Planar Visibility Counting*

M. Fischer M. Hilbig C. Jähn F. Meyer auf der Heide M. Ziegler,
Heinz Nixdorf Institute and University of Paderborn, GERMANY

Abstract

For a fixed virtual scene (=collection of simplices) $\mathcal{S}$ and given observer position $\vec{p}$, how many elements of $\mathcal{S}$ are weakly visible (i.e. not fully occluded by others) from $\vec{p}$? The present work explores the trade-off between query time and preprocessing space for these quantities in 2D: exactly, in the approximate deterministic, and in the probabilistic sense. We deduce the *existence* of an $\mathcal{O}(m^2/n^2)$ space data structure for $\mathcal{S}$ that, given $\vec{p}$ and time $\mathcal{O}(\log n)$, allows to approximate the ratio of occluded segments up to arbitrary constant absolute error; here m denotes the size of the Visibility Graph—which may be quadratic, but typically is just linear in the size n of the scene $\mathcal{S}$. On the other hand, we present a data structure *constructible* in $\mathcal{O}\big(n\cdot\log(n)+m^2\cdot\text{polylog}(n)/\ell\big)$ preprocessing time and space with similar approximation properties and query time $\mathcal{O}(\ell\cdot\text{polylog}\, n)$, where $1\leq\ell\leq n$ is an arbitrary parameter.

1 Motivation, Introduction, Overview

In computer graphics, occlusion culling algorithms exhibit both a) the potential of drastically accelerating the rendering of virtual scenes (namely in case when most objects are occluded and therefore need not be drawn) as well as b) an overhead for selecting those objects which are at least partially visible and thus have to be sent to the graphics hardware (conservative approach). Particularly for scenes and observer viewpoints where 'most' is visible, the benefit of a) may be rather low and due to b) result in a net performance *loss*. One would like to estimate in advance whether occlusion culling actually is going to pay off, and if not, turn it off. More generally our goal is to devise algorithms that automatically and optimally adapt, on a per-scene and per-frame basis, to the trade-off between spending more/less computational effort in filtering occluded objects on the one hand and the gain/loss in reduced rendering time on the other hand.

We propose the number of drawing primitives at least partly visible from a given observer position as a quantitative measure for estimating, in comparison with the total number of primitives (i.e. the *visibility ratio*), how much conservative occlusion culling can reduce the rendering complexity.

Definition 1 (visibility count) *For a scene* $\mathcal{S}=\{S_1,\dots,S_n\}$ *of 'geometric primitives'* $S_i\subseteq\mathbb{R}^d$, *a subset of* 'targets' $\mathcal{T}\subseteq\mathcal{S}$, *and an observer position* $\vec{p}\in\mathbb{R}^d$, *let* $\mathcal{V}(\mathcal{S},\vec{p},\mathcal{T}) := \{T\in\mathcal{T}\big|\exists\vec{q}\in T : \forall S\in\mathcal{S}\setminus\{T\} : [\vec{p},\vec{q}]^\circ\cap S=\emptyset\}$ *and denote by* $\mathrm{V}(\mathcal{S},\vec{p},\mathcal{T}) := \mathrm{Card}\,\mathcal{V}(\mathcal{S},\vec{p},\mathcal{T})$ *the number of objects in* $\mathcal{T}$ *visible (i.e. not fully occluded) from* $\vec{p}$ *through* $\mathcal{S}$. *Here* $[\vec{p},\vec{q}]^\circ := \{\lambda\vec{p}+(1-\lambda)\vec{q} : 0<\lambda<1\}$ *means the relatively open straight line segment from* $\vec{p}$ *to* $\vec{q}$.

