

FastV: From-point Visibility Culling on Complex Models

Abstract

We present an efficient technique to compute the potentially visible set (PVS) of triangles in a complex 3D scene from a viewpoint. The algorithm computes a conservative PVS at object space accuracy. Our approach traces a high number of small, volumetric frusta and computes blockers for each frustum using simple intersection tests. In practice, the algorithm can compute the PVS of CAD and scanned models composed of millions of triangles at interactive rates on a multi-core PC. We also use the visibility algorithm to accurately compute the reflection paths from a point sound source. The resulting sound propagation algorithm is 10 – 20X faster than prior accurate geometric acoustic methods.

1 Introduction

Visibility computation is a widely-studied problem in computer graphics and related areas. Given a scene, the goal is to determine the set of primitives visible from a single point (i.e. from-point visibility), or from any point within a given region (i.e. from-region visibility). At a broad level, these algorithms can be classified into object space and image space algorithms. The object space algorithms operate at object-precision and use the raw primitives for visibility computations. The image space algorithms resolve visibility based on a discretized representation of the objects and the accuracy typically corresponds to the resolution of the final image. These algorithms are able to exploit the capabilities of rasterization hardware and can render large, complex scenes composed of tens of millions of triangles at interactive rates using current GPUs.

In this paper, we primarily focus on from-point, object space conservative visibility, whose goal is to compute a superset of visible geometric primitives. Such algorithms are useful for walk-throughs, shadow generation, global illumination and occlusion computations. Another application for object space visibility algorithms is accurate computation of reflection paths for acoustic simulation or sound rendering. Given a point sound source, 3D models of the environment with material data, and the receiver’s position, geometric acoustic (GA) methods perform multiple levels of reflections from the obstacles in the scene to compute the impulse response (IR). Sample-based propagation algorithms, such as stochastic ray-tracing for GA can result in statistical errors or inaccurate IRs [Funkhouser et al. 2003; Lenhert 1993]. As a result, we need to use object space visibility techniques, such as beam tracing [Funkhouser et al. 1998; Laine et al. 2009], to accurately compute the propagation paths. However, current object space visibility algorithms work well on simple scenes with tens of thousands of triangles or with large convex occluders. There is a general belief that it is hard to design fast and practical object space visibility algorithms for complex 3D models [Ghali 2001].

Main Results: We present a novel algorithm (FastV) for conservative, from-point visibility computation. Our approach is general and computes a potentially visible set (PVS) of scene triangles from a given view point. The main idea is to trace a high number of 4-sided volumetric frusta and compute efficiently simple connected blockers within each frustum. We use the blockers to compute a far plane and cull away the non-visible primitives, as described in Section 3.

Our guiding principle is to opt for simplicity in the choice of different parts of the algorithm, including frustum tracing, frustum-intersection tests, blocker and depth computations. The main contribution of the paper is primarily in combining known algorithms (or their extensions) for these parts. Overall, FastV is the first prac-

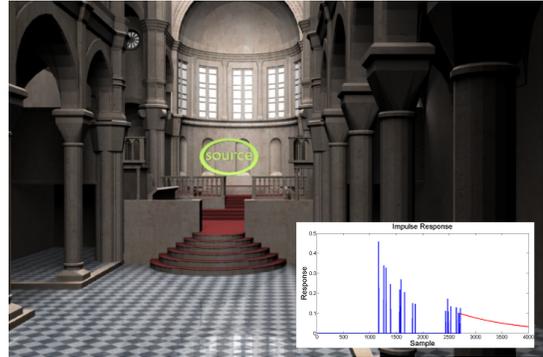


Figure 1: Fast Acoustic Simulation: We used FastV for accurate computation of reflection paths in this Cathedral model with 76.2K triangles. Our propagation algorithm performs three orders of reflections from the source (S) and compute the IR at the receiver (R) in less than 10 seconds. To the best of our knowledge, ours is the first efficient and accurate propagation algorithm to handle models of this complexity.

tical method for visibility culling in complex 3D models due to the following reasons:

1. Generality: Our approach is applicable to all triangulated models and does not assume any large objects or occluders. The algorithm proceeds automatically and is not susceptible to degeneracies or robustness issues.

2. Efficiency: We present fast algorithms based on Plücker coordinates to perform the intersection tests. We use hierarchies along with SIMD and multi-core capabilities to accelerate the computations. In practice, our algorithm can trace 101 – 200K frusta per second on a single 2.93 Ghz Xeon Core on complex models with millions of triangles.

3. Conservative: Our algorithm computes a conservative superset of the visible triangles at object-precision. As the frustum size is decreased, the algorithm computes a tighter PVS. We have applied the algorithm to complex CAD and scanned models with millions of triangles and simple dynamic scenes. In practice, we can compute conservative PVS, which is within a factor of 5 – 25% of the exact visible set, in a fraction of a second on a 16-core PC (as described in Section 5).

Accurate Sound Propagation: We use our PVS computation algorithm to accurately compute the reflection paths from a point sound source to a receiver, as described in Section 4. We use a two phase algorithm that first computes image-sources for scene primitives in the PVS computed for primary (or secondary) sources. This is followed by finding valid reflection paths to compute actual contributions at the receiver. We have applied our algorithm to complex models with tens of thousands of triangles. In practice, we observe a performance improvement of up to 20X over prior accurate propagation methods that use beam tracing.

