On the light side of geometric graphs

Eyal Ackerman* Rom Pinchasi†

Abstract

Let G be a geometric graph on n vertices in general position in the plane. Suppose that for every line ℓ in the plane the subgraph of G induced by the set of vertices in one of the two half-planes bounded by ℓ has at most k edges ($k \geq 1$ may be a function of n). Then G has at most $O(n\sqrt{k})$ edges. This bound is best possible.

Keywords: Geometric graphs, k-near bipartite.

1 Introduction

Let G be an n-vertex *geometric graph*. That is, a graph drawn in the plane such that its vertices are distinct points and its edges are straight-line segments connecting corresponding vertices. It is usually assumed, as we will assume in this paper, that the set of vertices of G is in general position in the sense that no three of them lie on a line.

Let ℓ be a line that does not contain any vertex of G (unless stated otherwise, we consider only such lines). Every edge of G either crosses ℓ, or is contained in one of the two half-planes bounded by ℓ.[1] We say that G has a *k-light side* with respect to ℓ, if one of these half-planes contains at most k edges of G. If G has a *k-light side* with respect to *every* line ℓ, then G is *k-near bipartite*. We consider the following problem: What is the maximum number of edges of an n-vertex k-near bipartite geometric graph?

We will think of k as a function of n, that is $k = k(n)$, so obviously this question is interesting only when $k(n) = o(n^2)$. The following simple construction shows an $n\sqrt{k}$ lower bound. Let G be the geometric graph whose vertices are the vertices of a regular n-gon P. We denote the vertices of P (and of G) by $v_0, \ldots, v_{n-1}$, indexed in a clockwise order. The *cyclic distance* between two vertices, v_i, v_j, $i < j$, is defined as $\min\{j-i, i+n-j\}$. The edge set of G consists of all edges (v_i, v_j) such that the cyclic distance between v_i and v_j is at least $\lfloor n/2 - \sqrt{k} \rfloor$. The number of edges in G is at least $n\sqrt{k}$ as each vertex has degree at least $2\sqrt{k}$. One can easily verify that each half-plane bounded by a line that passes through the center of P contains at most k edges of G. It follows that if ℓ is a line not passing through the center of P, then the half-plane that is bounded by ℓ and does not contain the center of P must contain at most k edges. Therefore, G is k-near bipartite.

Our main result shows that this construction is essentially best possible.

Theorem 1. *Let n and k be positive integers. Every n-vertex k-near bipartite geometric graph has at most $O(n\sqrt{k})$ edges.*

*Department of Mathematics, Physics, and Computer Science, University of Haifa at Oranim, Tivon 36006, Israel. ackerman@sci.haifa.ac.il.

†Mathematics Department, Technion—Israel Institute of Technology, Haifa 32000, Israel. room@math.technion.ac.il. Supported by BSF grant (grant No. 2008290).

[1]Since we only consider lines that do not contain vertices, it makes no difference if the half-planes are open or closed.

Remark. The condition that k should be positive is merely technical. The case $k = 0$ is equivalent to the case of geometric graphs in which there is no pair of disjoint edges. This is because once there are two disjoint edges in a geometric graph, this graph cannot be 0-near bipartite, as witnessed by any line separating the two disjoint edges. It is a well known classical result that such graphs in which there are no two disjoint edges contain at most n edges and that this bound can indeed be attained ([2, 5]).

Related work. It is a common technique when studying Turán-type problems in geometric graphs to split the edge set into ones that are crossed by a certain line and to ones that are not and then claim (usually by induction) that the number of edges not crossed by the line is small (see, e.g., [1, 3, 4, 6, 7, 8]). Fulek and Suk [1] studied geometric graphs that do not contain two disjoint copies of a certain geometric pattern. If there is a constant c such that an n-vertex geometric graph with at least cn edges must contain one copy of a certain geometric pattern, then a graph avoiding two disjoint copies of this pattern is cn-near bipartite and hence by Theorem 1 has $O(n^{3/2})$ edges. However, this bound is inferior to the $O(n \log n)$ bound found for this case in [1].