The present work describes a new randomized algorithm for approximating the visibility count $\mathrm{V}(\mathcal{S},\vec{p},\mathcal{S})$ for 2D scenes of non-intersecting line segments. We presume $\mathcal{S}$ to be arbitrary but fixed, whereas the observer position $\vec{p}$ is given as input. In view of the size $N=\mathrm{Card}(\mathcal{S})$ of the virtual scenes and the frame rates aimed at, running times must be sublinear $\mathcal{O}(N^\alpha)$, $\alpha<1$; after preprocessing $\mathcal{S}$ into data structures of almost linear space $\mathcal{O}(N^{1+\epsilon})$, $\epsilon\ll 1$. Here (time and) space complexity refers to the *number* of (operations on) unit-size real coordinates used (performed) by an algorithm—as opposed to e.g. rationals of varying bitlength. Our algorithm features a parameter $1\leq\ell<n$ to trade preprocessing space for query time.

2 Exact Visibility Counting

Visibility comprises a highly active field of research, both heuristically and in the sound framework of computational geometry [CCSD03]. Particularly the latter has proven combinatorially and algorithmically non-trivial already in the plane [ORou87, Ghos07]. Here the case of (simple) polygons is well studied [CAF07]; and so is point–point, point–segment, and segment–segment visibility for scenes $\mathcal{S}$ of n non-crossing line segments, captured e.g. in the Visibility Graph data structure. Our observer positions $\vec{p}$, however, are *not* restricted to segments in $\mathcal{S}$.

Fact 2 *Given a collection* $\mathcal{S}$ *of* n *non-crossing line segments in the plane and an observer position* $\vec{x}$, $\mathrm{V}(\mathcal{S},\vec{x},\mathcal{S})$ *can be calculated in time* $\mathcal{O}(n\cdot\log n)$ *and space* $\mathcal{O}(n)$.

*Supported by DFG projects Me872/12-1 within SPP 1307, by Zi1009/1-2, and by DFG Research Training Group GK-693 of the *Paderborn Institute for Scientific Computation* (PaSCo). The last author is grateful to 정지원 (né Otfried Schwarzkopf) for the opportunity to visit KAIST. A full version of this extended abstract including proofs and figures is available at arxiv.org/abs/0810.0052.

Also visibility *reporting* algorithms have to be ruled out for cases of linear output size. Our goal is *counting* in *sub*linear time. Even logarithmic time does become easily feasible based on the locus approach of storing *all* visibility counts in a Visibility Space Partition (VSP) [Schi01].

Fact 3 *a) For a collection $\mathcal{S}$ of n non-crossing line segments in the plane, there exists a partition of $\mathbb{R}^2$ into $\mathcal{O}(n^4)$ convex cells such that, for all observer positions $\vec{x} \in C$ within one cell C, $\mathcal{V}(\mathcal{S}, \cdot, \mathcal{S})$ is the same.*

b) The data structure indicated in a) and including for each cell its corresponding visibility count $\mathrm{V}(\mathcal{S}, C, \mathcal{S})$ uses storage $\mathcal{O}(n^4 \cdot \log n)$ and can be computed in time $\mathcal{O}(n^4 \cdot \log^2 n)$. Then, given an observer position $\vec{x}$, its corresponding cell $C \ni \vec{x}$ can be located, and the associated visibility count identified, in time $\mathcal{O}(\log n)$.

Proof. of Claim a): Draw lines through all $\binom{2n}{2}$ pairs of the $2n$ segment endpoints. It is easy to see that, in order for a near segment to appear in sight, the observer has to cross one of these $\mathcal{O}(n^2)$ lines; compare Lemma 7 below. Hence, within each of the $\mathcal{O}(n^4)$ cells they induce, the subset of segments weakly visible remains the same. □

2.1 Size of Visibility Space Partitions

The quartic size bound of Fact 3a) is of course prohibitive—and sharp in the worst case. In order to avoid trivialities, we restrict to *non*degenerate segment configurations $\mathcal{S}$ where i) any two segments meet only in their common endpoints, ii) no three endpoints share a common line, and iii) any two lines, defined by pairs of endpoints, do meet.

