications of Analytic Ray Tracer

To show the application of our ray tracer on fully general outdoor scenes, we generate a representative media profile from the stratified-plus-fluctuation model that is widely used in atmospheric acoustics[1]. The atmospheric acoustic index-of-refraction $n = c_0/c$, where c_0 is the reference sound speed, is modeled with a stratified component n_{str} and a fluctuation component n_{flu} so that $n = n_{str} + n_{flu}$. The stratified component follows a logarithmic profile of the altitude z: $n_{str}(z) = c_0/(c_0 + b \ln\left(\frac{z}{z_g} + 1\right))$, with parameters n_0, b, and z_g. c_0 is the sound speed at the ground surface, taken as the reference sound speed, and z_g is the roughness length of the ground surface. Typical values for parameter b are 1 m/s for a downward-refracting atmosphere and -1 m/s for an upward-refracting atmosphere. The fluctuation component models the three-dimensional random temperature variation and wind speed turbulence in the atmosphere. The value at position $\mathbf{x}$ can be computed as $n_{flu}(\mathbf{x}) = \sum_i G(\mathbf{k}_i) \cos(\mathbf{k}_i \cdot \mathbf{x} + \varphi_i)$[1], where $\mathbf{k}_i$ is the wave vector describing the *spatial* frequency of the fluctuation, φ_i is a random angle between $[0, 2\pi]$, and $G(\mathbf{k_i})$ is a normalization factor.

The Desert and Christmas benchmarks represent large-volume outdoor acoustic scenes that have complex surface geometry (e.g. terrains and buildings). The surface-primitive count and the volumetric expanse of the two scenes are listed in Table 3. Both scenes are visualized in Fig. 9, and we include illustrative ray paths for both upward-refractive and downward-refractive media conditions.

For each benchmark we trace 10K rays for up to three orders of reflections, and report the overall ray-tracing performance alongside a breakdown of the

17
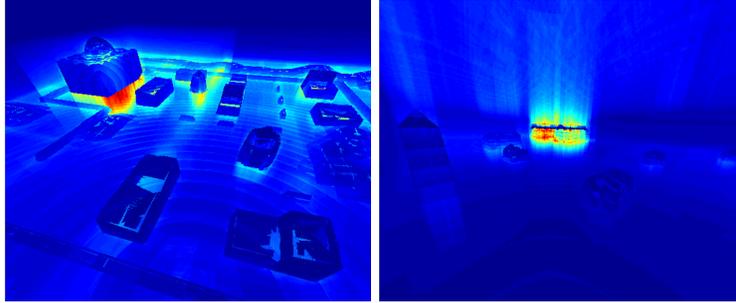
Figure 10: **Visualization of ray-surface intersection cost:** brighter colors indicate higher costs. It is shown that the computation concentrates in areas of dense surface geometry and around obstacle silhouettes, illustrating that the spatial culling of ray-surface intersections using BVH is effective. **Left:** Desert **Right:** Christmas village.

Scenes (m^3)	Surf. count	Avg. # segments	BV test	Surf. test	Frame time	RK4 time
Desert $(200 \times 200 \times 50)$	5,000	26.73 per path	31.88 (94%)	2.32 (1.47%)	446 ms	77 s
Christmas $(120 \times 80 \times 50)$	16,000	44.45 per path	34.31 (96.16%)	2.08 (1.22%)	917 ms	128 s

Table 3: **Breakdown of ray-tracing time:** the computation of analytic ray trajectories and ray properties takes less than 0.001% of frame time and thus is omitted here. We report the number of ray-BV and ray-surface intersection tests, and the respective percentage of frame time they take. The frame times of the analytic ray curve tracer are compared to the running time of Runge-Kutta 4 on the same scenes.

running times, as shown in Table 3. We see that the ray intersection computation dominates the ray-tracing cost, and that the majority of the intersections are with Bounding Volumes (BVs). The number of intersections with surface primitives is kept very low, which shows that the spatial culling using BVH is effective. A detailed visualization in Fig. 10 demonstrates this point further by showing that the intersection cost is concentrated in the vicinity of obstacles, and that for the majority of the scenes the ray intersections are resolved with relatively low cost. The analytic ray tracer achieves close to interactive performance for both benchmarks, which is two orders of magnitude faster than RK-4, and it scales well with the complexity in both media and the boundaries.