2 Proof of Theorem 1

First we show that every k-near bipartite graph has a subgraph in which the degree of every vertex is $O(k)$ and the number of edges is high. When $k(n) = \Omega(n)$ the graph itself satisfies this property. For $k(n) = o(n)$ we use the following lemma, whose proof we postpone.

Lemma 2.1. *There are constants $c, d > 0$ such that the following holds. Let $G = (V, E)$ be a geometric graph on n vertices that is k-near bipartite. Then there exists a subgraph of G that has at least $c|E| - O(n)$ edges and the degree of each of its vertices is at most dk.*

Theorem 1 follows immediately from Lemma 2.1 and the following theorem.

Theorem 2. *Let $d > 0$ be a constant and let G be an n-vertex k-near bipartite graph such that the degree of every vertex in G is at most dk. Then there is another constant $a = a(d)$ such that G has at most $an\sqrt{k}$ edges.*

Proof. Call a line ℓ *almost balanced* if each of the two half-planes bounded by ℓ contains at most $(d+1)k$ edges of G. Notice that if ℓ is almost balanced, then there are at most $k + (d+1)k = (d+2)k$ edges of G not crossing ℓ. We first show that there is an almost balanced line with any given slope, and that for every almost balanced line there is another almost balanced line separating almost the same subsets of vertices.

For a non-vertical line ℓ denote by $A(\ell)$ and $B(\ell)$ the vertices of G that are above and below ℓ, respectively. Let $e(U)$ denote the number of edges of G in the subgraph induced by $U \subseteq V$.

Proposition 2.2. *For every line ℓ there is a line ℓ' parallel to ℓ such that ℓ' is almost balanced.*

Proof. The proof is in fact just a continuity argument. Without loss of generality assume that ℓ is horizontal and that $e(B(\ell)) \leq k$. Start translating ℓ upwards keeping track of $e(B(\ell))$. Clearly, this number only increases and changes only when ℓ goes past a vertex of G. There is a first time where this number must be greater than k or else the number of edges of G is at most k and the lemma follows trivially (recall that $k = o(n^2)$). Assume therefore that $e(B(\ell))$ becomes greater than k as ℓ goes above a vertex x. Observe that at that point $e(B(\ell)) \leq (d+1)k$, since the degree of x is at most dk. On the other hand because G is k-near bipartite and $e(B(\ell)) > k$, it must be that $e(A(\ell)) \leq k$. Hence, we can take ℓ' to be this translation of the line ℓ. □

Let ℓ_1, ℓ_2 be an ordered pair of lines (not necessarily avoiding the vertices of G) and let o be their intersection point. The *double wedge* of (ℓ_1, ℓ_2), $\mathrm{dw}(\ell_1, \ell_2)$, is the set of vertices of G that meet the line ℓ_1 when it is being rotated counterclockwise about o until it coincides with ℓ_2.

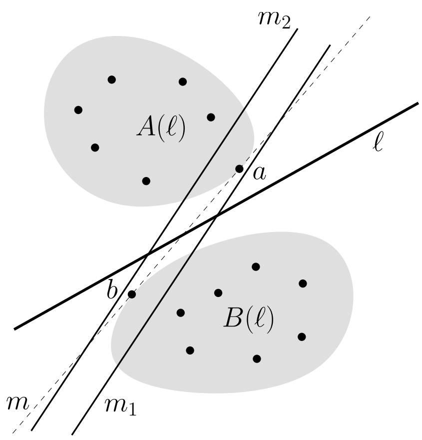

Figure 1: An illustration for the proof of Proposition 2.3.

Proposition 2.3. *Let ℓ be an almost balanced line. Then there exists an almost balanced line ℓ' such that $|dw(\ell, \ell')| = 1$.*

Proof. Let m be the common tangent to the convex hulls of $A(\ell)$ and $B(\ell)$ that separates them such that $|\mathrm{dw}(\ell, m)| = 2$, refer to Figure 1. Let $a \in A(\ell)$ and $b \in B(\ell)$ be the points that determine the line m (thus, $\mathrm{dw}(\ell, m) = \{a, b\}$). By slightly rotating m counterclockwise and translating it, one can obtain two lines m_1, m_2 such that $\mathrm{dw}(\ell, m_1) = \{b\}$ and $\mathrm{dw}(\ell, m_2) = \{a\}$, see Figure 1.

