Reconfiguring Triangulations with Edge Flips and Point Moves*

Greg Aloupis[1], Prosenjit Bose[2], and Pat Morin[2]

[1] School of Computer Science, McGill University
athens@cgm.cs.mcgill.ca
[2] School of Computer Science, Carleton University
{jit,morin}@scs.carleton.ca

Abstract. We examine reconfigurations between triangulations and near-triangulations of point sets, and give new bounds on the number of *point moves* and *edge flips* sufficient for any reconfiguration. We show that with $O(n \log n)$ edge flips and point moves, we can transform any geometric near-triangulation on n points to any other geometric near-triangulation on n possibly different points. This improves the previously known bound of $O(n^2)$ edge flips and point moves.

1 Introduction

An *edge flip* is a graph operation that is defined on (near)-triangulations [3]. An edge flip on a triangulation is simply the deletion of an edge, followed by the insertion of another edge such that the resulting graph remains a triangulation. The definition of an edge flip gives rise to several natural questions: Does there always exist a sequence of flips that reconfigures a given triangulation to any other triangulation? Are there bounds on the lengths of such sequences if they exist? Can these sequences be computed? These questions have been studied in the literature in many different settings. In particular, Wagner[19] proved that given any two n-vertex triangulations G_1 and G_2, there always exists a finite sequence of edge flips that reconfigures G_1 into a graph isomorphic to G_2. Subsequently, Komuro[10] showed that in fact $O(n)$ edge flips suffice. Recently, Bose et al.[2] showed that $O(\log n)$ simultaneous edge flips suffice and are sometimes necessary. This setting of the problem is referred to as the combinatorial setting since the triangulations are only embedded combinatorially, i.e. only the cyclic order of edges around each vertex is defined.

In the geometric setting, the graphs are embedded in the plane with edges represented by straight line segments. Pairs of edges can only intersect at their endpoints. Edge flips are still valid operations in this setting, except that now the edge that is added must be a line segment that cannot properly intersect any

* Research supported in part by the Natural Science and Engineering Council of Canada.

[3] A triangulation is a plane graph where every face is a triangle. In a *near*-triangulation, the outer face may not be a triangle.

of the existing edges of the graph. This implies that there are valid edge flips in the combinatorial setting that are no longer valid in the geometric setting. Lawson[12] showed that given any two geometric near-triangulations N_1 and N_2 embedded on the same n points in the plane, there always exists a finite sequence of edge flips that transforms the edge set of N_1 to the edge set of N_2. Hurtado, Noy and Urrutia[9] showed that $O(n^2)$ flips are always sufficient and that $\Omega(n^2)$ flips are sometimes necessary.

Note that in the geometric setting, only the near-triangulations that are defined on the *specified point set* can be attained via edge flips. For example, no planar K_4 can be drawn on a convex set of four points without introducing a crossing.

In order to resolve the discrepancy between the combinatorial and geometric settings, Abellanas et al.[1] introduced a geometric operation called a *point move*. A point move on a geometric triangulation is simply the modification of the coordinates of one vertex such that after the modification the graph remains a geometric triangulation. That is, the move is valid provided that after moving the vertex to a new position, no edge crossings are introduced . They also showed that with $O(n^2)$ edge flips and $O(n)$ point moves, any geometric triangulation on n points can be transformed to any other geometric triangulation on n possibly different points.

The question which initiated our investigation is whether or not $O(n^2)$ edge flips are necessary. In this paper, we show that with $O(n \log n)$ edge flips and point moves, we can transform any geometric near-triangulation on n points to any other geometric near-triangulation on n possibly different points. Next, we show that if we restrict our attention to geometric near-triangulations defined on a fixed point set of size n, the problem is just as difficult even with the use of point moves. Finally, we show that with a slightly more general point move, we can remove the extra log factor from our main result.

2 Results

In the remainder of the paper, all triangulations and near-triangulations are geometric. It is assumed that the outer face any given near-triangulation is convex, and that any two near-triangulations involved in a reconfiguration have the same number of points on the convex hull.

We assume that the n vertices of any given triangulation are in general position. It is not difficult to see that $O(n)$ point moves can reconfigure a triangulation to this form. We begin with some basic building blocks that will allow us to prove the main theorems.

