

Parity-constrained Triangulations with Steiner points

Victor Alvarez*

October 31, 2012

Abstract

Let $P \subset \mathbb{R}^2$ be a set of n points, of which k lie in the interior of the convex hull $\text{CH}(P)$ of P . Let us call a triangulation T of P even (odd) if and only if all its vertices have even (odd) degree, and pseudo-even (pseudo-odd) if at least the k interior vertices have even (odd) degree. On the one hand, triangulations having all its interior vertices of even degree have one nice property; their vertices can be 3-colored, see [1, 2, 3]. On the other hand, odd triangulations have recently found an application in the colored version of the classic “Happy Ending Problem” of Erdős and Szekeres, see [4].

In this paper we show that there are sets of points that admit neither pseudo-even nor pseudo-odd triangulations. Nevertheless, we show how to construct a set of Steiner points $S = S(P)$ of size at most $\frac{k}{3} + c$, where c is a positive constant, such that a pseudo-even (pseudo-odd) triangulation can be constructed on $P \cup S$. Moreover, we also show that even (odd) triangulations can always be constructed using at most $\frac{n}{3} + c$ Steiner points, where again c is a positive constant. Our constructions have the property that *every* Steiner point lies in the interior of $\text{CH}(P)$.

1 Introduction

Let $P \subset \mathbb{R}^2$ be a set of n points in general position and let $\Gamma : P \rightarrow \{0, 1\}$ be an assignment of parities to the elements of P , where 0 means even and 1 means odd. Let G be a straight-edge plane graph with vertex set P . We say that a vertex $v \in P$ of G is *happy*, with respect to G , if and only if the degree of v in G is of the parity assigned to v by Γ . If a vertex is not happy w.r.t. G we will say that v is *unhappy*. If the graph G is clear from context we will just say that vertices are happy or unhappy without referring to G .

Given P and Γ , the problem of finding plane graphs on P that maximize the number of happy vertices has recently received some attention. In [5] it was shown that one can always construct a tree, a 2-connected outerplanar graph, and a pointed pseudo-triangulation that makes all but at most three vertices happy. For the class of triangulations it was shown that one can always construct one that makes essentially $\frac{2n}{3}$ of its

*Fachrichtung Informatik, Universität des Saarlandes, alvarez@cs.uni-saarland.de.

vertices happy, but a configuration of points with parities was also shown where at least $\frac{n}{108}$ vertices remain unhappy regardless of the chosen triangulation. The construction of this lower bound requires the use of both parities, but the authors pointed out that there are two particular cases that might accept triangulations with as many as $n - o(n)$ happy vertices. These two particular cases are the ones where the parities assigned to the elements of P by Γ are either *all even* or *all odd* respectively. However, showing that in those particular cases $n - o(n)$ vertices can be made happy turned out to be very challenging. In the same paper the authors showed that in the *all-even* case, a triangulation that makes at least $\frac{2n}{3}$ vertices happy can always be constructed. They also showed that in the *all-odd* case a $\frac{10}{13}$ fraction of happy vertices can always be ensured.

These two special cases, all even or all odd, are of significant interest since they have interesting properties and/or applications. For example, it is well known that a connected graph G having all its vertices of even degree is Eulerian. If on top of it G happens to be a triangulation as well, then G is also 3-colorable, see [1, 2] for a general reference on 3-colorable planar graphs. Those two properties are usually considered “nice” in a graph, and they are characterized only by the parity of the degree of its vertices. For 3-colorability of triangulations a slightly weaker characterization is known: T is a triangulation having all its interior vertices of even degree if and only if T is 3-colorable, see [3] for example. The application related to the all-odd case is a little bit more intricate and we refer the reader to [4] where this application is shown.

Let P be as before. We will say that a triangulation T of P is *even*, or *odd*, if and only if the degree of *every* vertex of T is even, or odd respectively. If only the interior vertices of T are even, or odd, we will say that T is *pseudo-even*, or *pseudo-odd* respectively. This defines four kinds of triangulations of P : Even, pseudo-even, odd, pseudo-odd.

In this paper we attack the following problem: Given P and one kind \mathcal{T} of triangulations of the four mentioned above, construct a set $S = S(P, \mathcal{T})$ such that a triangulation of kind \mathcal{T} can *always* be constructed on $P \cup S$.

Thus, the problem attacked in this paper can be seen as the Steiner-point version of the ones regarding triangulations presented in [5]. With this idea in mind, the following lemma is worth noting:

Lemma 1. *There are arbitrarily large sets of points that, without the use of Steiner points, admit neither pseudo-even nor pseudo-odd triangulations.*

Proof. Let P be a set of points like the one shown to the left in Figure 1 where the size of the convex polygon Q shown in gray is $\equiv 1 \pmod 3$. Let $v \in P$ be the only point not in Q . It is clear that in *any* triangulation of P , point v must be adjacent to *every* vertex of Q , that is, without a triangulation of Q , every vertex of Q has degree three.

Now, it is well known that *every* triangulation of a polygon has at least two “ears”, *i.e.*, a triangle formed by three consecutive vertices of the polygon. This means that, regardless of what triangulation of Q we choose, there will always be a vertex of Q whose adjacencies are only its two neighbors in Q and v . Thus no pseudo-even triangulation of P exists. See to the right in Figure 1.




Figure 1: To the left we have a configuration in which all shown adjacencies are forced, and it accepts neither pseudo-even nor pseudo-odd triangulations. To the right we show in red one of the ears of the shown triangulation of Q .

To show that P does not admit a pseudo-odd triangulation either it suffices to show that regardless of what triangulation of Q we choose, there will always be at least one interior point of P having even degree. It is not hard to verify that in an even triangulation, the size of the outer face must be $\equiv 0 \pmod 3$, a proof can be found in [3]. Since $|Q| \equiv 1 \pmod 3$, then Q does not admit an even triangulation, so in *every* triangulation of Q there will be at least two vertices of odd degree, there must be an even number of them. So assume that there is a triangulation of Q in which the only two vertices of odd degree are the two neighbors u, w of v in $\text{CH}(P)$. Thus we could add a point p outside $\text{CH}(P)$, below the edge uw , and add the adjacencies pu, pw . This implies that the set of points $Q' = Q \cup \{p\}$ has an even triangulation, where upw is an ear. But $|Q'| \equiv 2 \pmod 3$, which clearly contradicts the necessary condition on the size of the outer face of an even triangulation. Therefore, in any triangulation of Q there must be at least one interior point $q \in P$ of odd degree. The force adjacency qv implies that q gets even degree in a triangulation of P , which is what we wanted to prove. ■

2 Our contribution

By Lemma 1 the use of Steiner points is sometimes *necessary* if we insist in constructing any of the four kinds of triangulations mentioned before (even, pseudo-even, odd, pseudo-odd). The relevant issue now is not whether we can construct the triangulations we are interested in, but rather with how many Steiner points can we achieve such constructions, the less, the better. The results we are going to show are the following:

Theorem 1. *Let $P \subset \mathbb{R}^2$ be a set of n points in general position where k of them are interior points. Then a set $S = S(P)$ of interior Steiner points of size at most $\lfloor \frac{k+2}{3} \rfloor + 2$ can always be constructed such that $P \cup S$ accepts a pseudo-even triangulation.*

Theorem 2. *Let P be as before. Then a set $S = S(P)$ of interior Steiner points of size at most $\lfloor \frac{n+1}{3} \rfloor + 1$ can always be constructed such that $P \cup S$ admits an even triangulation.*

Theorem 3. *Let $P \subset \mathbb{R}^2$ be a set of n points in general position where k of them are interior points. Then a set $S = S(P)$ of interior Steiner points of size at most $\lfloor \frac{k}{3} \rfloor + c$, with c a positive constant, can always be constructed such that $P \cup S$ accepts a pseudo-odd triangulation.*

Theorem 4. *Let P be as before. Then a set $S = S(P)$ of interior Steiner points of size at most $\lfloor \frac{n-1}{3} \rfloor + c$, with c a positive constant, can always be constructed such that $P \cup S$ admits an odd triangulation.*

The proofs of all theorems will be constructive. The rest of the paper is divided as follows: In Section 3 we start with some preliminaries. In Section 4 and 5 we show algorithms that imply Theorems 1, 2, and Theorems 3, 4 respectively. We close the paper in Section 6 we some observations and conclusions.

3 Pre-processing of P

Let us quickly recall that given a polygon \mathcal{P} , a vertex v of \mathcal{P} is called *reflex* if the internal angle at v is larger than π . We will call it convex otherwise. Also, by a suitable rotation of P we can assume that the vertex v of $\text{CH}(P)$ with the lowest y -coordinate is unique.

The following pre-processing of P will be done: Using v as a pivot we will sort each interior point of P by its angle with respect to v . Let p_1, \dots, p_k , be a labeling, from left to right with respect to this angular order, of the interior points of P . Let p_0, p_{k+1} be the left and right neighbors of v on $\text{CH}(P)$ respectively.

We construct a simple polygon \mathcal{P} from $P \setminus \{v\}$ as follows: We add each edge $p_i p_{i+1}$, for $0 \leq i \leq k$. We call this the lower part of \mathcal{P} and we will denote it by $L(\mathcal{P})$. Also, we consider the edges of $\text{CH}(P) \setminus \{p_0 v, p_{k+1} v\}$ and we call this the upper part of \mathcal{P} , which will be denoted by $U(\mathcal{P})$, see Figure 2.

Next we will scan $L(\mathcal{P})$ from left to right and we will consider the largest polygonal chains that can be formed using consecutive convex vertices of $L(\mathcal{P})$. Note that for each such polygonal chain, the left and right vertices must be reflex vertices of \mathcal{P} , see Figure 3. Now, for each chain, we will make adjacent its two endvertices, thus forming convex polygons Q_j , with $j \geq 0$, from those convex polygonal chains. These convex polygons can be thought as big “ears” hanging from $L(\mathcal{P})$. We will triangulate the rest of \mathcal{P} outside these Q_j ’s arbitrarily, see Figure 2. If there is no convex vertex of \mathcal{P} in $L(\mathcal{P})$, then this arbitrary triangulation of \mathcal{P} is a triangulation of all \mathcal{P} .

This is all the pre-processing that will be done. From here every algorithm will complete a triangulation of \mathcal{P} in its own way.

4 Even and pseudo-even triangulations

The following observation is very useful when working with 3-colorable triangulations:




Figure 2: To the left we have the polygon \mathcal{P} on $n - 1$ vertices in gray. The convex polygons formed by scanning $L(\mathcal{P})$ from left to right are shown dashed. Note that each pair of consecutive convex polygons shares at most one vertex. To the right we see a triangulation $T(\mathcal{P})$ of \mathcal{P} . The dashed edges are the only ones that are not arbitrary.

Observation 1. *Let T be a 3-colored triangulation with outer face C , not necessarily convex, and let u, v, w be three consecutive vertices of C . If the degree of v in T is even, then the color of u is different than the color of w . If the degree of v in T is odd, then both u, w have the same color.*

We can now continue with the algorithm. Let us triangulate each Q_j , if any exists, as follows: Take its vertex with the lowest y -coordinate, breaking ties arbitrarily, and join it with straight-line segments to all other vertices of Q_j , in case that Q_j has more than three vertices. This is usually called a “fan” triangulation. These fan triangulations along the arbitrary triangulation outside the Q_j ’s complete a triangulation $T(\mathcal{P})$ of polygon \mathcal{P} .

It is well-known that triangulation $T(\mathcal{P})$ can be 3-colored, see [6], thus we will do it, and observe that the only colorless vertex is v , the one we used in the beginning to sort P angularly. Clearly, if a triangulation of P is 3-colorable, then it must be *at least* pseudo-even, thus our idea now is to complete a 3-colorable triangulation of P using $T(\mathcal{P})$ and its 3-coloring. So at all time we will use a 3-coloring as a measure of the correctness of our algorithm.

From this point on, our construction is done by case analysis. We will assume without loss of generality that the colors at our disposal are $\{0, 1, 2\}$. Note that as $T(\mathcal{P})$ is already 3-colored, if all the interior vertices of P are colored by only two colors, say $0, 1$, we could use color 2 for v without violating the 3-coloring of $T(\mathcal{P})$, and hence, using the straight-line segments that connect v with each vertex of $L(\mathcal{P})$, we obtain a 3-colorable triangulation of P . Nevertheless, in general it is not going to happen that the interior vertices can be colored using only two colors, hence we need to do something else in such cases. We will proceed in a line-sweep fashion from p_0 to p_{k+1} with respect to the angular order around v .

Let us fix the color of v as the color of the smallest chromatic class in $L(\mathcal{P})$ using the 3-coloring of $T(\mathcal{P})$, and say that color is i^1 without loss of generality, $0 \leq i \leq 2$. Note that the points in $L(\mathcal{P})$ with color i are the ones causing trouble to complete the desired triangulation, hence we will handle those points depending on their kind in \mathcal{P} , namely if they are reflex or convex vertices of \mathcal{P} . We will keep the invariant that, by the time we are processing an interior point p_j , edge vp_j is present, and *all* interior points p_r , with $r < j$, have even degree already, except possibly for p_0 . Also note that by this time, if the degree of p_j is odd it is because p_{j+1} has color i , due to Observation 1.