6. Discussions

Our ray tracer can be extended in ways including but not limited to:

- coupling with flow models and using realistic media profiles as inputs;
- augmenting GA methods to account for inhomogeneous media;
- combining with other methods to form hybrid propagation algorithms.

The analytic ray curve tracer takes as input any general media profile, as long as that profile can be sampled for sound speed and gradient at any spatial location. The input can be an analytic function (e.g. the Monin-Obukhov similarity model[44, 45]) or a set of discrete sample points (e.g. a grid). Numerical methods have been coupled with flow models[46, 47, 48], and the detailed media conditions produced by the flow model as output can then serve as input for propagation simulation. Our ray tracer is also compatible with such a coupling.

GA methods model complex surface interactions[32, 31] such as diffuse reflections and surface BRDFs, assuming homogeneous media. By substituting our curved trajectory computation for the GA assumption of straight-line propagation, we can account for inhomogeneous media while still utilizing GA's surface handling techniques. Such augmented GA solutions can be applied to large indoor spaces such as auditoriums, and to dense outdoor urban areas.

The efficiency of our method makes it a good candidate for forming hybrid schemes with wave-based models. One possibility is to couple the ray curve tracer with Equivalent Source Method(ESM), similar to what Yeh et al.[49] proposed. The ray curve tracer could be focused on high-frequency and long-range propagation while ESM handles the rest. Alternatively, our ray tracer could be used for fast initial evaluations, and costly numerical methods such as FDTD can be applied economically to areas of interest [50, 51, 52].

The analytic ray curve tracer has several limitations. This method is based on the geometric ray model, which is a high-frequency approximation rather than a full wave solution. The efficiency depends on media coherence; if the media is chaotic and lacking spatial coherence, our method's performance will degenerate to about the same level as the numerical ray integration.

19

7. Conclusion

This paper addresses the challenge of efficient sound propagation in inhomogeneous, moving media and in large outdoor scenes with complex boundary surfaces. Our method adapts the geometric ray models to inhomogeneous media; its key performance improvement upon rectilinear rays is achieved by tracing as primitives analytic ray curves computed from local media gradients. Segments of ray curves are computed by sampling the media gradient on-the-fly, enabling us to account for both inhomogeneous and moving media without needing to precompute explicit cell structures. Acceleration based on BVH is readily adapted to speed up surface intersections of the ray curves, which enables logarithmic scaling with scene complexity and allows for dynamic scenes. The performance of this method is demonstrated in comparison to existing ray models, as well as on scenarios that are prohibitively expensive with existing methods. In addition to being an efficient stand-alone tool, this method also has the potential of complementing other GA and numerical methods.

[1] E. M. Salomons, Computational atmospheric acoustics, Kluwer Academic Publishers, 2001.

[2] K. Attenborough, K. M. Li, K. Horoshenkov, Predicting outdoor sound, CRC Press, 2006.

[3] F. B. Jensen, W. A. Kuperman, M. B. Porter, H. Schmidt, Computational ocean acoustics, Springer, 2011.

[4] H. Kuttruff, Room acoustics, CRC Press, 2009.

[5] A. D. Pierce, Acoustics: an introduction to its physical principles and applications, McGraw-Hill, 1981.

[6] B. G. Roberts, Horizontal-gradient acoustical ray-trace program TRIMAIN (1974).

[7] W. H. Watson, R. W. McGirr, RAYWAVE II: A propagation loss model for the analysis of complex ocean environments (1975).

[8] Q. Mo, H. Yeh, D. Manocha, Tracing analytic ray curves for light and sound propagation in non-linear media, arxiv preprint arxiv:1409.2235 (2014).

[9] F. R. DiNapoli, R. L. Deavenport, Theoretical and numerical Green's function field solution in a plane multilayered medium, Journal of Acoustical Society of America 67 (1) (1980) 92–105.

[10] M. K. Myers, G. L. McAninch, Parabolic approximation for sound propagation in the atmosphere, Journal of AIAA 16 (8) (1978) 836–842.

[11] K. E. Gilbert, M. J. White, Application of the parabolic equation to sound propagation in a refracting atmosphere, Journal of Acoustical Society of America 85 (2) (1989) 630–637.

