An Experimental Assessment of the 2D Visibility Complex

Hazel Everett* Sylvain Lazard* Sylvain Petitjean* Linqiao Zhang†

Abstract

We make an experimental assessment of the size of the 2D visibility complex of disjoint unit discs randomly distributed in the plane with density μ. We observe that the number of free bitangents is asymptotically linear in the number of discs and we study the dependence of the linear asymptote in terms of the density of the scene. Specifically, for a particular range of scene densities μ, we exhibit an approximation of the number of free bitangents in terms of μ and the number n of discs, for n larger than some function of μ. We also notice how our approximation gained for rather large densities can be used to guess the onset of the linear behavior for small densities.

1 Introduction

Many graphics applications, such as global illumination and interactive walkthroughs, require computing the visibility among objects in the scene. However, given that the number n of objects is naturally big in the real world, computing in a naive way the pairwise visibility, which is $O(n^2)$ in 2D and $O(n^4)$ in 3D, tends to be prohibitive.

Much research has been carried out to improve visibility computation. One approach uses the visibility complex data structure, which decomposes the 2D scene into a cellular structure such that each cell contains a set of maximal free segments which have the same visibility. The Greedy Flip Algorithm [2, 4] was developed to optimally compute the 2D visibility complex in $O(n \log n + m)$ time, where m is the size of the visibility graph of the scene, and $O(n)$ storage. This algorithm has been implemented by P. Angelier [1].

The visibility complex certainly facilitates visibility computations, but a further question should be investigated: what is the size of the output? For convex objects, the critical visibility events occur when crossing a bitangent to a pair of objects. The number of free bitangents thus dominates the size of the visibility graph. In the worst case, this number is $\Theta(n^2)$ in 2D. What about practical situations? In this paper, we carry out detailed experiments with the visibility complex package [1]. In the context of randomly distributed discs in the plane, our results suggest that the size of the visibility complex is asymptotically linear in the number of discs. This conclusion is supported by a theoretical analysis.

The rest of this paper is organized as follows. Section 2 presents the model used for generating disjoint unit discs in the plane and Section 3 outlines the experiments that were made. Section 4 gives some conclusions pertaining to the results found. In particular, we observe an asymptotic linearity for the number of bitangents, memory usage and time consumption. The linear asymptotic behavior is supported by a theoretical result summarized in Section 5.

†McGill University & LORIA, lzhang15@cs.mcgill.ca.

*LORIA - Université Nancy 2, INRIA Lorraine, and CNRS Firstname.Name@loria.fr.

2 Model

We first present the 2D scene model on which we run our experiments. We define the universe as a disc of radius R containing n pairwise distinct discs of unit radius. The density of the scene, denoted μ, is $\frac{n}{R^2}$. For scenes with constant density μ, we study the asymptotic behavior of the 2D visibility complex. Note that the constraint that the unit discs are disjoint is required in order to use the Greedy Flip Algorithm of the visibility complex package [1].

Defining a 2D scene includes defining variables n and μ, deriving $R = \sqrt{\frac{n}{\mu}}$, and generating n disjoint unit discs whose centers lie inside the disc of radius $R - 1$. The random distribution is achieved by generating n random points for the centers of the unit discs, with a constraint that the distance between each pair of points is larger than 2, so that the unit discs are disjoint. To satisfy this constraint, we check for each newly generated center if there is sufficient space in the universe for the corresponding disc. Note that this constraint will introduce some bias into the scene generating scheme, especially when the scene is dense. Thus the distribution of the discs will not necessarily be uniform.

In the sequel, we call a *bitangent* a maximal line segment tangent to two discs and intersecting the interior of no other.

3 Experiments

With the scene model defined in Section 2, we measure, for various densities μ between 0.0025 and 0.55, the number of bitangents of the scene with n unit discs.

We also measure the memory usage and the running time costs of computing this set of bitangents. We do not consider densities μ bigger than 0.55 because our scene generation scheme fails for large values of μ. As Figure 1 shows, density 0.55 already implies a fairly dense scene. (The densest possible scene has $\mu \approx 0.91$).

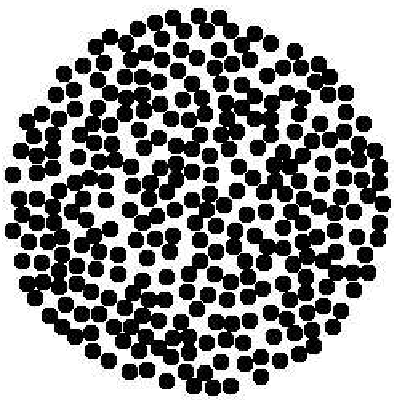

Figure 1: A scene of 300 discs with density $\mu = 0.55$.