Let us start now with our case analysis, we will assume that we are currently processing the interior point p_j , $1 \leq j < k$, so again, we will assume that the edge vp_j is already present in the partial triangulation of P constructed so far. We have the following cases:

- (1) Point p_{j+1} is of color i , just as v , and is a reflex vertex of \mathcal{P} . If points p_j and p_{j+2} are of different color then we can just add the edge $p_j p_{j+2}$, since p_{j+1} has already even degree in $T(\mathcal{P})$, see to the left in Figure 3. Thus we can add the edge vp_{j+2} and move to p_{j+2} . If on the other hand, p_j and p_{j+2} have the same color, we introduce one Steiner point s , below $L(\mathcal{P})$, of the third color different than the one of p_j and v , which will be adjacent to p_j, p_{j+1}, p_{j+2} and v , as shown in the middle in Figure 3. Thus we can again move to p_{j+2} and continue.




Figure 3: The point p_j is currently being processed. Point p_{j+1} is of the same color i of v . If p_j and p_{j+2} have the same color, then one Steiner point suffices to be able to move to p_{j+2} . To the right the convex polygon Q is shown in gray. Point p_{j+1} is the pivot of the fan triangulation of Q .

- (2) Point p_{j+1} is of color i again but a convex vertex of \mathcal{P} . If p_j and p_{j+2} have the same color we proceed just as before, introducing one Steiner point s , below $L(\mathcal{P})$ and of the third available color, which will be again adjacent to p_j, p_{j+1}, p_{j+2} and v . We move to p_{j+2} and continue, see in the middle in Figure 3.

¹In the figures, unless otherwise stated, we will use colors $\{i, i + 1, i + 2\} = \{\text{black}, \text{red}, \text{blue}\}$, with arithmetic modulo 3.

So let us assume that p_j and p_{j+2} have different colors, say w.l.o.g. $i + 1$ and $i + 2$ respectively. Note that, as p_{j+1} is a convex vertex of \mathcal{P} , it must be part of one of the convex polygons Q_1, \dots, Q_r , the big “ears”, we constructed in the pre-processing phase. Let us denote this one convex polygon simply by Q , and its rightmost vertex by p_l , $l \geq j + 2$, according to the angular order around v .

We have now the following sub-cases:

- (2.1) Vertex p_{j+1} was used as a pivot in the fan triangulation of Q , see to the right in Figure 3. This in particular means that p_{j+1} is the only vertex of color i in Q . If $l > j + 2$ then we can re-triangulate Q by constructing the fan triangulation of Q with pivot at p_{l-1} , this changes the 3-coloring of Q only between p_{j+1} and p_{l-1} , the former receives color $i + 2$ while the latter is the new unique vertex of Q of color i . Thus the only thing we did was to move the vertex of color i to the right, and hence we can join v to all vertices p_{j+1}, \dots, p_{l-2} using straight-line segments. If $l = j + 2$ then Q is actually a triangle, and everything is still valid, v is connected to $p_{l-2} = p_j$, see to the left in Figure 4.

We now further distinguish between the following cases:

- (2.1.1) Point p_l is of color $i + 1$, p_{l+1} is of color i and p_{l+2} is of color $i + 2$, see in the middle in Figure 4. We introduce two Steiner points s_1, s_2 of color $i + 2, i + 1$ respectively and we will make the following adjacencies: (1) s_1 gets adjacent to p_{l-1}, p_l, p_{l+1} and s_2 , and (2) s_2 gets adjacent to $p_{l-2}, p_{l-1}, s_1, p_{l+1}, p_{l+2}, v$. Observe that these adjacencies can always be done without introducing any crossing. Moreover, note that with two Steiner points we complete the even degree of each point in the region p_j, \dots, p_{l+2} in which there were originally two points of color i . Thus we can move to p_{l+2} and continue.
- (2.1.2) Point p_l and p_{l+1} as before, but p_{l+2} is of color $i + 1$. We will proceed as before except that this time the adjacencies of s_1, s_2 are as follows: (1) s_1 gets adjacent to $s_2, p_{l-1}, p_l, p_{l+1}, p_{l+2}$ and v , and (2) s_2 gets adjacent to p_{l-2}, p_{l-1}, s_1 and v . As before, we also introduce the adjacencies $p_{j+1}v, \dots, p_{l-2}v$ and $p_{l+2}v$. We can again move to p_{l+2} . See to the right in Figure 4 for the final configuration.
- (2.1.3) Point p_l as before but p_{l+1} is of color $i + 2$ instead. Note that in this case, from p_j to p_{l+1} there is only one vertex of color i , namely p_{l-1} , thus we will introduce only one Steiner point s . Also observe that since p_l is a reflex vertex of \mathcal{P} we can add the adjacency $p_{l-1}p_{l+1}$. Finally we make s adjacent to $p_{l-2}, p_{l-1}, p_{l+1}, v$, and we make, as before, v adjacent to p_{j+1}, \dots, p_{l-2} and p_{l+1} . See to the left in Figure 5 for the final configuration.

The following three cases are complementary:

- (2.1.4) Point p_l is of color $i + 2$, p_{l+1} is of color i and p_{l+2} is of color $i + 1$.




Figure 4: If p_{j+1} was used as a pivot to triangulate a convex polygon that can be cut from \mathcal{P} , then we can use p_{l-1} as the new pivot without changing the color of p_j or anything to its left. Note that p_l must be necessarily a reflex vertex of \mathcal{P} . In the middle we see the final configuration in the case that p_{l+1} is of color i and p_{l+2} is of color $i + 2$. To the right we see the final configuration when p_{l+1} is of color i and p_{l+2} is of color $i + 1$.




Figure 5: To the left we see the final configuration in the case that p_{j+1} was a pivot of color i and p_{l+1} is of color $i + 2$. In the middle and to the right we have that, if p_{j+1} of color i was not a pivot and its neighbors have different color from each other, then one of them must necessarily be a pivot, p_{j+2} in this case. So we have to go back and remove some adjacencies that will allow us to introduce the Steiner points appropriately. Quadrilateral \square is shown in gray.

(2.1.5) Point p_l and p_{l+1} as before and p_{l+2} is of color $i + 2$.

(2.1.6) Point p_l as before and p_{l+1} is of color $i + 1$.

However, these cases are the symmetric cases of (2.1.1), (2.1.2) and (2.1.3) respectively, where p_l gets the other possible color, and thus the constructions get shifted colors. The solution, as can easily be verified, is also of shifted colors. The two last figures of Figure 4 and the first one of Figure 5 can be used as reference if we shift their colors.

(2.2) In this case p_{j+1} of color i was not used as a pivot in the fan triangulation of Q . This means that p_j and p_{j+2} are adjacent, since they are part of Q . So

we have two cases depending on whether p_j or p_{j+2} is the pivot in the fan triangulation of Q .