Since G is k-near bipartite $e(A(\ell)) \leq k$ or $e(B(\ell)) \leq k$. Suppose that $e(A(\ell)) \leq k$. Then, since ℓ is almost balanced $e(B(\ell)) \leq (d+1)k$. Observe that m_1 separates $A(\ell) \cup \{b\}$ and $B(\ell) \setminus \{b\}$. Since the degree of b is at most dk it follows that m_1 has at most $(d+1)k$ edges on each of its sides and therefore it is almost balanced. Similarly, if $e(B(\ell)) \leq k$ then m_2 is almost balanced. $\square$

The strategy in the rest of the proof is to find almost balanced lines $\ell_1, \ldots, \ell_t$ with distinct directions, where t is at most $n/\sqrt{dk}$. The number of those edges that are not crossed by at least one of these lines is at most $(d+2)kt$. To estimate from above the number of those edges that cross all lines $\ell_1, \ldots, \ell_t$, observe that such edges have both of their vertices in two "opposite" unbounded faces of the arrangement of lines $\ell_1, \ldots, \ell_t$. We will choose the lines $\ell_1, \ldots, \ell_t$ so that the number of such edges will be small.

We choose the lines ℓ_i one by one. Initially, all the edges and vertices of G are colored blue. Recall that by Proposition 2.2 there is an almost balanced line with any given slope, and let ℓ_1 be an almost balanced line with a very small slope, such that there are no two vertices of G that determine a line with smaller slope than the slope of ℓ_1 (we assume, without loss of generality, that there is no vertical line containing two vertices of G). We recolor all the blue edges of G not crossing ℓ_1 with red.

Suppose that we have already chosen the lines $\ell_1, \ldots, \ell_i$. If the number of remaining blue edges is less than $(2d+2)k$, we stop. Otherwise we choose a new line ℓ_{i+1} such that the number of blue edges with one endpoint in each of the wedges of the double wedge $\mathrm{dw}(\ell_i, \ell_{i+1})$ is $\Theta(kn)$ (in a way that is specified below). All those edges are then colored green. The blue vertices in $\mathrm{dw}(\ell_i, \ell_{i+1})$ are colored red, as well as any blue edge that is adjacent to one of them (note that if such an edge was not colored green, then it does not cross ℓ_{i+1}). See Figure 2 for an example. It is not hard to see that the following invariants are maintained after ℓ_{i+1} is added:

(1) The endpoints of any remaining blue edge are blue.

(2) Every remaining blue edge crosses all the lines $\ell_1, \ldots, \ell_{i+1}$.

(3) Every red edge does not cross at least one of the lines $\ell_1, \ldots, \ell_{i+1}$.

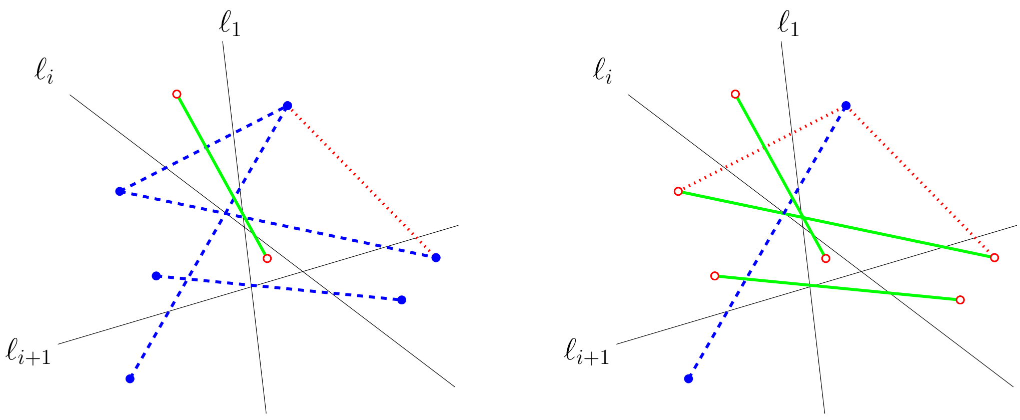

(a) before ℓ_3 is added (b) after ℓ_3 is added

Figure 2: An example for adding a new line. Solid edges are green, dashed edges are blue, and red edges are dotted. Red and blue vertices are represented by empty and full circles, respectively.