2 Previous Work

The problem of visibility has been extensively studied in computer graphics, computational geometry, acoustic simulation and related areas for more than four decades. We refer the readers to excellent recent surveys [Durand 1999; Cohen-Or et al. 2003]. Due to space limitations, we only give a brief overview of some object space and sampling-based methods.

97 **Object space visibility computations:** There is extensive work on
 98 object-precision algorithms, including methods for hidden surface
 99 removal [Ghali 2001] and exact visibility computation from a point
 100 using beam tracing [Heckbert and Hanrahan 1984; Funkhouser
 101 et al. 1998; Overbeck et al. 2007] or Plücker coordinates [Niren-
 102 stein 2003]. Many exact algorithms have also been proposed for
 103 region-based visibility [Durand 1999; Duguet and Drettakis 2002;
 104 Nirenstein 2003; Bittner and Wonka 2005]. There is consider-
 105 able literature on conservative visibility computations from a single
 106 viewpoint [Bittner et al. 1998; Coorg and Teller 1997; Hudson et al.
 107 1997; Luebke and Georges 1995] or from a region [Koltun et al.
 108 2000; Leyvand et al. 2003; Teller 1992]. Some of these algorithms
 109 have been designed for special types of models, e.g. architectural
 110 models represented as cells and portals, 2.5D urban models, scenes
 111 with large convex occluders, etc. It is also possible to perform con-
 112 servative rasterization [Akenine-Möller and Aila 2005] on current
 113 GPUs to compute an object-precision PVS from a point.

114 **Image space or sample-based visibility computations:** These
 115 methods either make use of rasterization hardware or ray-shooting
 116 techniques to compute a set of visible primitives [Cohen-Or et al.
 117 2003]. Most of these methods tend to be either approximate or
 118 aggressive [Nirenstein and Blake 2004; Wonka et al. 2006]. Cur-
 119 rent GPUs provide support for performing occlusion queries for
 120 from-point visibility and are used for real-time display of complex
 121 3D models on commodity GPUs [Klosowski and Silva 2000; Mat-
 122 tausch et al. 2008].

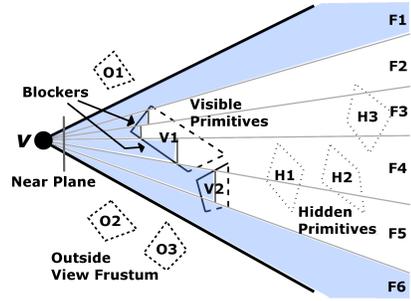
123 3 FastV: Visibility Computation

124 In this section, we present our conservative visibility computation
 125 algorithm. The inputs to our algorithm are: a view point ($\mathbf{v} \in \mathbb{R}^3$),
 126 a set of scene primitives (Π), and a viewing frustum (Φ), with an
 127 apex at \mathbf{v} . Our goal is to compute a subset of primitives $\pi \subseteq \Pi$
 128 such that every primitive $p \in \Pi$, which is hit by some ray $r \in \Phi$
 129 is included in the computed subset π . The subset π is called the
 130 potentially visible set (PVS). The smallest such PVS is the set of
 131 exactly visible primitives (π_{exact}). The subset π computed by our
 132 algorithm is conservative, i.e., $\pi \supseteq \pi_{exact}$. For the rest of the
 133 paper, we assume that the primitives are triangles, though our algo-
 134 rithm can be modified to handle other primitives. We also assume
 135 that the connectivity information between the scene triangles is pre-
 136 computed. We exploit this connectivity for efficient computation;
 137 however our approach is also applicable to polygon soup models. In
 138 order to perform fast intersection tests, we store the scene primitives
 139 in a bounding volume hierarchy (BVH) of axis-aligned bounding
 140 boxes (AABBs). This hierarchy is updated for dynamic scenes.

141 3.1 Overview

142 We trace pyramidal or volumetric beams from the viewpoint. Prior
 143 beam tracing algorithms perform expensive exact intersection and
 144 clipping computations of the beam against the triangles and tend
 145 to compute π_{exact} . Our goal is to avoid these expensive clipping
 146 computations, and rather perform simple intersection tests to com-
 147 pute the PVS. Moreover, it is hard to combine two or more non-
 148 overlapping occluders (i.e. occluder fusion) using object space
 149 techniques. This is shown in Figure 2, where object H_1 is occluded
 150 by the combination of V_1 and V_2 . As a result, prior conservative
 151 object space techniques are primarily limited to scenes with large
 152 occluders.

153 We overcome these limitations by tracing a high number of rela-
 154 tively small frusta and computing the PVS of each frustum inde-
 155 pendently. This makes it easy to parallelize our tracing algorithm on
 156 multi-core processors. We present very fast and simple algorithms
 157 to perform the intersection tests. In order to compute the PVS for
 158 each frustum, we try to compute a *blocker* that is composed of con-
 159 nected triangles (see Figure 3). The blockers are computed on the fly
 160 and need not correspond to a convex set or a solid object; rather



161 **Figure 2: Overview:** We divide the view-frustum with an apex
 162 at \mathbf{v} , into many small frusta. Each frustum is traced through the
 163 scene and its far plane is updated when it is blocked by a connected
 164 blocker. For example, frustum F_5 is blocked by primitives of object
 165 V_2 but frustum F_1 has no blockers. The objects V_1 and V_2 are part
 166 of the PVS and they block frusta F_2 to F_5 .