For each density μ varying from 0.025 to 0.55 with an increment of 0.025, we compute, for each μ, the visibility complex using the CGAL-based package [1] on scenes with a number of unit discs varying from 500 to $3,000$ with an increment of 100. And when μ varies from 0.0025 to 0.025 with an increment of 0.0025, we compute the visibility complex with n varying from 40 to $1,000$ with an increment of 40. For these small densities, we did not go above $1,000$ discs because of memory limitations in the software implementation.

For each test, we run 10 experiments and we only report the mean of the measure because the standard deviation is negligible. We report the number of oriented bitangents, the memory usage in units of kBs and the running time in units of 10^{-4} seconds (so that running time, number of bitangents and memory usage can be drawn on the same plot).

Note that the visibility complex package outputs directed bitangents, while it seems more intuitive to count non-directed bitangents. We therefore make a slight distinction between the two in what follows. Also, to ease our presentation, we denote the number of directed bitangents as $Dbts$, the number of non-directed bitangents as $Nbts$, memory usage as m, and running time as t. These symbols will be used in the rest of this paper.

All the experiments were made on a i686 machine with AMD Athlon 1.73 GHz CPU and 1 GB of main memory.

μ	0.025	0.150	0.275	0.400	0.525
$Dbts$	669.3	119.3	63.4	41.5	29.3
m	269.9	49.0	26.5	17.9	13.3
t	154.7	31.7	19.4	14.4	11.6

Table 1: Estimated slope for the asymptotic linear behavior of the number of bitangents, memory usage and time consumption, for selected values of μ in the range $[0.025, 0.55]$.

μ	0.025	0.150	0.275	0.400	0.525
$Dbts$	-185,464	-11,282	-3,482	-1,203	-243

Table 2: Estimated y-intercept for the asymptotic number of bitangents, for selected values of μ in the range $[0.025, 0.55]$.

4 Results and Interpretation

We now present our experimental results. We show in Figure 2 the outputs of our experiments for four representative values of the density μ: $0.0025, 0.005, 0.025$, and 0.55.

4.1 Asymptotic Linearity

Figure 2 shows quite clearly that $Dbts$, m and t, as functions of the number of discs of the scene, have a linear asymptotic behavior. We note that the slope is different for each μ and is a decreasing function of μ. We also observe that the linear behavior appears to start for smaller values of μ as μ increases.

We now focus on studying the relation of the slope of the linear asymptotic behavior versus the scene density μ. For each selected value of $\mu \in [0.025, 0.55]$, we use least-squares fitting to find the best line approximating the data for n between $1,000$ and $3,000$. The results for the slope of the asymptotic linear behavior are given in Table 1. We observe that the extracted slopes appear intimately related to the inverse of μ. We thus try to fit a function of the form

$$y = \frac{a}{\mu} + b$$

to the above. We indeed obtain a very good least-squares fitting. Since the constant b is very small, we drop it in what follows. For the three curves corresponding to $Dbts$, m, and t, we obtain the following estimated slopes:

$$S_{Dbts} = \frac{17.07}{\mu}, \quad S_m = \frac{6.92}{\mu}, \quad S_t = \frac{4.11}{\mu}.$$

In Figure 3, we plot the functions S_{Dbts}, S_m and S_t, along with the actual data points, and observe that the functions fit quite closely to the data.