(4) Each vertex in $\bigcup_{j=1}^{i} \mathrm{dw}(\ell_j, \ell_{j+1})$ is red. The rest of the vertices are blue.

The line ℓ_{i+1} is chosen using the following proposition.

Proposition 2.4. *Suppose that the number of remaining blue edges is at least $(2d+2)k$. Then there exists an almost balanced line ℓ_{i+1} such that the number of blue edges in the subgraph induced by the vertices in $dw(\ell_i, \ell_{i+1})$ is at least dk and at most $2dk$.*

Proof. We repeatedly apply Proposition 2.3 and find lines m_j, $j = 1, 2, \ldots$, that are almost balanced and at each step the number of vertices of G in $\mathrm{dw}(\ell_i, m_j)$ changes by one. Since the maximum degree is dk, it follows that the number of blue edges in the subgraph induced by $\mathrm{dw}(\ell_i, m_j)$ changes at each step by at most dk. Once the number of these blue edges is at least dk (and is therefore at most $2dk$), we stop and set $\ell_{i+1} = m_j$. Notice that upon stopping the number of blue vertices in $\mathrm{dw}(\ell_i, \ell_{i+1})$ is at least $\sqrt{dk}$. This is because the vertices of every blue edge are both blue.

It remains to show that, unless the number of blue edges is smaller than $(2d+2)k$, we indeed stop at some point and pick ℓ_{i+1}. Suppose we do not, then it follows from Proposition 2.3 that there is an index j such that m_j has a positive slope while m_{j+1} has a negative slope. Let j be the smallest index satisfying this.

Let E_i be the set of blue edges at that point. Notice that all the edges in E_i must cross ℓ_1 and ℓ_i by Invariant (2). Denote by $E_i^{'} \subseteq E_i$ the edges that do not cross m_j. Then $|E_i^{'}| \leq (d+2)k$ since m_j is almost balanced. Let $E_i^{''} \subseteq E_i$ be the blue edges with an endpoint in each of the two wedges of $\mathrm{dw}(\ell_i, m_j)$. Since m_j was not picked as the next line ℓ_{i+1} it follows that $|E_i^{''}| < dk$.

We claim that $E_i = E_i^{'} \cup E_i^{''}$. Suppose there is an edge $e \in E_i \setminus (E_i^{'} \cup E_i^{''})$. Then either e has one endpoint in each wedge of $\mathrm{dw}(\ell_1, \ell_i)$ or e has one endpoint in each wedge of $\mathrm{dw}(m_j, \ell_1)$. Suppose that the latter holds. If e has a negative slope, then its slope is smaller than the slope of ℓ_1, contradicting the choice of ℓ_1. If e has a positive slope, then so does the common tangent that separates $A(m_j)$ and $B(m_j)$ in the proof of Proposition 2.3, and so m_{j+1} should also have a positive slope. Suppose now that e has one endpoint in each of the wedges of $\mathrm{dw}(\ell_1, \ell_i)$, $i > 1$. Consider the left endpoint of e, denote it by v. Then $v \in B(\ell_1)$ and $v \in A(\ell_i)$. Therefore, there must be an index $1 \leq z < i$, such that $v \in B(\ell_z)$ and $v \in A(\ell_{z+1})$. But then $v \in \mathrm{dw}(\ell_z, \ell_{z+1})$ and should be colored red. We conclude that $E_i = E_i^{'} \cup E_i^{''}$ and therefore $|E_i| < (2d+2)k$ and we should have stopped picking lines after ℓ_i was picked. $\square$

We are now ready to complete the proof of Theorem 2. Suppose that the process described above stops after t lines have been chosen. Every edge of G is either blue, red, or green. The number of blue edges is at most $(2d+2)k$. The number of green edges is at most $2dkt$. The number of red edges is at most $(d+2)kt$, since each of the lines we choose is almost balanced and therefore there at most $(d+2)k$ edges that do not cross it. Because we color red at least $\sqrt{dk}$ vertices of G when adding a new line, it follows that $t \leq \frac{n}{\sqrt{dk}}$. Therefore, the number of edges of G is at most $(2d+2)k + (3d+2)\sqrt{\frac{k}{d}} \cdot n = O(\sqrt{k}n)$, since $k = o(n^2)$. This concludes the proof of Theorem 2. □

It remains to prove Lemma 2.1.