Lemma 1. **[2]** *A reconfiguration between two triangulations of the same point set that is in convex position can be done with $O(n)$ edge flips.*

Lemma 2. **[9]** *Let v_1, v_2 and v_3 be three consecutive vertices on the outer face of a near-triangulation T_1. Let C be the path from v_1 to v_3 on the convex hull of all vertices but v_2. A near-triangulation T_2 containing all edges of C may be* *constructed from T_1 with t edge flips, where t is the number of edges initially intersecting C in T_1.*

Lemma 3. *Given a near-triangulation T, any vertex $p \in T$ with degree $d > 3$ that is inside the convex hull of the vertices of T can have its degree reduced to 3 with $d - 3$ edge flips.*

Proof. Let P be the polygon that is the union of all triangles incident to p. By Meister's *two-ears theorem* [13], if P has more than three vertices, then it has at least two disjoint ears[4]. At most one of them can contain p. Therefore p and one of the ears form a convex quadrilateral. We may flip the edge from p to the tip of the ear, effectively cutting the ear from P and reducing the number of vertices of P by one. This process may be continued until P is reduced to a triangle that contains p as desired. □

Lemma 4. *Given a near-triangulation T, any vertex $p \in T$ with degree 3 that is inside the convex hull of the vertices of T can be moved to a new position in the triangulation along a straight path crossing t edges, using at most $2t$ edge flips and $2t + 1$ point moves, assuming the path does not cross through any vertices.*

Proof. Suppose that p is joined by edges to vertices v_1, v_2 and v_3. Without loss of generality, let edge v_2v_3 intersect the path that p must follow, and let this path continue into triangle $v_2v_3v_4$, as shown in Figure 1.

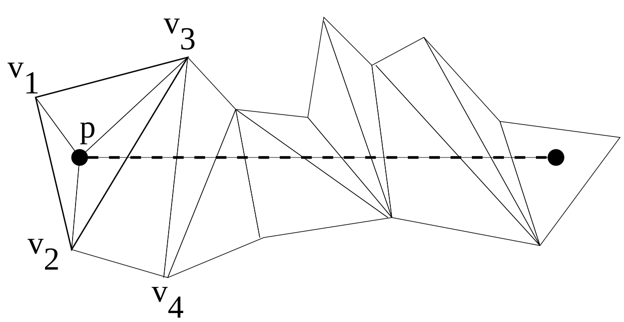

Fig. 1. A vertex p and a straight path that it must move along (dashed). p can pass through any edge with two edge flips.

Clearly p can be moved anywhere within triangle $v_1v_2v_3$ without the need of any edge flips. Then it can be moved along its path, as close to edge v_2v_3 as necessary, so that the quadrilateral $pv_2v_3v_4$ becomes convex. This allows edge v_2v_3 to be flipped into edge pv_4. Now p may continue along its path. As soon as it enters $v_2v_3v_4$, edge pv_1 may be flipped into v_2v_3. Now, with two edge flips and

[4] A triangle, defined by three consecutive vertices of a polygon, is an ear if it is empty and the vertices form a convex angle. The second vertex is the *tip* of the ear.

two point moves, p has crossed through the first edge intersecting its path, and still has degree 3. By the same argument, p may traverse its entire path with two edge flips and two point moves for each intersecting edge. One additional point move is required in the last triangle. Note that only three edges in the original and final triangulations will be different. □

Lemmata 3 and 4 imply the following result:

Lemma 5. *Given a near-triangulation T, any vertex in the interior of the convex hull of the vertices of T with degree d can be moved to a new position in the triangulation along a path crossing t edges, using $O(d + t)$ edge flips and point moves.*

Lemma 6. *An edge can be constructed between a convex hull vertex and any other vertex in a triangulation using $O(n)$ edge flips, with the aid of one moving point that is moved $O(n)$ times.*

Proof. Let v_1 be the hull vertex. First suppose that the second vertex is an interior point. Then it will play the role of the moving point, and we will label it p. We can move p directly towards v_1, until it is located within a triangle that has v_1 as a vertex. Now v_1 and p must be joined with an edge. Next we move p back along the same line to its original position, always maintaining edge v_1p. To do this, we consider the set of triangles that intersect p's path, as in Lemma 4. The point p can always enter a triangle intersecting the path back to its original location. The difference is that once it has crossed an intersecting edge, we do not restore the edge. This means that p will accumulate edge degree. An issue that needs to be taken care of is that of maintaining a triangulation when p is about to lose visibility to another vertex. This occurs when one of its incident edges is about to overlap with another edge in the triangulation, as shown in Figure 2.