We have already referred to (and implicitly employed in Fact 3 a refinement of) the Visibility Space Partition; so here finally comes the formal

Definition 4 *For two non-degenerate collections $\mathcal{S}$ and $\mathcal{T}$ of segments in the plane, partition all viewpoints $\vec{p} \in \mathbb{R}^2$ into classes having equal visibility $\mathcal{V}(\mathcal{S}, \vec{p}, \mathcal{T})$. Moreover let $\mathrm{VSP}(\mathcal{S}, \mathcal{T})$ denote the collection of connected components of these equivalence classes. The size of VSP is the number of line segments forming the boundaries of these components.*

Observe that $\mathrm{VSP}(\mathcal{S}, \mathcal{T})$ indeed constitutes a planar subdivision: a coarsening of the $\mathcal{O}(n^4)$ convex polygons induced by the arrangement of $\mathcal{O}(n^2)$ lines from the proof of Fact 3a). In fact a class of viewpoints of equal visibility can be disconnected and delimited by very many segments, hence merely counting the number of classes or cells does not reflect the combinatorial complexity. Fact 3a) and Proposition 5a) correspond to [Mato02, EXERCISE 6.1.7].

Proposition 5 *a) Even for a singleton target T, there exist a nondegenerate line segment configurations $\mathcal{S}$ such that $\mathrm{VSP}(\mathcal{S}, \{T\})$ has $\Omega(n^4)$ separate connected components.*

b) To each n, there exists a nondegenerate configuration $\mathcal{S}$ of at least n segments admitting a convex planar subdivision of complexity $\mathcal{O}(n)$ such that, from within each cell, the view to $\mathcal{S}$ is constant; i.e. $\mathrm{VSP}(\mathcal{S}, \mathcal{S})$ has linear size.

c) The size of $\mathrm{VSP}(\mathcal{S}, \mathcal{S})$ is at most quadratic in the size m of the Visibility Graph of $\mathcal{S}$.

d) A data structure as in Fact 3 can be calculated in time $\mathcal{O}(n \cdot \log n + m^2 \cdot \log^2 n)$ and space $\mathcal{O}(m^2)$.

Since the Visibility Graph itself can have at most quadratically more edges than vertices, Item c) strengthens Fact 3a). Empirically we have found that a 'random' scene typically induces a VSP of roughly quadratic size. This agrees with a 'typical' scene to have a linear size Visibility Graph according to [ELPZ07].

2.2 Single Target Visibility: Time versus Space

Intuitively it should be possible to reduce the memory consumption of Fact 3 at the expense of increasing the time bound. We achieve this for the case of one target, that is the decision version of visibility $\vec{x} \mapsto \mathrm{V}(\mathcal{S}, \vec{x}, \{T\}) \in \{0, 1\}$:

Theorem 6 *For each 2D scene $\mathcal{S}$, and line segment T, and $1 \le \ell \le n$, an $\mathcal{O}(n^4/\ell)$ size data structure can be computed in time $\mathcal{O}(n^4 \cdot \log^2 n/\ell)$ that allows to decide, given $\vec{p}$, whether T is weakly visible from $\vec{p}$ through $\mathcal{S}$, in query time $\mathcal{O}(\ell \cdot \log n)$.*

Note that Fact 3 is included as $\ell = 1$; whereas for $\ell := n^{1-\epsilon}$ we get arbitrarily close to cubic space while maintaining sublinear query time. On the other end, no value of ℓ recovers Fact 2.

Lemma 7 *Fix a collection $\mathcal{S} \uplus \{T\}$ of $n + 1$ non-crossing segments in the plane. Let $L_1, \ldots, L_k$ denote the $k = \binom{2n+2}{2}$ lines induced by the pairs of endpoints of segments in $\mathcal{S} \cup \{T\}$. For an observer moving in the plane, the weak visibility of T can change only as she crosses*

- *either one of the lines L_i intersecting T*
- *or someline supporting a segment $S \in \mathcal{S}$.*