Proof of Lemma 2.1: Let $G = (V, E)$ be a geometric graph on n vertices that is k-near bipartite. We will show that there exists a subgraph of G that has at least $|E|/20 - 4n$ edges and the degree of each of its vertices is at most $12k$ (in order to simplify the presentation we do not attempt to optimize these constants).

For every vertex x of G denote by $d(x)$ the degree of x in G. Divide the edges adjacent to x into two sets, those that go to the left and those that go to the right. Color red the $\lceil \frac{1}{10} d(x) \rceil$ edges going to the right from x that have the largest slopes, as well as the $\lceil \frac{1}{10} d(x) \rceil$ edges going to the right from x that have the smallest slopes. Do the same for the edges going to the left from x. The number of edges colored red is at most $\sum_x 4\lceil \frac{1}{10} d(x) \rceil \leq \sum_x 4(\frac{1}{10} d(x) + 1) = \frac{4}{5}|E| + 4n$. Remove all the red edges from G to obtain a subgraph $G_1 = (V, E_1)$ with at least $|E|/5 - 4n$ edges.

Let P denote the set of vertices whose degree in G_1 is at least $12k$ and let $Q = V \setminus P$. Of course, if P is empty, then we are done and G_1 is the desired subgraph.

Proposition 2.5. *There is no edge $(x, y) \in E_1$ such that $x, y \in P$.*

Proof. Assume to the contrary that $x, y \in P$ are connected by an edge e in G_1. Since $d_1(x), d_1(y) \geq 12k$ it follows that $d(x), d(y) \geq 20k$. Without loss of generality assume that x is to the left of y and the slope of e is positive. Because e was not colored red as an edge adjacent to x nor as an edge adjacent to y, we can conclude that in G there are at least $20k/10 = 2k$ edges adjacent to x going to the right with a greater slope than the slope of e and there are at least $2k$ edges adjacent to y going to the left with a greater slope than the slope of e. Consider the line ℓ containing e and slightly rotate it counterclockwise around the midpoint of e. Then there are at least $2k$ edges of G in each of the two half-planes bounded by ℓ. This is a contradiction to the assumption that G is k-near bipartite. □

We may assume, without loss of generality, that at least $1/4$ of the edges in G_1 that connect a vertex $p \in P$ and a vertex $q \in Q$ are such that p is to the left of q and the edge (p, q) has a positive slope. Thus, by removing all the other edges connecting a vertex in P and a vertex in Q, we obtain a subgraph $G_2 = (V, E_2)$ such that $|E_2| \geq |E_1|/4 \geq (|E|/5 - 4n)/4 = |E|/20 - n$.

We will show that by removing at most n edges from G_2, we obtain the desired graph. To this end, the following observation will be useful.

Proposition 2.6. *Let $(x, y), (x, y'), (x', y')$ be edges in G_2 such that $x, x' \in P$, $y, y' \in Q$, and the slopes of both (x, y) and (x', y') are greater than the slope of (x, y'). Then (x', y') has a greater slope than (x, y).*

Proof. Suppose that the slope of (x, y) is greater or equal to the slope of (x', y') (see Figure 3(a)). Note that if we slightly rotate clockwise the line containing (x', y') around the midpoint of (x', y') then the resulting line ℓ separates (red) edges in G going right from x with a slope that is greater than the slope of (x, y) and edges going right from x' with a slope is smaller than the slope of (x', y'). However, each of these two sets of edges contains at least $2k$ red edges since $x, x' \in P$, and therefore G does not have a k-light side w.r.t. ℓ, which is a contradiction. □