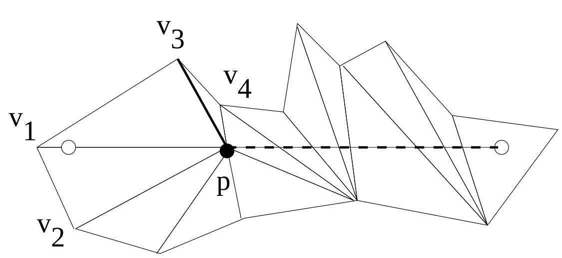

Fig. 2. Maintaining a triangulation while extending edge v_1p: p has moved from a position close to v_1 (shown white), and still has to traverse the dashed segment to its original position. Edge pv_3 causes a problem if p is to continue.

Suppose that edge pv_3 is about to overlap with edge v_3v_4. Vertices v_3 and v_4 cannot be on opposite sides of the remaining path that p must traverse, otherwise v_3v_4 may be flipped. The point p must share an edge with v_4 in this configuration. Points p and v_3 are also part of another triangle, along with some vertex v^* which may be anywhere on the path from v_1 to v_3. These two triangles must form a convex quadrilateral $pv^*v_3v_4$, otherwise p would have already lost visibility to v^*. Thus pv_3 may be flipped into v_4v^*, which means that v_3 is removed from the polygon that intersects p's path. The result is that when p reaches its original position, it leaves a *fan*[5] behind it, which includes edge v_1p. Overall one edge flip is used when p enters a new triangle, and at most one flip is used for every edge that attaches to p.

If both vertices of the edge that we wish to construct are on the hull, then we can take any point p within the hull and move it close to v_1 and onto the segment between the two hull vertices. p can then move along this segment to the second hull vertex until it is connected to both. At this moment, p may be perturbed so that the three vertices form a triangle. This triangle might contain other edges incident to p. Lemma 2 implies that these edges may be removed so that the desired edge can be constructed with $O(n)$ edge flips. □

2.1 Triangulations

With the basic building blocks in place, we now prove one of our main results.

Theorem 1. *With $O(n \log n)$ edge flips and point moves, we can transform any geometric triangulation on n points to any other geometric triangulation on n possibly different points.*

Proof. We transform one triangulation to another via a canonical configuration. As shown in Figure 3, the interior vertices form a *backbone* (i.e. their induced subgraph is a path). The top of the backbone is joined to the topmost hull vertex v_1, and all interior vertices are joined to the other two hull vertices, v_L and v_R.

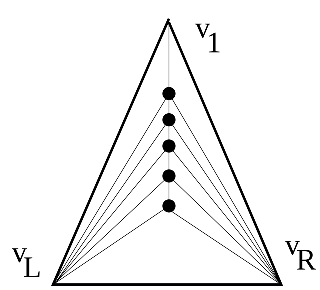

Fig. 3. The canonical configuration used for triangulations.

The canonical configuration is constructed in a divide-and-conquer manner. We perform a radial sweep from v_1, to find the median vertex interior to the

[5] A fan is a star-shaped polygon with a vertex as its kernel.

convex hull, v_M. After constructing edge $v_1 v_M$ we move v_M directly away from v_1 towards the base $v_L v_R$, maintaining $v_1 v_M$ until triangle $v_M v_L v_R$ contains no interior points. By Lemma 6, we use $O(n)$ operations to accomplish this. Now, we transform $v_1 v_M v_L$ and $v_1 v_M v_R$ into backbone configurations by induction since they are smaller instances of the same problem. The resulting configuration is shown in Figure 4.