167 they are objects that are homomorphism to a disk. Given a blocker
 168 for a frustum, we update the far plane associated with that frustum.

169 **Frustum Tracing:** We use a simple four-sided frustum, which is
 170 represented as a convex combination of four corner rays intersect-
 171 ing at the apex. Each frustum has a near plane, four side planes, and
 172 a far plane. The near plane and the four side planes of a frustum are
 173 fixed and the far plane is parallel to the near plane. The depth of
 174 the far plane from the view point is updated as the algorithm com-
 175 putes a new blocker for a frustum. Our algorithm sub-divides Φ
 176 into smaller frusta using uniform or adaptive subdivision and com-
 177 putes a PVS for each frustum. Eventually, we take the union of
 178 these different PVSs to compute a PVS for Φ .

179 **Algorithm:** The algorithm computes the PVS for each frustum in-
 180 dependently. We initialize the far plane associated with the frustum
 181 to be at infinity and update its value if any connected blocker is
 182 found. The algorithm traverses the BVH to efficiently compute the
 183 triangles that potentially intersect with a given frustum. We perform
 184 fast Plücker intersection tests between the frustum and a triangle to
 185 determine if the frustum is completely inside, completely outside,
 186 or partially intersecting the triangle. If the frustum is partially in-
 187 tersecting, we reuse the Plücker test from the frustum-triangle in-
 188 tersection step to quickly find the edges that intersect the frustum
 189 (see Section 3.2). We perform frustum-triangle intersection with
 190 the neighboring triangles that are incident to these edges. We re-
 191 peat this step of finding edges that intersect with the frustum and
 192 perform intersection tests with the triangles incident to the edge till
 193 the frustum is completely blocked by some set of triangles. If a
 194 blocker is found (see Section 3.3), we update the far plane depth of
 195 the frustum. Any triangles beyond the far plane of the frustum are
 196 discarded from the PVS. If there is no blocker associated with the
 197 frustum, then all the triangles overlapping with the frustum belong
 198 to the PVS.

199 3.2 Frustum Blocker Computation

200 We define a blocker for a frustum as a set of connected triangles
 201 such that every ray inside the frustum hits some triangle in the frus-
 202 tum blocker (see Figure 3(a)). When we intersect a frustum with a
 203 triangle, the frustum could partially intersect the triangle. In such
 204 a case, we walk to the neighboring triangles based on that intersec-
 205 tion and try to find a blocker for the frustum (see Figure 3). We
 compute all the edges of the triangle that intersect with the frustum.
 For every intersecting edge, we walk to the neighboring triangle in-
 cident to the edge and perform the frustum-triangle intersection test
 with the neighbor triangle.

The intersection and walking steps are repeated until one of the
 following conditions is satisfied:

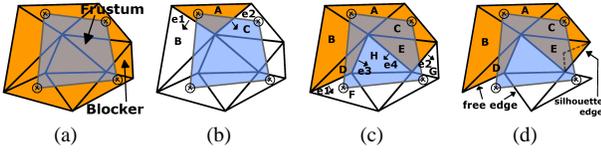


Figure 3: Frustum Blocker Computation: (a) Example of a blocker with connected triangles. (b)-(c) Intersection and Walking: Identify intersecting edges (e_1 , e_2 , e_3 , and e_4) and walk to the adjacent triangles (denoted by arrows from edge to the triangle). (d) Abort frustum blocker computation if a free-edge or a silhouette-edge is found.

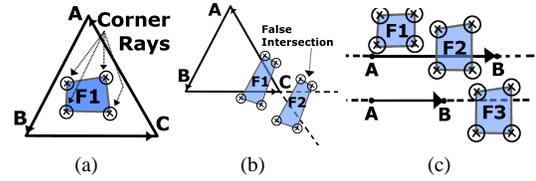


Figure 4: Conservative Plücker Tests: (a) All four corner rays of the frustum F_1 have the same orientation as seen along every directed edge of the triangle ABC . Thus, F_1 is completely-inside ABC . (b) Intersection between a frustum and a triangle can be conservative. F_2 will be classified as partially intersecting. (c) Different cases of frustum-edge intersections: F_1 does not intersect the edge AB , F_2 intersects AB . F_3 is falsely classified as intersecting AB by the test.

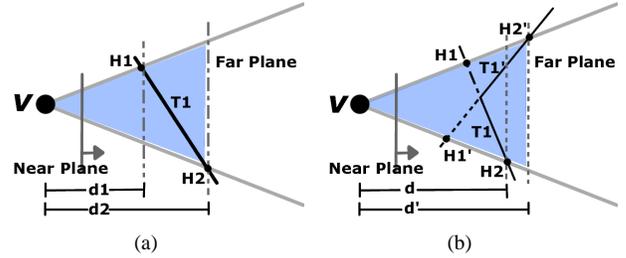


Figure 5: Updating Far Plane Depth: (a) Frustum lies completely inside triangle T_1 . The depth of the far plane is set to the maximum of d_1 and d_2 . (b) Triangles T_1 and T_1' constitute the blocker. We compute the far plane depths of each triangle and use the maximum value.