Proof. Standard continuity argument: Let $\vec{p}$ denote the observer's position and suppose point $\vec{x} \in T$ is visible, i.e. the segment $[\vec{p}, \vec{x}]$ does not intersect $S \in \mathcal{S}$. Now move $\vec{p}$ until $\vec{x}$ is just about to become hidden behind $S \in \mathcal{S}$. Then start moving $\vec{x}$ on T such as to remain visible. Keep moving $\vec{p}$ and adjusting $\vec{x}$: this is possible (at least) as long as the line through $\vec{p}$ and $\vec{x}$ avoids all endpoints of $\mathcal{S} \cup \{T\}$. □

Proof. [Theorem 6] Consider, as in the proof of Fact 3, the $\mathcal{O}(n^2)$ lines induced by pairs of segment endpoints of $\mathcal{S}$. Consider the intersections of these lines with T (if any). Partition T into $\mathcal{O}(\ell)$ subsegments $T_1, \ldots, T_\ell$, each intersecting $\mathcal{O}(n^2/\ell)$ of the above lines. For each piece T_i, take the arrangement $\mathcal{A}_i$ of size $\mathcal{O}\big((n^2/\ell + n)^2\big)$ induced by those lines intersecting T_i, and all $\mathcal{O}(n)$ lines through one endpoint of T_i and one of some $S \in \mathcal{S}$, and all $\mathcal{O}(n)$ lines supporting segments from $\mathcal{S}$. By Lemma 7, within each cell C of $\mathcal{A}_i$, the weak visibility of T_i is constant (either `yes` or `no`) and can be stored with C: Doing so for each $\mathcal{A}_i$ ($1 \leq i \leq \ell \leq n$) and each of the $\mathcal{O}(n^4/\ell^2 + n^3/\ell + n^2)$ cells C of $\mathcal{A}_i$ uses memory of order $\mathcal{O}(n^4/\ell + n^3 + n^2\ell) = \mathcal{O}(n^4/\ell)$ as claimed; and corresponding time according to Fact 3b).

Then, given a query point $\vec{p} \in \mathbb{R}^2$, locating $\vec{p}$ in each arrangement $\mathcal{A}_i$ takes total time $\mathcal{O}(\ell \cdot \log n)$; and yields the answer to whether T_i is weakly visible from $\vec{p}$ or not. Now T itself is of course visible iff some T_i is: a disjunction computable in another $\mathcal{O}(\ell)$ steps. □

Similar to Proposition 5d), we can improve Theorem 6 to $\mathcal{O}(n \cdot \log n + m^2 \cdot \log^2 n/\ell)$ preprocessing time and $\mathcal{O}(m^2/\ell)$ space and query time $\mathcal{O}(\ell \cdot \log n)$, where m denotes the size (number of edges) of the Visibility Graph of $\mathcal{S}$.

3 Approximate Visibility Counting

Lacking deterministic exact algorithms for calculating visibility counts satisfying both time and space requirements, we now resort to approximations: of $\mathrm{V}(\mathcal{S}, \vec{x}, \mathcal{S})$ up to prescribable absolute error $k \in \mathbb{N}$ or, equivalently, of the visibility *ratio* $\mathrm{V}(\mathcal{S}, \vec{x}, \mathcal{S}) / \operatorname{Card}(\mathcal{S})$ up to absolute error $\epsilon = k/\operatorname{Card}(\mathcal{S})$. Our main result presents a randomized approximation within sublinear query time using almost cubic preprocessing space and time in the worst-case; almost linear space and quadratic preprocessing time in the 'typical' case.

3.1 Deterministic Approach: Relaxed VSPs

Visibility space partitions, and the algorithms based upon them, are so memory expensive because they distinguish (i.e. introduces separate arrangement cells for) observer positions whose visibility differs by as little as one; recall Definition 4. Considerably more space efficient algorithms are feasible by partitioning observer space into (or merely covering it by) more coarse classes:

Definition 8 *Fix $k \in \mathbb{N}$ and collections $\mathcal{S}$ and $\mathcal{T}$ of non-intersecting segments in the plane. Some covering $\{C_1, \ldots, C_I\}$ of $\mathbb{R}^2$ is a k-relaxed VSP of $(\mathcal{S}, \mathcal{T})$ if*

$$\forall 1 \leq i \leq I \; \forall \vec{p}, \vec{q} \in C_i : \quad \mathrm{V}(\mathcal{S}, \vec{p}, \mathcal{T}) - \mathrm{V}(\mathcal{S}, \vec{q}, \mathcal{T}) \;\leq\; k.$$

In the sequel we shall restrict to k-relaxed VSPs which constitute planar subdivisions (i.e. each C_i being a simple polygon); and refer to their size in the sense of Definition 4. Indeed, such VSPs allow for locating a given observer position $\vec{x}$ in logarithmic time to yield a cell $C_i \ni \vec{x}$ which, during preprocessing, had been assigned a value $\mathrm{V}(\mathcal{S}, \vec{q}, \mathcal{T})$ approximating $\mathrm{V}(\mathcal{S}, \vec{x}, \mathcal{T})$ up to absolute error at most k.

Observe that, for $\operatorname{Card} \mathcal{T} \leq k$, the trivial planar subdivision $\{\mathbb{R}^2\}$ is a k-relaxed VSP of $(\mathcal{S}, \mathcal{T})$. In particular the quartic lower size bound of Proposition 5b) applies only to 0-relaxed VSPs but breaks down for $k \geq 1$. This indicates that much smaller sizes become feasible when considering k-relaxed VSPs for, say, $k \approx \sqrt{n}$ or even $k \approx n/\log n$.

Proposition 9 *a) To $k < n - 1$ there exists a non-degenerate family $\mathcal{S}$ of n segments in the plane such that any k-relaxed VSP has size $\Omega(n^2/k)$.*

b) There also exist such families such that any k-relaxed VSP has size at least $\Omega(n^4/k^4)$.

c) Let $\mathcal{S}$ be a non-degenerate family of n segments in the plane and N the size of its VSP. Then there exists a k-relaxed VSP of size $\lfloor N/(k+1) \rfloor$.

d) There also exists a k-relaxed VSP of size $\mathcal{O}(m^2/k^2)$, where m denotes the size of the Visibility Graph of $\mathcal{S}$.

Recall that $N \leq m^2 \leq n^4$, thus leaving quadratic gap between a) and d) for k small; and between b) and d) for k large. Item c) succeeds over d) in cases where N asymptotically does not exceed m^2/k.

Proposition 9d) is an application of the Cutting Lemma of Chazelle, Friedman, and Matoušek [Mato02, LEMMA 4.5.3]. Notice that we only claim the *existence* of such small data structures. In order to *construct* them, our proofs of Proposition 9c+d) both proceed by first calculating the 0-relaxed VSP and then coarsening it. Specifically, although a Triangular Cutting can be found in asymptotically nearly optimal time [Agar90], 'filling' it with visibility counts seems to incur quartic preprocessing time, again—not to mention the impractically large constants hidden in asymptotic big-Oh notation.

3.2 Random Sampling

We now proceed to the simple, generic, randomized

Algorithm 10

i) Guess a sample target $\mathcal{T} \subseteq \mathcal{S}$ of size m.

ii) Calculate the count $\mathrm{V}(\mathcal{S}, \vec{x}, \mathcal{T})$ of objects in $\mathcal{T}$ visible through $\mathcal{S}$.

iii) Return the ratio $\mathrm{V}(\mathcal{S}, \vec{x}, \mathcal{T}) / \operatorname{Card}(\mathcal{T})$;

iv) and hope that it does not deviate too much from the 'true' ratio $\mathrm{V}(\mathcal{S}, \vec{x}, \mathcal{S}) / \operatorname{Card}(\mathcal{S})$.