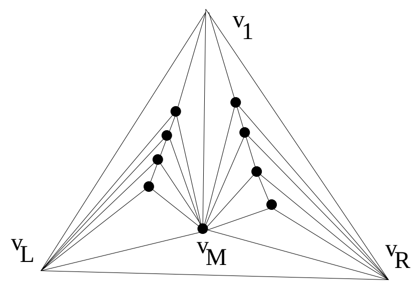

Fig. 4. The configuration of a triangulation prior to merging the backbones on each side of the median vertex v_M.

We now show that the two sides may be merged using $O(n)$ operations. As shown in Figure 5a, we first move the lowest vertex of a backbone into a position that is close to the base and is along the extension of edge $v_1 v_M$. This requires one edge flip. The vertices on the left/right backbones are processed in ascending order, and are always moved just above the previous processed vertex, as shown in Figure 5b. Each vertex will require two point moves and one edge flip. Thus $v_1 v_L v_R$ is reconfigured into canonical form, and by a simple recurrence the number of edge flips and point moves used is $O(n \log n)$. It is trivial to move a canonical triangulation to specific coordinates using n point moves. Thus the transformation between any two triangulations may be completed. □

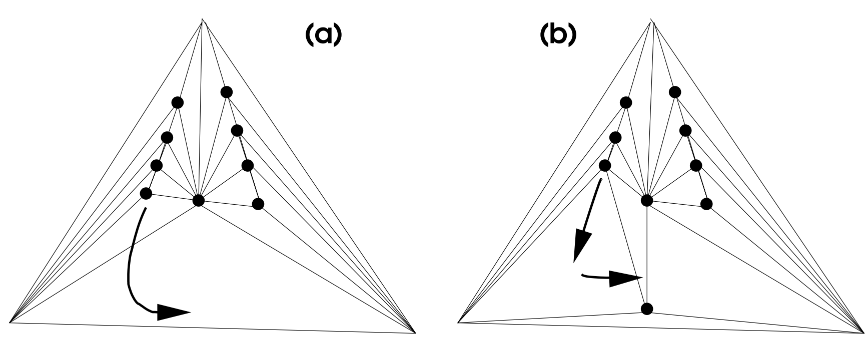

Fig. 5. Merging two backbones into one.

2.2 Near-triangulations

If the initial graph is a near-triangulation, Theorem 1 does not directly apply. Some care must be taken to handle a non-triangular outer face. Details are given in the proof of the following theorem:

Theorem 2. *With $O(n \log n)$ edge flips and point moves, we can transform any geometric near-triangulation on n points to any other geometric near-triangulation on n possibly different points.*

Proof. As in the case with triangulations, we transform one near-triangulation to another via a canonical configuration. In the primary canonical configuration, shown in Figure 6, one chosen hull vertex (v_1) is joined by chords to all other hull vertices. Thus v_1 is in the kernel of a convex fan. Every triangle of the fan, except for one, is empty. All interior vertices, located in the non-empty triangle T, are in the canonical configuration of a triangulation.

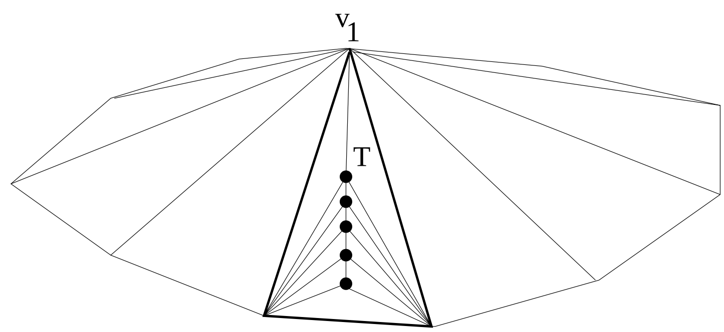

Fig. 6. The primary canonical configuration used for near-triangulations.

We first construct all edges of the top-level fan configuration, leaving interior vertices in their original positions. Then within each triangle of the fan, we rearrange the interior vertices into a canonical triangulation. Finally, we merge all triangles of the fan, so that all interior points move to a single triangle and are in canonical form.