[12] V. E. Ostashev, D. K. Wilson, L. Liu, D. F. Aldridge, N. P. Symons, Equations for finite-difference, time-domain simulation of sound propagation in moving inhomogeneous media and numerical implementation, Journal of Acoustical Society of America 117 (2) (2005) 503–517.

[13] D. K. Wilson, L. Liu, Finite-difference, time-domain simulation of sound propagation in a dynamic atmosphere, Cold Regions Research and Engineering Lab, Hanover NH (No. ERDC/CRREL-TR-04-12.).

[14] T. Nomura, K. Takagi, S. Sato, Finite element simulation of sound propagation concerning meteorological conditions, International Journal for Numerical Methods in Fluids 64 (10-12) (2010) 1296–1318.

[15] C. P. Vendhan, G. C. Diwan, S. K. Bhattacharyya, Finite-element modeling of depth and range dependent acoustic propagation in oceanic waveguides, Journal of Acoustical Society of America 127 (6) (2010) 3319–3326.

[16] C. A. Brebbia, Boundary element methods in acoustics, Computational Mechanics Publications, 1991.

[17] M. Hornikx, R. Waxler, J. Forssén, The extended Fourier pseudospectral time-domain method for atmospheric sound propagation, Journal of Acoustical Society of America 128 (4) (2010) 1632–1646.

[18] M. Hornikx, R. Waxler, The extended Fourier pseudospectral time-domain (pstd) method for fluid media with discontinuous properties, Journal of Acoustical Society of America 18 (04) (2010) 297–319.

[19] J. Hofmann, K. Heutschi, Simulation of outdoor sound propagation with a transmission line matrix method, Applied Acoustics 68 (2) (2007) 158–172.

[20] G. Guillaume, P. Aumond, B. Gauvreau, G. Dutilleux, Application of the transmission line matrix method for outdoor sound propagation modelling-part 1: Model presentation and evaluation, Applied Acoustics 76 (2014) 113–118.

[21] P. Aumond, G. Guillaume, B. Gauvreau, C. Lac, V. Masson, M. Bérengier, Application of the transmission line matrix method for outdoor sound propagation modellingpart 2: Experimental validation using meteorological data derived from the meso-scale model Meso-NH, Applied Acoustics 76 (2014) 107–112.

[22] J. B. Allen, D. A. Berkley, Image method for efficiently simulating small-room acoustics, Journal of Acoustical Society of America 65 (4) (1979) 943–950.

[23] J. Borish, Extension of the image model to arbitrary polyhedra, Journal of Acoustical Society of America 75 (6) (1984) 1827–1836.

[24] M. Vorländer, Simulation of the transient and steadystate sound propagation in rooms using a new combined raytracing/imagesource algorithm, Journal of Acoustical Society of America 86 (1) (1989) 172–178.

[25] M. Taylor, A. Chandak, Q. Mo, C. Lauterbach, C. Schissler, D. Manocha, Guided multiview ray tracing for fast auralization, IEEE Transactions on Visualization and Computer Graphics 18 (11) (2012) 1797–1810.

[26] M. Taylor, A. Chandak, L. Antani, D. Manocha, Resound: interactive sound rendering for dynamic virtual environments, in: Proceedings of 17th ACM international conference on Multimedia, 2009.

[27] T. Funkhouser, I. Carlbom, G. Elko, G. Pingali, M. Sondhi, J. West, A beam tracing approach to acoustic modeling for interactive virtual environments, in: Proceedings of 25th Conference on Computer Graphics and Interactive Techniques, 1998.

[28] A. Chandak, L. Antani, M. Taylor, D. Manocha, FastV: From-point visibility culling on complex models, Computer Graphics Forum 28 (4) (2009) 1237–1246.

[29] M. Bertram, E. Deines, J. Mohring, J. Jegorovs, H. Hagen, Phonon tracing for auralization and visualization of sound, IEEE Visualization.

[30] E. Deines, M. Bertram, J. Mohring, J. Jegorovs, F. Michel, H. Hagen, G. M. Nielson, Comparative visualization for wave-based and geometric acoustics, IEEE Transactions on Visualization and Computer Graphics 12 (5) (2006) 1173–1180.