(2.2.1) Point p_j is the pivot in Q . Then consider the points two positions ahead, that is p_{j+3}, p_{j+4} . Let us assume w.l.o.g. that they exist, otherwise $p_{j+2} = p_{k+1}$, which is the right neighbor of v on $\text{CH}(P)$, and p_j, p_{j+1}, p_{j+2} are the last three points that the algorithm will process. It is then easy to verify that two interior Steiner points suffice to finish.

Now, consider the pattern of colors of p_j, \dots, p_{j+4} . The first three colors are fixed $i+1, i, i+2$. If we see $i, i+1$ next, then the pattern matches what we see in case (2.1.4)², which we know how to solve using two Steiner points, but we would also get rid of two points of color i . If we see next $i, i+2$, then we see the same as in case (2.1.5). If p_{j+3} is of a color different than i then it must necessarily be of color $i+1$, since p_{j+2} , of color $i+2$, and p_{j+3} are adjacent. But since p_j , of color $i+1$, is the pivot of Q , then p_{j+3} cannot be a vertex of Q , so p_{j+2} is actually a reflex vertex of \mathcal{P} , and we would find ourselves seeing what we see in case (2.1.6).

(2.2.2) Point p_{j+2} is the pivot in Q . This case is trickier since for the points p_j, \dots, p_{j+4} we could see the pattern of colors $i+1, i, i+2, i+1, i$, which we do not know how to solve locally using only two Steiner points. What we will do is to not look ahead but to see behind.

Since the edge $p_j v$ is currently in the triangulation being built, there is *exactly* one triangle using this edge. Let $q \notin \{p_j, v\}$ be the third vertex of such triangle. Note that q lies to the left of the edge $p_j v$ and hence it already has even degree, moreover, the color of q is $i+2$, since p_j, v have color $i+1, i$ respectively. Now we have the following two cases:

- Vertex q is a Steiner point, or the quadrilateral $\square = q, p_j, p_{j+1}, v$ is convex. Let us consider only the case that \square is convex, if that is not the case then q is a Steiner point and it can be moved as pleased to make \square convex without affecting anything.

We will flip the edge $p_j v$ for the edge qp_{j+1} and introduce one Steiner point s of color $i+1$ that will be adjacent to q, p_{j+1}, p_{j+2}, v , see in the middle of Figure 5.

- If q is not a Steiner point *and* \square is non-convex, then it is not hard to see that the only possibility is $q = p_{j-2}$, and p_{j-1} must then be a reflex vertex of \mathcal{P} of color i . Note then that the edge $e = p_{j-2} p_j$ must be present in the triangulation, by case (1), and thus p_{j-1} is adjacent to no Steiner point. Hence we will remove e and we will introduce one Steiner point s_1 of color $i+2$ that is adjacent to $p_{j-1}, p_j, p_{j+1}, s_2$, where s_2 is another Steiner point of color $i+1$ that is adjacent to

²Which in turn is the symmetric case of (2.1.1).

$p_{j-2}, p_{j-1}, s_1, p_{j+1}, p_{j+2}, v$. We can now move to p_{j+2} and continue. See to the right in Figure 5.

Note that the color i of v was chosen as the color of the smallest chromatic class in $L(\mathcal{P})$, so its cardinality can be at most $\lfloor \frac{k+2}{3} \rfloor$. Now, it is important to observe that we assumed that the point p_j that we process is neither p_0 nor p_{k+1} of $\text{CH}(P)$, since the algorithm started with $j \geq 1$. So it could happen that at least one of those extreme points is of the same color i of v , which would give a conflict in the 3-coloring of the triangulation we are constructing. Let us see how can we deal with this kind of situation. Assume without loss of generality that p_0 is of color i . Then, before start processing the first interior points p_1 , introduce one Steiner point v' inside $\text{CH}(P)$, so close to v that the angular order p_1, \dots, p_k w.r.t. v is also kept by v' . This Steiner point v' will replace v in the algorithm, so it will get color i as well. Now symbolically delete v and run the algorithm. When the algorithm ends we will still have the conflict of the monochromatic edge p_0v' , we could simultaneously have the same conflict with edge $v'p_{k+1}$.




Figure 6: Polygon \mathcal{P} shown in light gray. In the figures color $i = \text{black}$, and the color white means that those points are somehow 3-colored without conflicting with the black points. The visibility region of v is shown in dark gray.

Put back v and remove the conflicting edges from the construction. We will proceed depending on what v sees, having the edges of the construction as obstacles. That is, if edge p_0v' is the only one with conflicts, then v sees the polygonal chain p_0, p_1, v', p_{k+1} , see to the left in Figure 6. If the edge $v'p_{k+1}$ also has conflicts then v sees the polygonal chain $p_0, p_1, v', p_k, p_{k+1}$, see to the right in Figure 6. The solution will depend on whether p_0v' is the only conflict or not, and whether v' has even degree, in the triangulation constructed by the algorithm, or not. The four cases, shown in solid lines, and their solutions, shown with dashed lines, can be seen in Figure 7. Observe that at most one more Steiner point s is used for the solutions, and that v', s are both interior Steiner points. So s can be charged to p_0 , which is one of the $\lfloor \frac{k+2}{3} \rfloor$ points of $L(\mathcal{P})$ of color i , and v' would be simply *one* interior Steiner point that cannot be charged to anything.

Finally, in the analysis of cases (2.1.1) to (2.1.6) we always assumed that point p_{l+1} always existed. This might not always be the case since we could have $p_l = p_{k+1}$, but in this case we could safely assume that $p_{l+1} = v$ or $p_{l+1} = v'$, depending on conflicts on $\text{CH}(P)$, so the second Steiner point we introduce in those cases is also an interior




Figure 7: Polygon \mathcal{P} shown in gray. On the top part we have the solution for the cases where p_0v' is the only conflict and the degree of v' is odd, left, or even, right. Below we have the solutions for the case when both edges $p_0v', v'p_{k+1}$ are in conflict and the degree of v' is odd, left, or even, right. In the figures color $i = \text{black}$.

Steiner point that cannot be charged to anything. Hence the construction used at most $\lfloor \frac{k+2}{3} \rfloor + 2$ interior Steiner points, and Theorem 1 follows.

4.1 Extension to even triangulations

Theorem 1 guarantees a 3-colorable triangulation, which will be at least pseudo-even, but it might not necessarily be completely even, and this is because when we choose an arbitrary triangulation for a part of \mathcal{P} , some vertices of $\text{CH}(P)$ might get odd degree. Thus in order to construct an even triangulation we have to do some more work.

As we mentioned before, in an even triangulation the size of the outer face must be multiple of three. Thus the first thing we will do, if necessary, is to complete $\text{CH}(P)$ using at most two Steiner points so that we fulfill $|\text{CH}(P)| \equiv 0 \pmod 3$.