To construct the fan chords, we always divide the problem into two roughly equal parts. We begin by constructing two chords as follows: perform a radial sweep from v_1 to successive hull vertices v_i $\{2 \leq i \leq n-1\}$, always keeping fewer than $\frac{n}{3}$ vertices in the swept region. Let v_j be the last hull vertex for which this holds. Construct chords v_1v_j and v_1v_{j+1}. The unswept region not including triangle $v_1v_jv_{j+1}$ contains fewer than $\frac{2n}{3}$ vertices. The swept region contains fewer than $\frac{n}{3}$ vertices. Triangle $v_1v_jv_{j+1}$ may contain an arbitrary number of vertices, but this is not a sub-problem (we will not look at this region again during the construction of the fan). Now we can continue a new sweep on each side of $v_1v_jv_{j+1}$. Construction of the two chords could take $O(n)$ edge flips and point moves, as described in Lemma 6. However the even split of the subproblems ensures that the total number of operations is $O(n \log n)$.

Each fan triangle $v_1 v_i v_{i+1}$, containing k_i interior points, can be reconfigured into a backbone structure with $O(k_i \log k_i)$ operations, by Theorem 1. Thus the total number of edge flips and point moves used to reconfigure all triangles of the fan into backbone structures is $O(n \log n)$.

Now we are left only with the task of merging the fan triangles so that only one of them will contain all interior points. We can add k_i interior points of a canonical triangulation to an adjacent canonical triangulation using $O(k_i)$ edge flips and point moves. The k_i points are processed in descending order and are always added to the top of the adjacent triangulation, as shown in Figure 7.

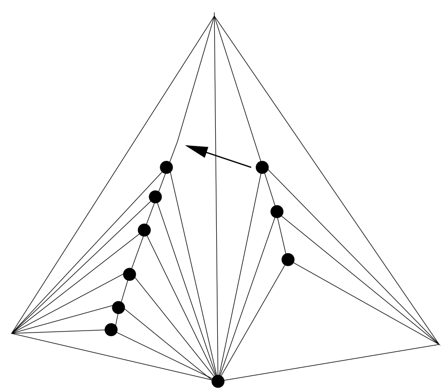

Fig. 7. Merging two adjacent fan triangles.

Thus we obtain one triangle in canonical form next to an empty triangle. It is just as easy to merge two canonical triangles separated by an empty triangle. If we encounter two or more adjacent empty fan triangles, we may use Lemma 1 to reconfigure them so that they will not affect the fan-merging process (see Figure 8). By the above arguments, once we select the triangle that is to finally contain all of the interior points (the median triangle is a good choice), we can iteratively merge its neighboring triangles onto it using a total of $O(n)$ edge flips and point moves.

Finally we are left with a single triangle containing all interior points in canonical form. On either side, we may have an arbitrary triangulation (resulting from handling multiple adjacent empty fan triangles), but the vertices will be in convex position. By Lemma 1 they may be moved to our desired configuration using $O(n)$ edge flips.

We must still show that this primary canonical configuration can be moved to specific coordinates. This can be done with $O(n)$ point moves , though space restrictions prevent us from going into any detail. $\square$

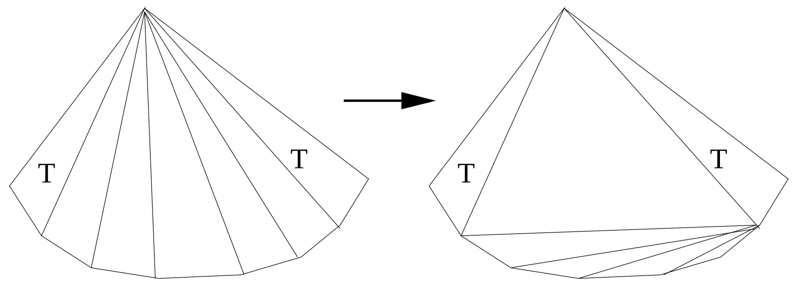

Fig. 8. Handling multiple adjacent empty fan triangles. Triangles marked (T) contain triangulations.