Frustum-Edge Intersection: When a frustum partially intersects with a triangle, we can quickly determine which edges of the triangle intersect the frustum. We reuse the Plücker test between the frustum and the triangle to find the edges of the triangle that intersect the frustum. As shown in Figure 4(c), a frustum intersects with an edge if all four corner rays of the frustum do not have the same orientation as seen along an edge. This test may falsely classify an edge as intersecting even if the frustum does not intersect the edge, as shown in Figure 4(c) and thereby make our algorithm conservative. This test is also used in Section 3.3 to compute a set of triangles that may block the frustum completely.

Far Plane Depth Update: The far plane associated with a frustum is updated whenever a blocker is found. The blocker may correspond to a single triangle or multiple triangles. If a frustum lies completely inside a triangle, the triangle blocks the frustum. We, therefore, mark the triangle as visible and update the depth of the far plane of the frustum as shown in Figure 5(a). The frustum intersects the triangle at points h_1 and h_2 , and d_1 and d_2 are the projected distances of $|Vh_1|$ and $|Vh_2|$ along the near plane normal. We set the far plane depth of the frustum as the maximum of the projected distances. In other cases, the blocker may be composed of multiple triangles. We update the far plane depth of the frustum as shown in Figure 5(b). We compute the far plane depth for every triangle in the frustum blocker, assuming the frustum is completely inside the triangle. In Figure 5(b), d and d' are the far plane depths for triangles T_1 and T_1' , respectively, of the frustum blocker. The far plane depth of the frustum is set to the maximum of far plane depths computed for the triangles in the frustum blocker, which is d' in this case.

3.4 Frustum Subdivision

Our algorithm implicitly assumes that the size of connected blockers is larger than the cross-section of the frusta. The simplest al-

- a All triangles incident to every intersecting edge found during the frustum blocker step have been processed. This implies that we have found a blocker.
- b A *free-edge*, i.e. an edge with only one incident triangle, or a *silhouette edge*, i.e. an edge with incident triangle facing in opposite directions as seen from the viewpoint, intersects with the frustum. In this case, we conclude that the current set of intersecting triangles does not constitute a blocker.

Note that our termination condition (b) for blocker computation is conservative. It is possible that there may exist a frustum blocker with a silhouette edge, but we need to perform additional computations to identify such blockers [Navazo et al. 2003; Laine 2006]. In this case, we opt for simplicity, and rather search for some other blocker defined by a possibly different set of triangles. Or we subdivide the frustum and the current object will become a blocker for a smaller sub-frustum.

If we terminate the traversal test due to condition (a), we have successfully found a frustum blocker. All triangles in the frustum blocker are marked visible and the far plane depth associated with the frustum is updated. Note that the depth of the far plane of the frustum is chosen such that all triangles in the frustum blocker lie in front of the far plane. If we terminate due to condition (b), then the algorithm can't guarantee the existence of a frustum blocker. All triangles processed during step are still marked visible but the far plane depth is not updated.

3.3 Frustum Intersection Tests

A key component of the algorithm is performing the intersection tests of the scene primitives with a frustum. The algorithm traverses the BVH and performs intersection tests between a frustum and the AABBs associated with the BVH. We use the technique proposed by Reshetov et al. [2005] to perform fast intersection tests between a frustum and an AABB. For every leaf node of the hierarchy we perform the intersection test with the frustum and triangle(s) associated with that leaf node. In order to perform the intersection test efficiently, we represent the corner rays of a frustum and the oriented edges of the triangle using Plücker coordinates [Shoemake 1998]. The orientation of a ray as seen along the edges of a triangle governs the intersection status of the ray with the triangle (see Figure 4(a)). Similarly, the orientation of four corner rays of the frustum as seen along the edges of a triangle governs the intersection status of the frustum with the triangle. We can determine with object-precision accuracy whether the frustum lies completely inside the triangle, completely outside the triangle, or partially intersects the triangle [Chandak et al. 2008].

In practice, the Plücker test is conservative and it can wrongly classify a frustum to be partially intersecting a triangle even if the frustum is completely outside the triangle (as shown in Figure 4(b)). This can affect the correctness of our algorithm as we may wrongly classify an object as a blocker due to these conservative intersection tests. We add a post-processing phase after each blocker computation to identify such cases.

289 gorithm subdivides a frustum in a uniform manner. This approach
 290 is simpler to implement and also simpler to parallelize on current
 291 multi-core and many-core architectures, in terms of load balanc-
 292 ing. However, many complex models (e.g. CAD models) have a
 293 non-uniform distribution of primitives in 3D. In that case, it may
 294 be more useful to perform adaptive subdivision of the frusta. In
 295 that case, we use the AD-FRUSTUM representation [Chandak et al.
 296 2008], which uses a quadtree data structure. We use the follow-
 297 ing criteria to perform subdivision. If the PVS associated with a
 298 frustum is large, we recursively compute the PVS associated with
 299 each sub-frustum. Whenever the algorithm only computes a partial
 300 blocker of connected triangles using the intersection tests, we
 301 estimate its cross-section area and use that area to compute the sub-
 302 frusta. There are other known techniques to estimate the distribu-
 303 tion of primitives [Wonka et al. 2006] and they can be used to
 304 guide the subdivision. As compared to uniform subdivision, adap-
 305 tive techniques reduce the total number of frusta traced for PVS
 306 computation. Moreover, we use spatial coherence to reduce the
 307 number of intersection tests between the parent and child frusta.