Let v again the *unique* vertex of $\text{CH}(P)$ with the smallest y -coordinate, and sort all $P \setminus \{v\}$ angularly, from left to right, around v . Let p_1, \dots, p_{n-1} be the points of $P \setminus \{v\}$ in this sorted order.

The main idea behind the construction is to enclose P in a bigger polygon \mathcal{Q} so that all p_0, \dots, p_{n-1} are interior points, and then run the algorithm of Theorem 1, which will guarantee that all p_0, \dots, p_{n-1} will have even degree. The construction is done in such a way that $\text{CH}(P)$ appears in the construction, at the end we can complete the degree of v to an even degree, if necessary, and such that by removing \mathcal{Q} we do not destroy

any parity, since \mathcal{Q} , upon deletion, will take with it an even number of adjacencies per “affected” point of $\text{CH}(P)$, so the degree of those affected points will remain even at the end. So let us jump to the actual construction.

Enclose $\text{CH}(P)$ in a bigger polygon \mathcal{Q} with the following properties: (1) The size of \mathcal{Q} is $|\text{CH}(P)| + 1$, (2) The vertex of \mathcal{Q} with smallest y -coordinate is also v and is unique, (3) The polygon \mathcal{C} formed by $\mathcal{Q} \cup \text{CH}(P)$ can be triangulated using a “zig-zag” starting at v . That is, if we denote the vertices of $\text{CH}(P)$ in clockwise order starting with v by $q_1 = v, q_2, \dots, q_k$, with $k \geq 3$ and $k \equiv 0 \pmod{3}$, and we denote the vertices of \mathcal{Q} the same way by $q'_1 = v, q'_2, \dots, q'_k, q'_{k+1}$, then the “zig-zag” is $q_1 = q'_1 = v, q'_2, q_2, q'_3, \dots, q'_k, q_k, q'_{k+1}, v = q'_1 = q_1$, which along the edges of \mathcal{C} complete a triangulation of \mathcal{C} itself. See to the left in Figure 8.




Figure 8: To the left: The polygon \mathcal{Q} is the outer face of the construction shown. Observe that it does not have to necessarily be convex. The convex hull of P , $\text{CH}(P)$, is shown in dark gray, and \mathcal{C} is shown in light gray, along with its zig-zag triangulation. To the right: The particular 3-coloring of the zig-zag triangulation of \mathcal{C} using colors $\{0, 1, 2\} = \{\text{black}, \text{blue}, \text{red}\}$.

This construction can always be achieved and has the following properties: (1) All points of $P \setminus \{v\}$ lie in the interior of \mathcal{Q} , (2) In the “zig-zag” triangulation of \mathcal{C} , every vertex of $\text{CH}(P)$, and every vertex of \mathcal{Q} , except for q'_2, q'_{k+1} , has even degree, (3) In the 3-coloring of the triangulation of \mathcal{C} that starts with color 0 at $q'_2 \in \mathcal{Q}$, and color 1 at $q_2 \in \text{CH}(P)$, we have that $q_k \in \text{CH}(P)$ receives color 0, vertex $q'_{k+1} \in \mathcal{Q}$ receives color 1, and v receives color 2. See to the right in Figure 8.

Now, pre-process $\mathcal{Q} \cup \mathcal{P}$ as explained in § 3. This pre-processing will construct in a first step the polygon \mathcal{P} used in the algorithm of Theorem 1. The reader should see that this time \mathcal{P} is the polygon whose upper part $U(\mathcal{P})$ is $\mathcal{Q} \setminus \{vq'_2, q'_{k+1}v\}$, and whose lower part $L(\mathcal{P})$ is formed by the adjacencies $p_i p_{i+1}$, for $1 \leq i \leq n - 2$, along with the adjacencies $q'_2 q_2, q'_{k+1} q_k$. In a second step, the “ears” formed by consecutive convex

vertices of $L(\mathcal{P})$ will be computed. Recall that these “ears” are the convex polygons Q_j ’s in the pre-processing phase. Here, by suitably locating $q'_2, q'_{k+1} \in \mathcal{Q}$ we can assume that both $q_2, q_k \in \text{CH}(P)$ are reflex vertices of \mathcal{P} , see Figure 9. Moreover, observe that *any* other vertex $q_i \in \text{CH}(P)$, $2 < i < k$, must also be a reflex vertex of \mathcal{P} . This implies that the convex polygons Q_j ’s are contained in $\text{CH}(P)$. This is important because in a third step of the pre-processing we compute an arbitrary triangulation of \mathcal{P} minus those Q_j ’s. At this point we will choose the zig-zag triangulation of \mathcal{C} as part of this “arbitrary” triangulation, adding other arbitrary edges inside $\text{CH}(P)$ if the zig-zag triangulation of \mathcal{C} does not complete a triangulation of \mathcal{P} minus the Q_j ’s, the important thing here is that the zig-zag triangulation of \mathcal{C} , which contains $\text{CH}(P)$, appears.




Figure 9: Polygon \mathcal{P} shown in gray. The dash lines delimit the “ears” that are constructed by the algorithm. They are contained in $\text{CH}(P)$ since every vertex of $\text{CH}(P)$ is a reflex vertex of \mathcal{P} . To the left we can see a whole 3-colored triangulation of \mathcal{P} , where the zig-zag triangulation of \mathcal{C} appears. The 3-coloring is extended from the 3-coloring of \mathcal{C} .

If we now execute the algorithm of Theorem 1, the first thing it will do is to compute a fan triangulation of each Q_j , if any exists, see to the right in Figure 9. This will complete a triangulation $T(\mathcal{P})$ of polygon \mathcal{P} . The second thing it will do is to compute a 3-coloring of $T(\mathcal{P})$. Clearly we can extend the particular 3-coloring of the zig-zag triangulation of \mathcal{C} , explained before, to a 3-coloring of $T(\mathcal{P})$. Here again the important thing is that this particular 3-coloring of the triangulation of \mathcal{C} appears. The third thing the algorithm will do is to re-color v with the color of the smallest chromatic class of the 3-coloring of $T(\mathcal{P})$. Here we have two cases:

- Vertex v stays of color 2. Observe that since $q'_2, q'_{k+1} \in \mathcal{Q}$, $q_2, q_k \in \text{CH}(P)$, the neighbors of v in \mathcal{Q} and $\text{CH}(P)$ respectively, have colors 0 and 1, then v does not create any conflict with the 3-coloring of $T(\mathcal{P})$. Thus we will keep the adjacencies

q'_2v, q_2v, vq_k and vq'_{k+1} , and we will execute the algorithm to the end. Since the colors of $q'_2, q'_{k+1}, q_2, q_k, v$ are never changed by the algorithm, we end up having a 3-colored triangulation, where v also has even degree, that uses at most $\lfloor \frac{n+1}{3} \rfloor + 1^3$ interior Steiner points. Note that these interior Steiner points are also interior w.r.t. $\text{CH}(P)$, since the algorithm of Theorem 1 would introduce the Steiner points inside \mathcal{Q} , but strictly below the lower part $L(\mathcal{P})$ of \mathcal{P} . Since the adjacencies q'_2v, q_2v, vq_k and vq'_{k+1} are present in the output triangulation, the Steiner points fall strictly inside $\text{CH}(P)$.