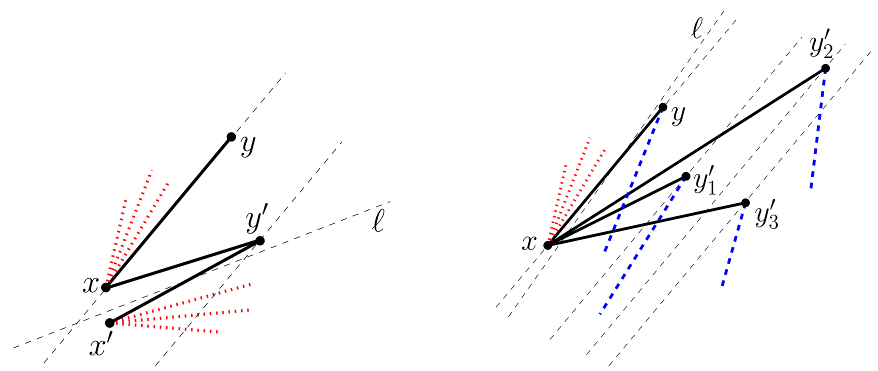

(a) If the slope of (x, y) is greater or equal to the slope of (x', y') then G does not have a k-light side w.r.t. ℓ.

(b) If $d_3(x) \geq 12k$ then G does not have a k-light side w.r.t. ℓ.

Figure 3: Illustrations for the proof of Lemma 2.1.

Next, for every vertex x in G_2 we color blue the edge with the greatest slope that is adjacent to x. Let $G_3 = (V, E_3)$ be the subgraph we obtain by removing all the blue edges from G_2. Then $|E_3| \geq |E_2| - n \geq |E|/20 - 2n$. Denote by $d_3(x)$ the degree in G_3 of a vertex x. We claim that $d_3(x) \leq 12k$ for every vertex x and therefore G_3 is the desired graph.

Suppose that G_3 contains a vertex x such that $d_3(x) \geq 12k$. Therefore, $x \in P$. Let (x, y) be the edge with the greatest slope that is adjacent to x in G_3, and let (x, y') be a different edge (with a smaller slope). It follows from Proposition 2.6 that every edge (x', y') in G_2 has a greater slope than the slope of (x, y). Therefore, for every neighbor y' of x there is at least one blue edge (x', y') whose slope is greater than the slope of (x, y). If we slightly rotate counterclockwise the line containing (x, y) around the midpoint of (x, y), then the resulting line ℓ separates these edges and red edges in G that are going right from x and whose slope is greater than the slope of (x, y) (see Figure 3(b)). However, each of these two sets of edges contains at least $2k$ edges, therefore G does not have a k-light side w.r.t. ℓ, which is a contradiction. This concludes the proof of Lemma 2.1. $\square$

Acknowledgment. We thank anonymous referees for several helpful suggestions for improving the presentation of the paper.

References

[1] R. Fulek and A. Suk, On disjoint crossing-families in geometric graphs, *Thirty Essays on Geometric Graph Theory*, J. Pach (Ed.), Springer, forthcoming.

[2] H. Hopf and E. Pannwitz, Aufgabe Nr. 167, *Jahresbericht Deutsch. Math.-Verein.* **43** (1934) p.114.

[3] J. Pach, R. Pinchasi, G. Tardos, G. Tóth, Geometric graphs with no self-intersecting path of length three, *European J. Combinatorics* **25** (2004), no. 6, 793–811.

[4] J. Pach, F. Shahrokhi, and M. Szegedy, Applications of the crossing number, *Algorithmica* **16** (1996), 111–117.

[5] J.W. Sutherland, Lösung der Aufgabe 167, *Jahresbericht Deutsch. Math.-Verein.* **45** (1935) 33–35.

[6] G. Tóth and P. Valtr, Geometric graphs with few disjoint edges, *Discrete Comput. Geom.* **22** (1999), no. 4, 633–642.

[7] P. Valtr, Graph drawings with no k pairwise crossing edges, In G. D. Battista, editor, *Graph Drawing*, volume 1353 of *Lecture Notes in Computer Science*, 205–218, Springer, 1997.

[8] P. Valtr, On geometric graphs with no k pairwise parallel edges, *Discrete Comput. Geom.* **19** (1998), no. 3, 461–469.