Item iv) is justified for instance by the following application of the Chernoff–Hoeffding Bound [AlSp00]:

Lemma 11 *Fix $\vec{x} \in \mathbb{R}^d$ and $\delta > 0$, then choose $\mathcal{T} \subseteq \mathcal{S}$ as m independent identically distributed random draws from $\mathcal{S}$. It holds*

$$\boldsymbol{Prob}_{\mathcal{T}}\Big[\big|\tfrac{\mathrm{V}(\mathcal{S},\vec{x},\mathcal{T})}{m} - \tfrac{\mathrm{V}(\mathcal{S},\vec{x},\mathcal{S})}{n}\big| \geq \delta\Big] \leq 2 \cdot e^{-2m\cdot\delta^2}$$

In other words: In Algorithm 10 taking m (quadratic in the aimed *absolute* accuracy δ but) *constant* with respect to the scene size n suffices to achieve the desired approximation with constant probability; slightly increasing it further amplifies exponentially the chance for success.

3.3 The VC-Dimension of Visibility

Note that the random experiment $\mathcal{T}$ and the probability analysis of its properties in Lemma 11 holds for each $\vec{x}$ but not uniformly in $\vec{x}$. This means for our purpose to re-sample $\mathcal{T} \subseteq \mathcal{S}$ at every frame. On the other hand, Lemma 11 does not exploit any geometry. An important connection between combinatorial sampling and geometric properties is captured by the Vapnik–Chervonenkis Dimension [AlSp00]:

Fact 12 *Let X be a set and $\mathcal{R}$ a collection of subsets $R \subseteq X$. Denote* $d := \mathrm{VCdim}(X, \mathcal{R}) :=$

$$\max\big\{\operatorname{Card} Y \big| Y \subseteq X, \{Y \cap R : R \in \mathcal{R}\} = 2^Y\big\}. \quad (1)$$

Let $Y \subseteq X$ be randomly distributed at uniform of $\operatorname{Card}(Y) \geq \Omega\big((d \cdot \log\frac{d}{\delta} + \log\frac{1}{p})/\delta^2\big)$. Then with probability at least $1-p$ it holds for each $R \in \mathcal{R}$: $\big|\operatorname{Card}(X \cap R)/\operatorname{Card}(X) - \operatorname{Card}(Y \cap R)/\operatorname{Card}(Y)\big| \leq \delta$.

Lemma 13 *Fix a collection $\mathcal{S}$ of n noncrossing $(d-1)$–dimensional simplices in $\mathbb{R}^d$.*

a) Define $X := \mathcal{S}$ and $\mathcal{R} := \{\mathcal{V}(\mathcal{S}, \vec{x}, \mathcal{S}) : \vec{x} \in \mathbb{R}^d\}$. Then $\mathrm{VCdim}(X, \mathcal{R}) \leq d^2 \cdot \big(\log n + \mathcal{O}(1)\big)$.

b) In $\mathbb{R}^2$ there exist non-degenerate collections $\mathcal{S} = X$ of n line segments such that $\mathrm{VCdim}(X, \mathcal{R}) \geq \log n - \mathcal{O}(\operatorname{loglog} n)$.

Proof. [Claim a] Fact 3a) straight forwardly generalizes to yield $\operatorname{Card}\mathcal{R} \leq \mathcal{O}(n)^{d^2}$. Hence, in Equation (1), $2^Y = \{Y \cap R : R \in \mathcal{R}\}$ requires $2^{\operatorname{Card} Y} \leq \operatorname{Card}\mathcal{R}$ and therefore $\operatorname{Card} Y \leq d^2 \cdot \big(\log n + \mathcal{O}(1)\big)$. □

3.4 Main Result

Lemma 13a), together with Fact 12, enhances Lemma 11: For $d := 2$, a sample $\mathcal{T}$ of size $m := \mathcal{O}(\operatorname{polylog} n)$ has visibility ratio $\mathrm{V}(\mathcal{S}, \vec{x}, \mathcal{T})/m$ close to the 'true' on $\mathrm{V}(\mathcal{S}, \vec{x}, \mathcal{S})/n$ with high probability with respect to *all* viewpoints $\vec{x}$! In particular we may preprocess the visibility of each $T \in \mathcal{T}$ separately according to Theorem 6 and conclude in combination with Proposition 5d):