The reader can verify that if we now remove $\mathcal{Q} \setminus \{v\}$, along with all the adjacencies that it takes with it, we are left with the desired even triangulation.

- Vertex v changes color. Then we will assume without loss of generality that v gets color 0. If v got color 1 we would have a symmetric conversation. Remember that by the particular 3-coloring of \mathcal{C} we have that $q'_2, q'_{k+1} \in \mathcal{Q}$ have colors 0, 1 respectively, and $q_2, q_k \in \text{CH}(P)$ have color 1, 0 respectively, see Figure 8. Thus the algorithm will introduce the Steiner point v' of color 0, to take the place of v , and will symbolically delete v . This time we can charge v' to q'_2 which is one of the $\lfloor \frac{n+1}{3} \rfloor$ points of $L(\mathcal{P})$ of color 0. Now introduce the adjacency $v'p_{n-2}$ and execute the algorithm until it finishes. There, point p_{n-2} would be the last processed point. So we arrive at the configuration shown in the left upper corner of Figure 10.

In this configuration we necessarily have that the degree of $q_2 = p_1 \in \text{CH}(P)$ is odd, since it is adjacent to $q'_2, v' \in \mathcal{Q}$ which have both color 0. The degree of p_{n-2} is also odd since it is adjacent to $v', q_k = p_{n-1} \in \text{CH}(P)$, and both are colored 0 as well. So the only options we have now are whether v' has even or odd degree between q_2 and p_{n-2} . This is equivalent to consider whether p_{n-2} is of color 2 or 1 respectively. Those two configurations are shown in the right upper corner, and in the middle respectively, in Figure 10 in solid.

Now let us put v back and we will finish the construction as shown in Figure 10 with dashed lines. We might require the use of another Steiner point s , as shown to the right in the upper corner of the same figure, which can be charged to $q_k = p_{n-1}$. In the end v gets even degree as well. It is not hard to verify that the total number of interior Steiner points is again at most $\lfloor \frac{n+1}{3} \rfloor + 1$, that all Steiner points are interior to $\text{CH}(P)$, and that the removal of $\mathcal{Q} \setminus \{v\}$ does not destroy any parity. Hence Theorem 2 follows.

5 Pseudo-odd and odd triangulations

Working locally with (pseudo-)odd triangulations is slightly easier. The following observation was already pointed out in [7]:

³It is not $\lfloor \frac{n+1}{3} \rfloor + 2$ this time since there is no conflict in $\text{CH}(P)$.




Figure 10: Polygon \mathcal{P} is shown in light gray. Using colors $\{0, 1, 2\} = \{\text{black, blue, red}\}$, if we made as if point p_{n-2} was the last point, we arrive at the configuration shown in the left upper corner. The white color of point p_{n-2} means that we do not care about its real color at this time. If we put v back, and we color it with 3, we can add the dashed adjacencies shown in the middle and the right upper corner, depending on the actual color of p_{n-2} . The color of v will conflict with the color of $q'_3, q'_k \in \mathcal{Q}$, but this is not a problem since in the end we will remove $\mathcal{Q} \setminus \{v\}$.

Observation 2. *Let Δ be a triangle in a triangulation. Then at most seven interior Steiner points suffice to obtain an odd-triangulation of Δ .*

The way this odd triangulation of Δ is obtained is shown in Figure 11. Observe that, in general, no similar statement can be done for (pseudo-)even triangulations.

For simplicity we will right now focus on pseudo-odd triangulations only. The algorithm to construct them is essentially the same as for pseudo-even triangulations, however, this time we will not be able to use a 3-coloring as guide. That 3-coloring played an important role in the upper bound on the number of interior Steiner points that we used: We did not have to explicitly track whether we introduced at most one Steiner point every three interior points, but we just had to take care of solving conflicts with the coloring as they appeared. The 3-coloring then ensured such conflicts to appear at most one-third of the time. So for (pseudo-)odd triangulations we will have to explicitly take care of introducing at most one Steiner point per every three interior points.

The pre-processing of P will be again as explained in § 3. Let $v \in P$ be again the unique vertex with smallest y -coordinate, the one we use to sort the interior points




Figure 11: All interior points are Steiner points. Gray vertices are of even parity before the introduction of Steiner points, and all black vertices on the boundary of Δ are of odd parity. The middle case spawns two sub-problems, of the kind shown to the right, by introducing one Steiner point. This gives in total $1 + 3 + 3 = 7$ Steiner points for an odd-triangulation of Δ .

angularly around. The pre-processing phase ends with the construction of the convex polygons Q_j 's, the big ears hanging from the lower part $L(\mathcal{P})$ of the therein created polygon \mathcal{P} . For pseudo-even triangulations we would fan-triangulate those Q_j 's, if any. For pseudo-odd we will choose different triangulations, the main idea is the following: If *every* vertex of a convex polygon Q_j is of even degree in a triangulation $T(\mathcal{P})$ of \mathcal{P} , then we could just directly make them adjacent to v and leave them all odd. Unfortunately this might not always be possible. We will try nonetheless to achieve something similar, and the following result shown in [5] will be very useful:

Lemma 2 (O. Aichholzer *et al.*). *Let Q be a convex polygon where each of its vertices has a parity assigned. Let p, q, r be any three consecutive vertices of Q . Then there exists a triangulation of Q that makes all vertices of Q happy with the possible exception of p, q, r .*

After the pre-processing there is an arbitrary triangulation of \mathcal{P} minus the Q_j 's. We will complete a triangulation $T(\mathcal{P})$ of \mathcal{P} as follows: for each one of the Q_j 's set the parity of its vertices appropriately so that we can apply Lemma 2 and make them all even in $T(\mathcal{P})$, with the possible exception of its last three vertices, w.r.t. the angular order of the interior points around v . The triangulations of those Q_j 's, however, might change during the algorithm depending on what conflicts we encounter.