4 Geometric Sound Propagation

309 In this section, we describe our sound propagation algorithm. Given
 310 a point sound source, the CAD model of the scene with material
 311 properties (i.e. the acoustic space), and the receiver position, the
 312 goal is to compute the impulse response (IR) of the acoustic space.
 313 Later the IRs are convolved with the audio signal to reproduce the
 314 sound. We use our PVS computation algorithm described above for
 315 fast image-source computation that only takes into account specular
 316 reflections [Allen and Berkley 1979; Funkhouser et al. 2003; Laine
 317 et al. 2009]. In practice, this approach is only accurate for high
 318 frequency sources.

319 Each image source radiates in free space and considers secondary
 320 sources generated by mirroring the location of the input source over
 321 each boundary element in the environment. For each secondary
 322 source, the specular reflection path can be computed by performing
 323 repeated intersections of a line segment from the source position
 324 to the position of the receiver. In order to accurately compute all
 325 propagation paths, the algorithm creates image-sources (secondary
 326 sources) for every polygon in the scene. This step is repeated for
 327 all the secondary sources upto some user specified (say k) orders
 328 of reflection. Clearly, the number of image sources are $O(N^{k+1})$,
 329 where N is the number of triangles in the scene. This can become
 330 expensive for complex models.

331 We use our PVS computation algorithm to accelerate the compu-
 332 tation for complex scenes. We use a two stage algorithm. In the
 333 first stage, we use our conservative visibility culling algorithm and
 334 compute all the secondary image sources up to the specified orders
 335 of reflection. Since we overestimate the set of visibility triangles,
 336 we use the second stage to perform a validation step. For the first
 337 stage, we use a variant of Laine et al.'s [2009] algorithm and only
 338 compute the secondary image-sources for the triangles that are vis-
 339 ible from the source. Specifically, we shoot primary frusta from the
 340 sound source. For every primary frustum we compute its PVS. We
 341 then reflect the primary frustum against all visible triangles to cre-
 342 ate secondary frusta, which is similar to creating image-sources for
 343 visible triangles. This step is repeated for secondary frusta upto k
 344 orders of reflection. In second stage, we construct paths from the
 345 listener to the sound source for all the frusta which reach the lis-
 346 tener. As our approach is conservative, we have to ensure that this
 347 path is a valid path. To validate the path, we intersect each seg-
 348 ment of the path with the scene geometry and if an intersection is found
 349 the path is discarded.

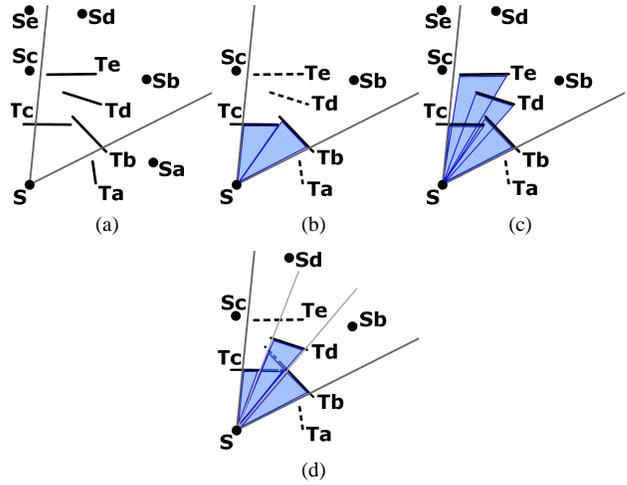


Figure 6: Geometric sound propagation: Comparison: Given a sound source, S , and triangles T_a, T_b, T_c and T_d the image source method (see 6a) creates image-sources of S against all primitives in the scene. Beam tracing algorithms [Funkhouser et al. 1998] (see 6b) compute image-sources for only exactly visible triangles, T_b and T_c in this case. Accelerated beam tracing [Laine et al. 2009] approach computes image-sources for all triangles inside the beam volume (see 6c), i.e., T_b, T_c, T_d , and T_e in this case. Our algorithm (see 6d) computes image-sources for triangles T_b, T_c , and T_d in the PVS by our FastV algorithms.

5 Results

351 In this section, we present our results on from-point conservative
 352 visibility (Section 5.1) and accurate geometric sound propagation
 353 (Section 5.2). Our results were generated on a 16-core 64-bit Intel
 354 X7350@2.93 GHz. We used SSE instructions to accelerate frustum
 355 intersection tests and use OpenMP to parallelize on multiple cores.