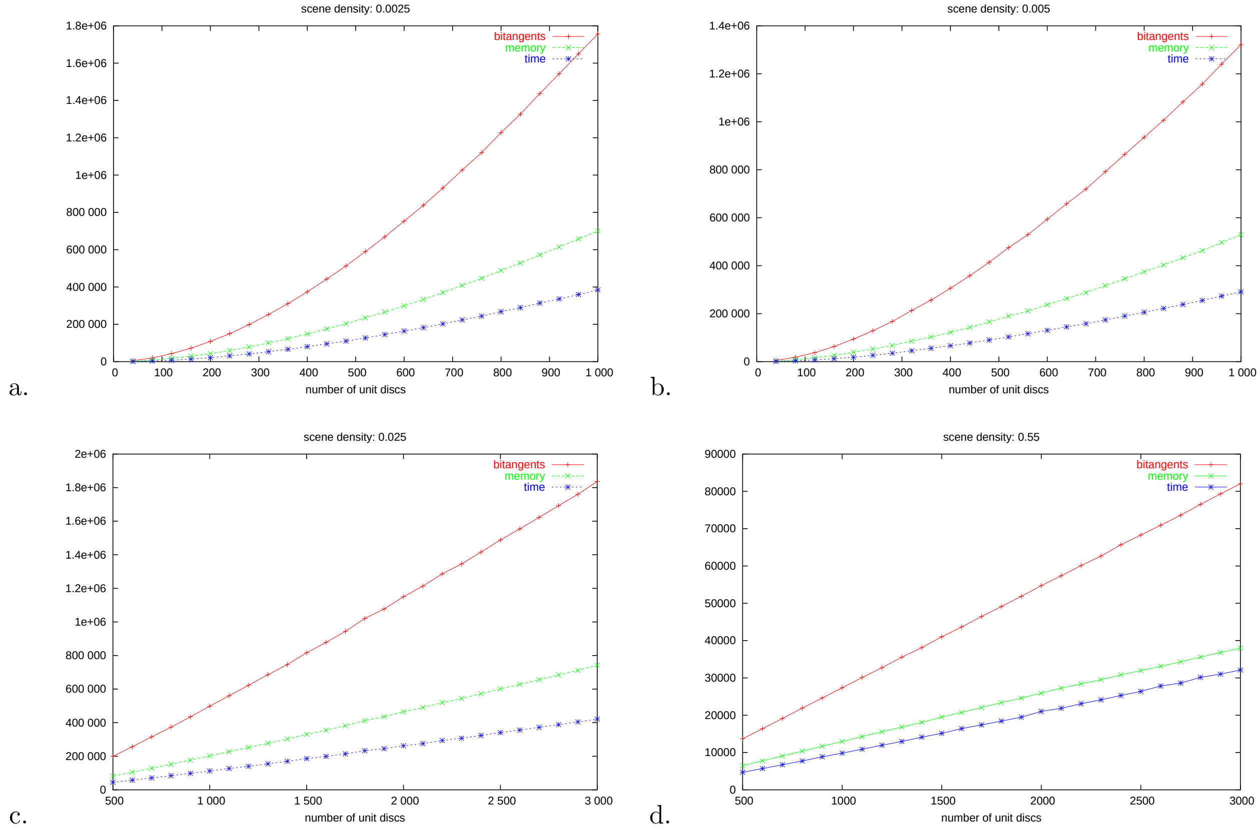

Figure 2: Plots of the number of directed bitangents, memory usage, and running time versus the number of unit discs, when scene density is a. $\mu = 0.0025$, b. $\mu = 0.005$, c. $\mu = 0.025$, d. $\mu = 0.55$. The unit of the memory usage is kBs, that of the running time is 10^{-4} seconds.

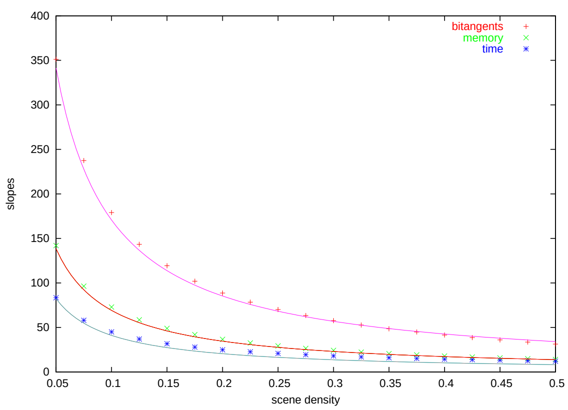

Figure 3: The slopes of the asymptotic linear behaviors of the number of bitangents, memory usage, and running time in terms of the scene density μ, with both the experimental data points and the extracted fits of the form $\frac{a}{\mu}$.

For large values of n, the observed number of non-oriented bitangents $Nbts$, memory usage m and running time t can be written as

$$\begin{aligned} Nbts &= \frac{1}{2}(S_{Dbts}n + C_{Dbts}), \\ m &= S_m n + C_m, \\ t &= S_t n + C_t, \end{aligned}$$

where C_{Dbts}, C_m and C_t are functions of μ. The least-squares fitting done above also produced samples of these functions for specific values of μ (see Table 2 for the samples of C_{Dbts}). Now we can try to fit simple functions to the given data. For the function C_{Dbts}, Figure 4 shows that fitting with a function $\frac{a}{\mu}$ does not give good results, but adding an inverse quadratic term improves things dramatically. We obtain the following fit:

$$C_{Dbts} = -\frac{63.6}{\mu^2} - \frac{2,232}{\mu} + 5,110 \quad (1)$$

4.2 Starting Point of Linear Behavior

An interesting issue is to determine, as a function of μ, the value of n at which the linear asymptotic behavior starts.

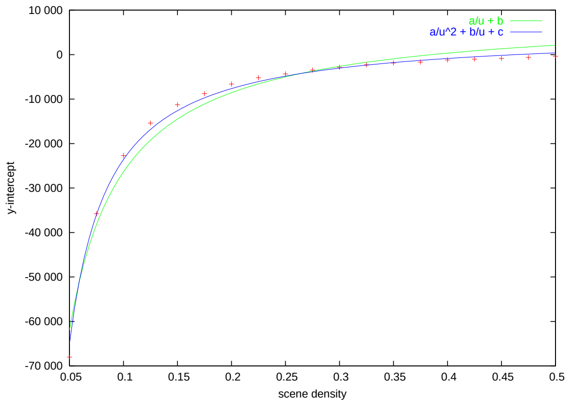

Figure 4: The y-intercept of the asymptotic linear behavior of the number of bitangents versus scene density μ, with the experimental data points, a "linear" fit of the form $\frac{a}{\mu} + b$ and a "quadratic" fit of the form $\frac{a}{\mu^2} + \frac{b}{\mu} + c$.