2.3 Remarks

If two triangulations have the same point set, the problem is no easier than the general problem. Suppose that there exists an algorithm that can transform a triangulation T_1 on a given n-point set to a triangulation T_2 on the same point set using $F_n = o(n \log n)$ edge flips and point moves. Then this algorithm can be used to transform a triangulation on one point set to any other triangulation on a possibly different point set with $F_n + O(n)$ edge flips and point moves. This argument is summarized in Figure 9. Let Figure 9a be the input triangulation. With F_n flips and moves, move to the triangulation in Figure 9b where every vertex is adjacent to the lower left vertex v_ℓ of the outer face.

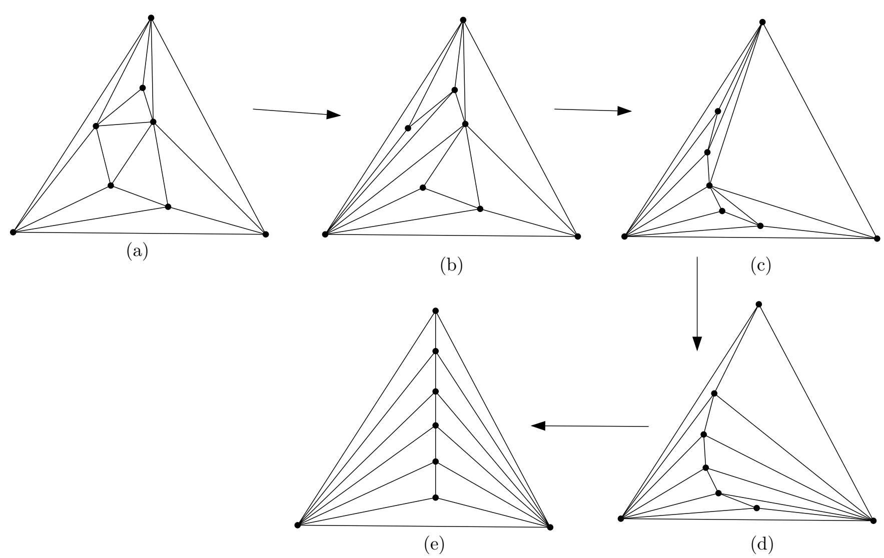

Fig. 9. Problem on fixed point set is not easier.

Now consider the triangulated polygon, P, that consists of edges not adjacent to v_ℓ. Notice that if we perform a radial sweep from v_ℓ, the boundary of P is monotonic. At least two of the triangles in P are disjoint ears, which means there must exist an ear tip that is an interior vertex and is also joined to v_ℓ by an edge in the original triangulation. We may move this point directly towards v_ℓ and cut the ear from P. This still leaves a monotone polygon P'. By continuously locating such ears, and moving them to a predefined convex position, we can obtain the configuration illustrated in Figure 9c. The monotonicity of P (and its descendants) and the convexity of the final configuration of interior points guarantee that no edge crossings will occur. This process requires a linear number of point moves.

Next, by Lemma 1, we can use $O(n)$ edge flips to obtain the triangulation where the lower right vertex of the outer face is adjacent to every vertex, as illustrated in Figure 9d. From here, it is trivial to move to the canonical configuration.

We conclude with the following:

Theorem 3. *If an algorithm exists that can reconfigure between any two geometric triangulations of the same point set with $o(n \log n)$ edge flips and point moves, then we can also transform any geometric triangulation on n points to any other geometric triangulation on n different points with $o(n \log n)$ flips and moves.*

It is tempting to try to find a fast algorithm that will construct a monotone path, as illustrated in the transition from Figure 9a to Figure 9b. Consider the polygon that is the union of all triangles incident to the lower left vertex of Figure 9b. By continuously cutting ears of this polygon, we may get to a triangulation that is *similar* to that of Figure 9a, using $O(n)$ edge flips. The similarity is that all neighbors of the lower left vertex will be in convex position. However, we have little control over the resulting positions of the remaining edges if we use only $O(n)$ operations. It is possible to create triangulations for which the reversal of this ear-cutting technique is not possible. In fact, Figure 9c serves as an example, if we add a few more vertices inside the large triangle. In this figure none of the edges directly visible from the lower left vertex can be flipped, so there is no obvious way to achieve a monotone path with fewer than $O(n \log n)$ operations.