[31] B. Kapralos, M. Jenkin, E. Milios, Sonel mapping: a probabilistic acoustical modeling method, Building Acoustics 15 (4) (2008) 289–313.

[32] C. Schissler, R. Mehra, D. Manocha, High-order diffraction and diffuse reflections for interactive sound propagation in large environments, ACM Transactions on Graphics 32 (4) (2014) 39.

[33] N. Tsingos, J. D. Gascuel, A general model for the simulation of room acoustics based on hierarchical radiosity, ACM SIGGRAPH sketch.

[34] S. Siltanen, T. Lokki, S. Kiminki, L. Savioja, The room acoustic rendering equation, Journal of Acoustical Society of America 122 (3) (2007) 1624–1635.

[35] R. M. Jones, J. P. Riley, T. M. Georges, HARPO: A versatile three-dimensional hamiltonian ray-tracing program for acoustic waves in an ocean with irregular bottom (1986).

[36] M. B. Porter, The BELLHOP manual and users guide: Preliminary draft (2011).

[37] S. M. Reilly, G. Potty, M. Goodrich, WaveQ3D: Fast and accurate acoustic transmission loss (TL) eigenrays, in littoral environments, for active sonar simulation/stimulation systems (2014).
URL http://oalib.hlsresearch.com/Rays/USML/usml_frontpage

[38] R. M. Jones, J. P. Riley, T. M. Georges, HARPA: A versatile three-dimensional hamiltonian ray-tracing program for acoustic waves in the atmosphere above irregular terrain (1986).

[39] U. O. of Naval Research, Ocean acoustics library: Ray models/software (2014).
URL http://oalib.hlsresearch.com/Rays/index.html

[40] P. Wang, L. Zhang, V. O. K. Li, A stratified acoustic model accounting for phase shifts for underwater acoustic networks, Sensors 13 (5) (2013) 6183–6203.

[41] B. D. Dushaw, J. A. Colosi, The eigenray acoustic ray tracing code (2013).
URL http://909ers.apl.washington.edu/twiki/bin/view/Main/EigenRay

[42] V. Červený, Seismic ray theory, Cambridge University Press, 2005.

[43] V. Havran, Heuristic ray shooting algorithms, Ph.D. thesis, Faculty of Electrical Engineering, Czech Technical University, Prague (2000).

[44] H. A. Panofsky, J. A. Dutton, Atmospheric turbulence. Models and methods for engineering applications, Wiley, 1984.

[45] R. B. Stull, An introduction to boundary layer meteorology, Vol. 13, Springer, 1988.

[46] R. Blumrich, D. Heimann, A linearized Eulerian sound propagation model for studies of complex meteorological effects, Journal of Acoustical Society of America 112 (2) (2002) 446–455.

[47] Z. C. Zheng, W. Li, Numerical stabilities and boundary conditions in time-domain eulerian simulations of acoustic wave propagations with and without background flow, Applied Mathematics and Computation 202 (1) (2008) 146–161.

[48] T. Oshima, M. Imano, Y. Hiraguri, Y. Kamoshida, Linearized euler simulations of sound propagation with wind effects over a reconstructed urban terrain using digital geographic information, Applied Acoustics 74 (12) (2013) 1354–1366.

[49] H. Yeh, R. Mehra, Z. Ren, L. Antani, D. Manocha, M. C. Lin, Wave-ray coupling for interactive sound propagation in large complex scenes, ACM Transactions on Graphics 32 (6) (2013) 165.

[50] T. V. Renterghem, E. M. Salomons, D. Botteldooren, Efficient FDTD-PE model for sound propagation in situations with complex obstacles and wind profiles, Acta Acustica united with acustica 91 (4) (2005) 671–679.

[51] T. V. Renterghem, E. M. Salomons, D. Botteldooren, Parameter study of sound propagation between city canyons with a coupled FDTD-PE model, Applied Acoustics 67 (6) (2006) 487–510.

[52] B. Bergen, B. V. Genechten, T. V. Renterghem, B. Pluymers, D. Vandepitte, D. Botteldooren, W. Desmet, Comparison of numerical prediction techniques for sound propagation in complex outdoor environments, in: Proceedings of 15th International Congress on Sound and Vibration, 2008.