Having this particular triangulation $T(\mathcal{P})$ of \mathcal{P} we will again scan $L(\mathcal{P})$ from the left to right, following the angular order p_1, \dots, p_k of the k interior points of \mathcal{P} around v , so again p_0 and p_{k+1} are the neighbors of v in $\text{CH}(\mathcal{P})$. We will again assume that by the time we are processing point p_j , the adjacency $p_j v$ is already present and every point p_r , $r < j$, is of odd degree in the current construction. Nevertheless, observe that since we want to add roughly $\frac{k}{3}$ interior Steiner points, processing p_j actually means processing p_j, p_{j+1}, p_{j+2} , so we have different cases depending on the local situation. That is, if we

are currently stuck at p_j it is because its current degree is even, otherwise we could just continue. So the cases we have to study are the triples of parities: (e, e, e) , (e, e, o) , (e, o, e) and (e, o, o) , where they correspond entry-wise to the current parities of (p_j, p_{j+1}, p_{j+2}) , and e, o stand for even and odd respectively. It is very important to keep in mind that in the triple (p_j, p_{j+1}, p_{j+2}) , the only point adjacent to v is p_j . Also, in *all* cases we will assume that p_{j+3} exists. If that is not the case then one of p_j, p_{j+1}, p_{j+2} is p_{k+1} , and we would find ourselves with a problem of constant size that can be solved using a constant number of interior Steiner points, due to Observation 2. We will now jump to the case distinction.

- (1) (e, e, e) . Regardless of whether p_j, p_{j+1}, p_{j+2} are reflex or convex vertices of \mathcal{P} , the configuration is shown in solid to the left in Figure 12 and its solution is shown dashed.




Figure 12: In the figures, the colors represent the parity of the vertices *before* the adjacencies of the solution are added. Gray color means even degree, black color means odd degree, and white means that we do not necessarily take care of that point at this step. The original configurations are shown in solid black while their solutions are shown dashed. Point s is a Steiner point.

- (2) (e, o, e) . Regardless of whether p_j, p_{j+1}, p_{j+2} are reflex or convex vertices of \mathcal{P} , look the middle configuration of Figure 12.
- (3) (e, e, o) . If p_{j+2} is a reflex vertex of \mathcal{P} , then the situation is shown to the right in Figure 12.

If p_{j+2} is a convex vertex of \mathcal{P} , then p_{j+2} is part of some convex polygon Q we applied Lemma 2 on. Observe that Q cannot be a triangle, otherwise p_{j+2} would be the middle vertex and therefore p_{j+2} would have even degree, since its two neighbors in Q would be adjacent. Thus Q must have size at least four. This implies that p_j cannot be part of Q either, because otherwise we could safely assume that p_j , as any other vertex of Q to the left of p_j , is happy due to Lemma 2. This means that p_{j+1} is necessarily a reflex vertex of \mathcal{P} , and thus the leftmost vertex of

Q . Finally, this all means that Q cannot actually have size larger than four, *i.e.*, Q must necessarily be a convex quadrilateral. If Q had size larger than four we could safely assume that the degree of p_{j+2} in the triangulation of Q is even due to Lemma 2 and the fact that neither p_{j+1} nor p_{j+2} have been modified by the algorithm before. Therefore the situation is as pictured in the upper left corner of Figure 13 and its solution is shown right below.




Figure 13: The colors represent the parity of the vertices *before* the adjacencies of the solution are added. Gray color means even degree, black color means odd degree, and white means that we do not necessarily take care of that point at this step. The original configurations are shown in solid black while their solutions are shown dashed. Point s is a Steiner point.

- (4) (e, o, o) . This case is the hardest of all. If p_{j+1} is a reflex vertex of \mathcal{P} , then the situation is shown in the second figure in the upper part of Figure 13, and whose solution is shown right there dashed. If p_{j+1} is a convex vertex of \mathcal{P} , then observe that p_{j+2} cannot be a convex vertex as well, since otherwise p_j, \dots, p_{j+3} are part of the same convex polygon, and p_j would be happy due to Lemma 2. So p_{j+2} is a reflex vertex of \mathcal{P} . More, if the degree of p_{j+3} in $T(\mathcal{P})$ is odd then the solution is like shown to the right in Figure 13.

Thus we enter a case where p_{j+2} is a reflex vertex of \mathcal{P} and p_{j+3} is of even degree in $T(\mathcal{P})$. This case cannot be solved locally, considering only p_j, \dots, p_{j+3} , and there is more than one way we can deal with it. The way we will do it here is the following: We will go ahead until p_{j+6} and we will solve p_j, \dots, p_6 with at most *two* Steiner points, one of them will be charged to the triple p_j, p_{j+1}, p_{j+2} , and the other to the triple $p_{j+3}, p_{j+4}, p_{j+5}$. At the end of the construction we will be left with the adjacency $p_{j+6}v$, which is where the algorithm will continue. So for all what follows we will assume that p_{j+4}, \dots, p_{j+6} exist, otherwise $p_{j+3} = p_{k+1}$, and we find ourselves with the last four points the algorithm would process. This can be solved with a constant number, larger than one, of Steiner points, which turns into a constant overhead overall.

Since we are assuming that p_{j+3} is currently even, we have four main sub-cases that depend on the pair of parities (\cdot, \cdot) of (p_{j+4}, p_{j+5}) in $T(\mathcal{P})$. The configurations and solutions are pictured in Figure 14.




Figure 14: The colors represent the parity of the vertices *before* the adjacencies of the solution are added. Gray color means even degree, black color means odd degree, and white means that we do not necessarily take care of that point at this step. The original configurations are shown in solid black while their solutions are shown dashed. Points s_1, s_2 are Steiner points.

- (4.1) (e, e) . The configuration is shown in the left upper corner of Figure 14 in solid and its solution is shown dashed. Observe that for the solution it does not play a role whether $p_{j+3}, p_{j+4}, p_{j+5}$ are reflex or convex vertices of \mathcal{P} .
- (4.2) (e, o) . The configuration is shown in the second figure on the upper part of Figure 14 in solid and its solution is shown dashed. Again, it does not matter whether $p_{j+3}, p_{j+4}, p_{j+5}$ are reflex or convex vertices of \mathcal{P} .
- (4.3) (o, o) . If p_{j+4} is a reflex vertex of \mathcal{P} , the configuration is shown in solid to the right on the upper part of Figure 14, and its solution is shown dashed.

Now let us argue that p_{j+4} cannot be a convex vertex. If p_{j+4} were a convex vertex it would be part of a convex polygon Q we applied Lemma 2 on in the beginning. This polygon Q cannot be a triangle, because then p_{j+4} would be the middle point and it would then have even degree in $T(\mathcal{P})$. Also, Q cannot have at least five sides because by Lemma 2, again, p_{j+4} could be

made of even degree in a triangulation of Q , without affecting the parity of p_{j+3} in the same triangulation. Thus the only case that kind of makes sense is that Q is a convex quadrilateral. So, there are two possible options for Q . Either $Q = p_{j+2}, p_{j+3}, p_{j+4}, p_{j+5}$ or $Q = p_{j+3}, p_{j+4}, p_{j+5}, p_{j+6}$. The former is not possible because p_{j+2} would be part of Q and the algorithm would have chosen in the beginning to leave p_{j+2} of even degree, instead of odd, using a triangulation of Q . This is achievable using Lemma 2. The latter is also not possible because although p_{j+3} is currently of even degree, p_{j+4}, p_{j+5} are of odd degree in the current triangulation of Q . Since p_{j+4}, p_{j+5} are the two middle vertices of Q , and Q has only two triangulations, then at least one of them would get even degree, and again, we would be discussing a different case. Therefore p_{j+4} must necessarily be a reflex vertex of \mathcal{P} .