5.1 Visibility Results

357 We demonstrate our results on computing from-point object space
 358 conservative PVS on a variety of models ranging from simple mod-
 359 els (like soda hall, armadillo, blade) to complex models (like power
 360 plant and thai statue) to a dynamic model (flamenco animation).
 361 These models are shown in Figure 7. Our results are summarized in
 362 Table 1. We are not aware of any prior method that can compute the
 363 exact visible set on these complex models. Therefore, we compute
 364 an approximation to π_{exact} . by shooting frusta at $4K \times 4K$ resolu-
 365 tion and compute the PVS for that resolution. The *PVS-ratio* refers
 366 to: (size of PVS) / (size of π_{exact}), and is a measure of how con-
 367 servative is the answer. In all benchmarks, we are able to compute
 368 a conservative approximation to the PVS at interactive rates on the
 369 multi-core PC. The frame sequences used for generating average
 370 results are shown in accompanying video. Further, we show that
 371 our approach converges well to π_{exact} as we shoot higher number
 372 of frusta (see Figure 8). Detailed results on convergence for each
 373 model are provided in the Appendix.



Figure 7: Benchmarks: Left to right: (a) Armadillo (345K triangles). (b) Blade (1.8M triangles). (c) Thai Statue (10M triangles). (d) Soda Hall (1.5M triangles). (e) PowerPlant (12M triangles). (f) Flamenco (dynamic scene)

Model			PVS Ratio	PVS Size	Time (ms)
Name	Tris	Type			
Armadillo	345K	scan	1.16	98K	59
Blade	1.8M	scan	1.05	190K	179
Thai	10M	scan	1.06	210K	132
SodaHall	1.5M	cad	1.22	2.1K	30
PowerPlant	12M	cad	1.25	15K	259
Flamenco	40K	dynamic	1.11	7K	31

Table 1: Main Results: Our results on from-point conservative visibility for models of varying complexities. All the timing were computed on a 16-core 64-bit Intel X7350@2.93 GHz. The algorithm first performs view frustum culling and uses FastV only for occlusion culling. The PVS ratio provides a measure of how conservative is the computed answer with respect to occlusion culling.

5.2 Geometric Sound Propagation Results

We present our results on accurate geometric sound propagation in section. Table 2 summarizes our results. We perform geometric sound propagation on models of varying complexity from 438 triangles to 212K triangles. We used three benchmarks presented in accelerated beam tracing (ABT) algorithm [Laine et al. 2009]. We also used two additional complex benchmarks with 78K and 212K triangles. We are not aware of any implementation of accurate geometric propagation that can handle models of such complexity.

Model	Tris	Time (msec)	Speed Up (ABT)
Simple Room	438	10	10.1
Regular Room	1190	58	22.2
Complex Room	5635	406	11.8
Sibenik	78.2K	4500	–
Trade Show	212K	13600	–

Table 2: Accurate sound propagation: We highlight the performance of sound propagation algorithms on four benchmarks. We observe 10 – 20 speedup on the simple model.

6 Comparison and Analysis

In this section we analyze our algorithm and compare it with prior techniques. The accuracy of our algorithm is governed by the accuracy of the intersection tests, which exploit the IEEE floating-point hardware. Our approach is robust and general, and not prone to any degeneracies.

Conservative approach: We compute a conservative PVS for every frustum. This follows from our basic approach to compute the blockers and far planes for each frustum. In practice, our approach can be overly conservative in some cases. The underlying blocker

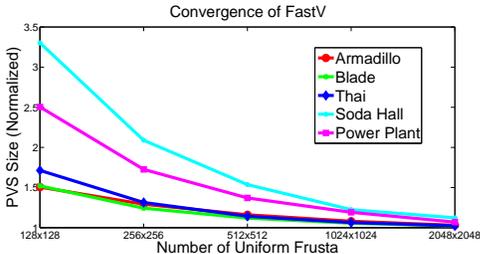


Figure 8: PVS ratio vs. #Frusta: As the number of frusta increase, the PVS computed by our answer converges to π_{exact} . This graph shows the rate of convergence for different benchmarks. The CAD models have a higher ratio as compared to scanned models.

computation algorithm is conservative. Moreover, we don't consider the case when the union of two or more objects can serve as a blocker. This is shown in Fig. 2) with two disjoint occluders, V_1 and V_2 . Instead of using more sophisticated algorithms for blocker computation, we found it cheaper to subdivide the frustum into sub-frusta and compute blockers for them. As a result, we can make our approach less conservative by using more frusta and the PVS (π) converges well to π_{exact} (see Figure 8).

Model connectivity and triangle soup models: Our algorithm exploits the connectivity information in the model to compute the blockers, which are formed using connected triangles. If the connectivity information is not available, then the algorithm would subdivide the frustum such that each blocker would consist of only one triangle.

6.1 Visibility Computations

Our approach performs volumetric tracing, which is similar to beam tracing. However, we don't perform exact clipping operations to compute an exact representation of the visible surface. Rather we only estimate the triangles belonging to the PVS by identifying the blockers for each frustum. None of the triangles in the scene are subdivided. Beam tracing algorithms can also be accelerated by using spatial data structures [Funkhouser et al. 1998; Overbeck et al. 2007; Laine et al. 2009], but they have mostly been applied to scenes with large occluders (e.g. architectural models). In practice, beam tracing can be considerably more expensive than our approach. On the other hand, the PVS computed by our algorithm tends to be more conservative than that computed by beam tracing.