We finally consider the following more powerful point move as an alternative to the point move studied so far. In this more powerful point move, we can delete an interior vertex of degree three (and all its incident edges), and create a new vertex of degree three inside another triangle of the triangulation. With this type of move we can reconfigure triangulations using $O(n)$ operations. We simply select a triangle incident to a hull edge and create a backbone inside. This is done by continuously selecting a vertex of constant degree from outside the triangle, reducing its degree to three, and moving it to the lower end of the backbone.

References

1. Abellanas, M., Bose, P., Garcia, A., Hurtado, F., Ramos, P., Rivera-Campo, E., Tejel, J.: On local transformations in plane geometric graphs embedded on small grids. In Proceedings of the International Workshop on Computational Geometry and Applications (CGA) **2** (2004) 22–31
2. Bose, P., Czyzowicz, J., Gao, Z., Morin, P., Wood, D.: Parallel diagonal flips in plane triangulations. Tech. Rep. TR-2003-05, School of Computer Science, Carleton University, Ottawa, Canada (2003)
3. Brunet, R., Nakamoto, A., Negami, S.: Diagonal flips of triangulations on closed surfaces preserving specified properties. J. Combin. Theory Ser. B **68(2)** (1996) 295–309
4. Cortés, C., Grima, C., Marquez, A., Nakamoto, A.: Diagonal flips in outer-triangulations on closed surfaces. Discrete Math. **254(1-3)** (2002) 63–74
5. Cortés, C., Nakamoto, A.: Diagonal flips in outer-torus triangulations. Discrete Math. **216(1-3)** (2000) 71–83
6. Galtier, J., Hurtado, F., Noy, M., Pérennes, S., Urrutia, J.: Simultaneous edge flipping in triangulations. Internat. J. Comput. Geom. Appl. **13(2)** (2003) 113–133
7. Gao, Z., Urrutia, J., Wang, J.: Diagonal flips in labelled planar triangulations. Graphs Combin. **17(4)** (2001) 647–657
8. Hurtado F., Noy, M.: Graph of triangulations of a convex polygon and tree of triangulations. Comput. Geom. **13(3)** (1999) 179–188
9. Hurtado, F., Noy, M., Urrutia, J. Flipping edges in triangulations. Discrete Comput. Geom. **22(3)** (1999) 333–346
10. Komuro, H.: The diagonal flips of triangulations on the sphere. Yokohama Math. J. **44(2)** (1997) 115–122
11. Komuro, H., Nakamoto, A., Negami, S.: Diagonal flips in triangulations on closed surfaces with minimum degree at least 4. J. Combin. Theory Ser. B **76(1)** (1999) 68–92
12. Lawson, C.: Software for c_1 surface interpolation. In J. Rice, ed., Mathematical Software III, pp. 161–194, Academic Press, New York (1977)
13. Meisters, G.: Polygons have ears. American Mathematical Monthly **82** (1975) 648–651
14. Nakamoto, A., Negami, S.: Diagonal flips in graphs on closed surfaces with specified face size distributions. Yokohama Math. J. **49(2)** (2002) 171–180
15. Negami, S.: Diagonal flips in triangulations of surfaces. Discrete Math.**135(1-3)** (1994) 225–232
16. Negami, S.: Diagonal flips in triangulations on closed surfaces, estimating upper bounds. Yokohama Math. J. **45(2)** (1998) 113–124
17. Negami, S.: Diagonal flips of triangulations on surfaces, a survey. Yokohama Math. J. **47** (1999) 1–40
18. Negami, S., Nakamoto, A.: Diagonal transformations of graphs on closed surfaces. Sci. Rep. Yokohama Nat. Univ. Sect. I Math. Phys. Chem. **40** (1993) 71–97
19. Wagner, K.: Bemerkung zum Vierfarbenproblem. Jber. Deutsch. Math.-Verein. **46** (1936) 26–32
20. Watanabe, T., Negami, S.: Diagonal flips in pseudo-triangulations on closed surfaces without loops. Yokohama Math. J. **47** (1999) 213–223