We focus on the number of bitangents. For specific values of $\mu \in [0.015, 0.200]$, we estimate the value n_0 of n such that, for all $n > n_0$, the difference between the computed number of bitangents and that estimated by the linear asymptote is less than some threshold, here $\varepsilon = 2\%$. The result is displayed in Figure 5. Again, the data seem to exhibit a familiar look, and we can try to fit a function of the form $n_0 = \frac{a}{\mu} + b$. The result is quite satisfactory (Figure 5). The fitted function is given by:

$$n_0 = \frac{22.15}{\mu} + 61 \qquad (2)$$

This function can be used to guess the onset of the linear behavior for small densities. For instance, when $\mu = 0.0025$ (Figure 2.a), Equation (2) predicts that the linear behavior is only reached for $n > n_0 \approx 8,921$. For such a low density, we however stopped our experiments at $n = 3,000$ because the number of bitangents is large and memory consumption is prohibitive due to memory leaks in the visibility complex package.

5 Conclusion

We made an experimental assessment of the size of the 2D visibility complex of a scene consisting of n randomly distributed disjoint unit discs. We showed that the number of free non-oriented bitangents is well approximated by

$$\frac{8.54}{\mu}n - \frac{31.8}{\mu^2} - \frac{1,116}{\mu} + 2,555 \quad \text{for} \quad n > \frac{11.08}{\mu} + 30.5$$

where μ denotes the density of the scene.

Theoretical results obtained in a slightly different context support the experimental conclusions of this paper. In [3], the authors show that, amongst n uniformly

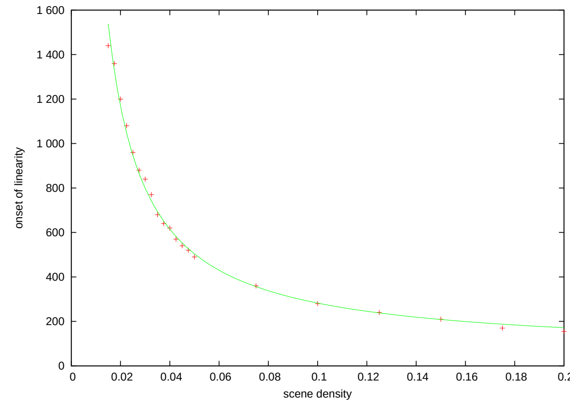

Figure 5: Plot of the value n_0 of n at which the linearity "starts" for specific values of μ and fitting of the available data with a function of the form $n_0 = \frac{a}{\mu} + b$.

distributed (non-necessarily disjoint) unit balls in $\mathbb{R}^3$, the expected number of maximal non-occluded line segments tangent to four balls is linear. This result certainly carries over to the case of unit discs in the plane, so that the expected number of non-occluded bitangents to n uniformly distributed unit discs is linear. The major difference with our experimental setting is that we assume the discs to be disjoint and do not guarantee that the distribution is uniform. Still, the asymptotic linearity appears to carry over to our setting.

The merit of our experiments is to give a good hold on the asymptotic behavior of the number of bitangents, while only a very rough upper bound on the slope of the asymptote could be given in [3]. And the fact that the estimated slope is reasonably small in our random setting gives indications that the size of the visibility complex might be tractable in practical, real-world applications.

References

[1] P. Angelier and M. Pocchiola. CGAL-based implementation of visibility complexes. Technical Report ECG-TR-241207-01, Effective Computational Geometry for Curves and Surfaces (ECG), 2003. `http://mapage.noos.fr/angelier/`.

[2] P. Angelier and M. Pocchiola. A sum of squares theorem for visibility complexes and applications. In B. Aronov, S. Basu, J. Pach, and M. Sharir, editors, *Discrete and Computational Geometry*, volume 25 of *Algorithms and Combinatorics*, pages 79–139. Springer-Verlag, 2003.

[3] O. Devillers, V. Dujmović, H. Everett, X. Goaoc, H.-S. Na, and S. Petitjean. The expected number of 3D visibility events is linear. *SIAM Journal on Computing*, 32(6):1586–1620, 2003.

[4] M. Pocchiola and G. Vegter. Topologically sweeping visibility complexes via pseudo-triangulations. *Discrete and Computational Geometry*, 16(4):419–453, 1996.