Most of the prior object space conservative visibility culling algorithms are designed for scenes with large occluders [Bittner et al. 1998; Coorg and Teller 1997; Hudson et al. 1997; Luebke and Georges 1995]. These algorithms can work well on special types of models, e.g. architectural models represented using cells and portals or urban scenes. In contrast, our approach is mainly designed for general 3D models and doesn't make any assumption about large occluders.

It is possible to perform conservative rasterization using current GPUs [Akenine-Möller and Aila 2005]. However, it has the overhead of rendering additional triangles and CPU-GPU communication latency. It may be possible to accelerate conservative rasterization by using hierarchical methods [Mattausch et al. 2008]. The resulting approach could be faster than FastV in some cases, but may compute a more conservative PVS. This could result in a slower sound propagation algorithm.

It is hard to make a direct comparison with image space approaches because of their accuracy. In practice, image space approaches can exploit the rasterization hardware or fast ray-tracing techniques [Reshetov et al. 2005] and would be faster than FastV. Moreover, image space approaches also perform occluder fusion and in some cases may compute a smaller set of visible primitives than FastV. However, the main issue with the image space approaches is deriving any tight bounds on the accuracy of the result. This is highlighted in the appendix, where we used ray tracing to approximate the visible primitives. In complex models like the powerplant, we need a sampling resolution of at least $32K \times 32K$ to compute a good approximation of the visible primitives. At lower resolutions, the visible set computed by the algorithm doesn't seem to converge well.

6.2 Sound Propagation Algorithm

Most accurate geometric acoustic methods can be described as variants of the image-source method. Figure 6 compares different accurate geometric sound propagation methods. The main difference between these methods is in terms of which image-sources they choose to compute [Funkhouser et al. 1998; Laine et al. 2009;

Schröder and Lentz 2006; Antonacci et al. 2008]. A naïve image source method computes image sources against all triangles in the scene for a sound source (Figure 6(a)) [Allen and Berkley 1979]. Beam tracing methods compute the image-sources for exactly visible triangles from a source (Figure 6(b)) and this method is applied recursively. Recent methods based on beam tracing, like accelerated beam tracing [Laine et al. 2009], compute image-sources for every triangle inside the beam volume (Figure 6(c)). Our approach, shown in Figure 6(d), finds the conservative PVS from a source and compute image-sources for the triangles in the conservative PVS. As a result, for a given model our approach considers more image-sources as compared to exact beam tracing. It is an efficient compromise between the expensive step to compute exactly visible triangles in beam tracing vs. computing additional image-sources in accelerated beam tracing. We observe 10–20X speedups over prior accurate methods. Recently, Chandak et al. [2008] also used adaptive frustum tracing for geometric sound propagation. However, that algorithm performs discrete clipping and intersection tests at the boundary of the frustum and therefore, it is hard to derive any good bounds on the accuracy of impulse responses.

6.3 Limitations

Our approach has some limitations. Since we don't perform occluder fusion, the PVS computed by our algorithm can be overly conservative sometimes. If the scene has no big occluders, we may need to trace a large number of frusta. Our intersection tests are fast, but the conservative nature of the blocker computation can result in a larger PVS. The model and its hierarchy are stored in main memory, and therefore our approach is limited to in-core models. Our algorithm is easy to parallelize and works quite well, but is still slower than image space approaches that perform coherent ray tracing or use GPU rasterization capabilities.

7 Conclusions and Future Work

We present a fast and simple visibility culling algorithm and demonstrate its performance on complex models. The algorithm is general and works well on complex 3D models. To the best of our knowledge, this is the first from-point object space visibility algorithm that can handle complex 3D models with millions of triangles at almost interactive rates.

There are many avenues for future work. We will like to implement the algorithm on a many-core GPU or upcoming Larrabee processor to further exploit the high parallel performance of these commodity processors. This could provide capability to design more accurate rendering algorithms based on object-precision visibility computations on complex models (e.g. shadow generation). We can use temporal coherence between successive frames along with adaptive subdivision to further improve the runtime performance. We will also like to evaluate the trade-offs in terms of using more sophisticated blocker computations algorithms [Navazo et al. 2003; Laine 2006]. In terms of sound propagation, our approach can be extended to compute edge diffraction based on uniform theory of diffraction (UTD).

References

AKENINE-MÖLLER, T., AND AILA, T. 2005. Conservative and tiled rasterization using a modified triangle set-up. *Journal of graphics tools* 10, 3, 1–8.

ALLEN, J. B., AND BERKLEY, D. A. 1979. Image method for efficiently simulating small-room acoustics. *The Journal of the Acoustical Society of America* 65, 4 (April), 943–950.

ANTONACCI, F., FOCO, M., SARTI, A., AND TUBARO, S. 2008. Fast tracing of acoustic beams and paths through visibility lookup. *Audio, Speech, and Language Processing, IEEE Transactions on* 16, 4 (May), 812–824.

BITTNER, J., AND WONKA, P. 2005. Fast exact from-region visibility in urban scenes. *Eurographics Symposium on Rendering*, 223–230.

BITTNER, J., HAVRAN, V., AND SLAVIK, P. 1998. Hierarchical visibility culling with occlusion trees. *Computer Graphics International, 1998. Proceedings* (Jun), 207–219.