- (4.4) (o, e) . If p_{j+4} happens to be a reflex vertex of \mathcal{P} , then the configuration is shown to the left on the lower part of Figure 14, and its solution is shown dashed.

If p_{j+4} is a convex vertex of \mathcal{P} , then its part of a convex polygon Q we applied Lemma 2 on in the beginning. By the same arguments as in case (o, o) we know that Q can be neither a triangle nor a convex polygon larger than four. Nevertheless, this time Q can indeed be the quadrilateral $p_{j+3}, p_{j+4}, p_{j+5}, p_{j+6}$. The other possibility for Q is in this case not possible either due to the same argument. Since Q is a convex quadrilateral and p_{j+4} is of odd degree in the current triangulation of Q , then the only possibility is that the diagonal $p_{j+4}p_{j+6}$ is present. The configuration is shown to the right on the lower part of Figure 14 in solid and its solution is shown dashed.

This concludes the case analysis.

From all cases discussed it is clear that the algorithm introduces at most one interior Steiner point per triple of interior points of P , and at the end it might require to brute-force a configuration of constant size, that due to Observation 2 can be solve using a constant number c of interior Steiner points as well. Hence overall the algorithm makes use of at most $\lfloor \frac{k}{3} \rfloor + c$ interior Steiner points and Theorem 3 follows.

5.1 Extension to odd triangulations

The extension of the algorithm of Theorem 3 to odd triangulations is now very easy. We again enclose the set of points P in a bigger polygon, just as we did for even triangulations. What the configuration looks like can be seen to the left in Figure 8 on page 12. We then run the algorithm for pseudo-odd triangulations, and at the end the only thing that can happen is that the pivot $v \in P$ gets even degree. However, by Observation 2, pivot v can be made odd by locally adding a constant number of Steiner points in one of its adjacent triangles. Thus, in total, the number of used interior Steiner points is at most $\lfloor \frac{n-1}{3} \rfloor + c$, for some constant c . The odd triangulation is finally obtained by removing the bigger polygon enclosing P . Therefore Theorem 4 follows.

6 Conclusions

In this paper we have presented algorithms that construct, with help of Steiner points, (pseudo-)even and (pseudo-)odd triangulations of a given set of points P . The number of Steiner points that the algorithms use is roughly $\frac{k}{3}$ for the pseudo variants, where k is the number of interior points of P , and roughly $\frac{n}{3}$ for the other cases. It is important to observe that our algorithms *do not* modify the convex hull of P , and therefore they preserve extent measures of P , such as diameter, width, among others. If we do not care about modifying $\text{CH}(P)$, or about the position of the Steiner points, then the task is in general significantly easier. For example, *only two* Steiner points far away from $\text{CH}(P)$ would suffice to construct a pseudo-even triangulation, say one at ∞ and the other at $-\infty$. Albeit being this construction possible, we do not know why it would be interesting to use it, since the output set of points does not look anything like the one that was given as the input.

We could believe that the techniques shown here could be pushed further to improve the overall number of Steiner points, say to go from one-third to one-sixth, but this will imply a significantly larger number of cases to analyze, we see this really as a secondary interesting improvement. What we really see as the primary open problems are the following:

- Is it possible to always construct (pseudo-)even or (pseudo-)odd triangulations of a given set of points P using only a *constant* number of (interior) Steiner points? In other words, how big is the lower bound on the number of (interior) Steiner points that are required to *always* construct such triangulations? We have failed so far trying to prove a (sub-)linear lower bound, which is the natural guess when working on these kind of problems.
- Moving away from using Steiner points, would it be possible to construct even or odd triangulations of a given set of points P where at most a constant number of points remain unhappy? This question was posed by Aichholzer *et al.* in [5]. In that same paper they proved a lower bound of roughly $\frac{n}{108}$ unhappy vertices when the assignment of parities is not uniform. Thus, as they pointed out, the interesting cases are all even and all odd.

Note that although these two questions look similar, they might not be equivalent. If we are interested in an even triangulation of P and we construct one where a constant number of vertices remain unhappy, those unhappy vertices might be far from each other. This means that we might have to bring them together, at least in pairs, using Steiner points. We have to get rid of them in pairs, at least, since the number of odd vertices is always even. But to get them close to each other we might require more than a constant number of Steiner points. Thus the techniques to solve those two open problems might be rather different. We see a real challenge there.

7 Acknowledgement

The author would like to thank Jorge Urrutia and Marco Heredia for suggesting the problem and for helpful comments in the initial stage of this work, and to Raimund Seidel for valuable feedback.

This research was partially supported by CONACYT-DAAD of México.

References

- [1] P. J. Heawood, “On the four-color map theorem,” *Quart. J. Pure Math.*, vol. 29, pp. 270–285, 1898. 1, 2
- [2] R. Steinberg, “The state of the three color problem,” in *Quo Vadis, Graph Theory? A Source Book for Challenges and Directions* (J. W. K. John Gimbel and L. V. Quintas, eds.), vol. 55 of *Annals of Discrete Mathematics*, pp. 211 – 248, Elsevier, 1993. 1, 2
- [3] K. Diks, L. Kowalik, and M. Kurowski, “A new 3-color criterion for planar graphs,” in *WG* (L. Kucera, ed.), vol. 2573 of *Lecture Notes in Computer Science*, pp. 138–149, Springer, 2002. 1, 2, 3
- [4] O. Aichholzer, T. Hackl, C. Huemer, F. Hurtado, and B. Vogtenhuber, “Large bichromatic point sets admit empty monochromatic 4-gons,” *SIAM J. Discrete Math.*, vol. 23, no. 4, pp. 2147–2155, 2010. 1, 2
- [5] O. Aichholzer, T. Hackl, M. Hoffmann, A. Pilz, G. Rote, B. Speckmann, and B. Vogtenhuber, “Plane graphs with parity constraints,” in *WADS* (F. K. H. A. Dehne, M. L. Gavrilova, J.-R. Sack, and C. D. Tóth, eds.), vol. 5664 of *Lecture Notes in Computer Science*, pp. 13–24, Springer, 2009. 1, 2, 16, 21
- [6] S. Fisk, “A short proof of chvátal’s watchman theorem,” *J. Comb. Theory, Ser. B*, vol. 24, no. 3, p. 374, 1978. 5
- [7] A. Pilz, “Parity properties of geometric graphs,” Master’s thesis, Graz University of Technology, 2009. 14