Theorem 14 *Given $0 < \delta < 1$, a collection $\mathcal{S}$ of n non-crossing segments in the plane ($d = 2$), and $1 \leq \ell \leq n \leq m$ where m denotes the size of the Visibility Graph of $\mathcal{S}$. Then a randomized algorithm can preprocess $\mathcal{S}$ within time $\mathcal{O}\big(n \cdot \log n + m^2 \cdot \operatorname{polylog} n \cdot \log\frac{1}{\delta}/(\ell\cdot\delta^2)\big)$ and space $\mathcal{O}\big(m^2 \cdot \operatorname{polylog} n \cdot \log\frac{1}{\delta}/(\ell\cdot\delta^2)\big)$ into a data structure having with high probability the following property: Given $\vec{x} \in \mathbb{R}^2$, one can approximate the visibility ratio $\mathrm{V}(\mathcal{S}, \vec{x}, \mathcal{S})/\operatorname{Card}(\mathcal{S})$ up to absolute error at most δ in time $\mathcal{O}\big(\ell \cdot \operatorname{polylog} n \cdot \log\frac{1}{\delta}/\delta^2\big)$.*

Recall the trade-off between space and time gauged by parameter ℓ; and m being 'typically' linear in n.

4 Conclusion and Perspective: Dimension > 2

For constant δ, $k := \delta \cdot n$ in Proposition 9d) yields logarithmic query time in worst-case quadratic ('typical' case even linear) preprocessing space: this beats Theorem 14. On the other hand, the simplicity of Algorithm 10 makes it very practical; and preprocessing according to Theorem 6 is reasonable as well. In fact we have implemented the algorithm underlying Theorem 14 and are currently evaluating it for the purpose of adaptive culling as mentioned in the introduction.

The present work restricts to the planar case of line segments. Many virtual scenes are $2\frac{1}{2}$-dimensional: objects of various heights rooted on a common plane. In 3D scenes of triangles, however, visibility immediately becomes 3SUM-complete and thus unlikely tractable in time $o(n^2)$, see [GaOv95, Section 6.1].

References

[Agar90] P. Agarwal: "Partitioning Arrangements of Lines I: An Efficient Deterministic Algorithm", pp.449–483 in *Discrete Comput. Geom.* vol.**5** (1990).

[AlSp00] N. Alon, J.H. Spencer: "*The Probabilistic Method*", 2nd Edition, Wiley (2000).

[CAF07] S. Charneau, L. Aveneau, L. Fuchs: "Exact, Robust and Efficient Full Visibility Computation in Plücker Space", pp.773–782 in *Visual Comput.* vol.**23** (2007).

[CCSD03] D. Cohen-Or, Y.L. Chrysanthou, C.T. Silva, F. Durand: "A Survey of Visibility for Walkthrough Applications", pp.412–431 in *IEEE Transactions on Visualization and Computer Graphics* vol.**9:3** (2003).

[ELPZ07] H. Everett, S. Lazard, S. Petitjean, L. Zhang: "On the Expected Size of the 2D Visibility Complex", pp.361–381 in *Int. J. Comput. Geom.* vol.**17:4** (2007).

[GaOv95] A. Gajentaan, M. Overmars: "On a Class of $\mathcal{O}(n^2)$ Problems in Computational Geometry", pp.165–185 in *Computat. Geom.: Theory & Applications* vol.**5:3** (1995).

[Ghos07] S.K. Ghosh: "*Visibility Algorithms in the Plane*", Cambridge University Press (2007).

[Mato02] J. Matoušek: "*Lectures on Discrete Geometry*", Springer Graduate Texts in Mathematics vol.**212** (2002).

[ORou87] J. O'Rourke: "*Art Gallery Theorems and Algorithms*", Oxford University Press (1987).

[Schi01] R.D. Schiffenbauer: "A Survey of Aspect Graphs", *TR-CIS-2001-01* Brooklyn University (2001).