CHANDAK, A., LAUTERBACH, C., TAYLOR, M., REN, Z., AND MANOCHA, D. 2008. Ad-frustum: Adaptive frustum tracing for interactive sound propagation. In *Proc. IEEE Visualization*.

COHEN-OR, D., CHRYSANTHOU, Y., SILVA, C., AND DURAND, F. 2003. A survey of visibility for walkthrough applications. *Visualization and Computer Graphics, IEEE Transactions on* 9, 3 (July-Sept.), 412–431.

COORG, S., AND TELLER, S. 1997. Real-time occlusion culling for models with large occluders. In *SIGGRAPH '97: Proceedings of the 1997 symposium on Interactive 3D graphics*, ACM, New York, NY, USA, 83–ff.

DUGUET, F., AND DRETTAKIS, G. 2002. Robust epsilon visibility. *Proc. of ACM SIGGRAPH*, 567–575.

DURAND, F. 1999. *3D Visibility, Analysis and Applications*. PhD thesis, U. Joseph Fourier.

FUNKHOUSER, T., CARLBOM, I., ELKO, G., PINGALI, G., SONDHI, M., AND WEST, J. 1998. A beam tracing approach to acoustic modeling for interactive virtual environments. In *Proc. of ACM SIGGRAPH*, 21–32.

FUNKHOUSER, T., TSINGOS, N., AND JOT, J.-M. 2003. Survey of Methods for Modeling Sound Propagation in Interactive Virtual Environment Systems. *Presence and Teleoperation*.

GHALI, S. 2001. A survey of practical object space visibility algorithms. In *SIGGRAPH Tutorial Notes*.

HECKBERT, P. S., AND HANRAHAN, P. 1984. Beam tracing polygonal objects. In *Proc. of ACM SIGGRAPH*, 119–127.

HUDSON, T., MANOCHA, D., COHEN, J., LIN, M., HOFF, K., AND ZHANG, H. 1997. Accelerated occlusion culling using shadow frusta. In *Proc. of ACM Symposium on Computational Geometry*, 1–10.

KLOSOWSKI, J., AND SILVA, C. 2000. The prioritized-layered projection algorithm for visible set estimation. *IEEE Trans. on Visualization and Computer Graphics* 6, 2, 108–123.

KOLTUN, V., CHRYSANTHOU, Y., AND COHEN-OR, D. 2000. Virtual occluders: An efficient intermediate pvs representation. In *Eurographics Workshop on Rendering*, 59–70.

LAINE, S., SILTANEN, S., LOKKI, T., AND SAVIOJA, L. 2009. Accelerated beam tracing algorithm. *Applied Acoustic* 70, 1, 172–181.

LAINE, S. 2006. *An Incremental Shaft Subdivision Algorithm for Computing Shadows and Visibility*. Master's thesis, Helsinki University of Technology.

LENHART, H. 1993. Systematic errors of the ray-tracing algorithm. *Applied Acoustics* 38.

LEYVAND, T., SORKINE, O., , AND COHEN-OR, D. 2003. Ray space factorization for from-region visibility. *Proc. of ACM SIGGRAPH*, 595–604.

LUEBKE, D., AND GEORGES, C. 1995. Portals and mirrors: Simple, fast evaluation of potentially visible sets. In *ACM Interactive 3D Graphics Conference*.

MATTAUSCH, O., BITTNER, J., AND WIMMER, M. 2008. Chc++: Coherent hierarchical culling revisited. *Proc. of Eurographics*, 221–230.

NAVAZO, I., ROSSIGNAC, J., JOU, J., AND SHARIF, R. 2003. Shieldtester: Cell-to-cell visibility test for surface occluders. In *Proc. of Eurographics*, 291–302.

NIRENSTEIN, S., AND BLAKE, E. 2004. Hardware accelerated visibility preprocessing using adaptive sampling. In *Eurographics Workshop on Rendering*.

NIRENSTEIN, S. 2003. *Fast and Accurate Visibility Preprocessing*. PhD thesis, University of Cape Town, South Africa.

OVERBECK, R., RAMAMOORTHY, R., AND MARK, W. R. 2007. A Real-time Beam Tracer with Application to Exact Soft Shadows. In *Eurographics Symposium on Rendering*, 85–98.

RESHETOV, A., SOUPIKOV, A., AND HURLEY, J. 2005. Multi-level ray tracing algorithm. *ACM Trans. Graph.* 24, 3, 1176–1185.

SCHRÖDER, D., AND LENTZ, T. 2006. Real-Time Processing of Image Sources Using Binary Space Partitioning. *Journal of the Audio Engineering Society (JAES)* 54, 7/8 (July), 604–619.

SHOEMAKE, K. 1998. Plücker coordinate tutorial. *Ray Tracing News* 11, 1.

TELLER, S. J. 1992. *Visibility Computations in Densely Occluded Polyhedral Environments*. PhD thesis, CS Division, UC Berkeley.

WONKA, P., WIMMER, M., ZHOU, K., MAIERHOFER, S., HESINA, G., AND RESHETOV, A. 2006. Guided visibility sampling. *Proc. of ACM SIGGRAPH*, 494–502.