Some thoughts about decomposition of a polygon into two congruent pieces

UNFINISHED DRAFT

G. Rote, TU Graz, Austria
rote@opt.math.tu-graz.ac.at

June 24, 1997

This draft was prompted by the paper of Kimmo Eriksson, Splitting a polygon into two congruent polygons, *American Mathematical Monthly* **103** (Mai 1996), 393–400. In this paper, while all the lemmas and statements which are proved may be correct, there are a lot of errors in the side remarks.

1. Kimmo claims that his algorithm works also for polygons with holes. At the end I present a counter-example.

2. Kimmo claims that his algorithm has a complexity of $O(n^3)$ steps. (a) This refers to "pencil-and-paper" complexity, where, in one step, one can draw a line segment and check whether it intersects some part of the boundary. On a computer, this checking would have to be supported by some geometric "ray-shooting" data structure and would incur at least an $O(\log n)$ overhead. (b) More seriously, in the case of mirror symmetry, the number of steps of the algorithm cannot be bounded in terms of n alone. Consider the 9-gon $(0,1)$, $(1,0)$, $(2,0)$, $(2,1)$, $(3,2)$, $(3,2k)$, $(1,2k)$, $(1,2k-1)$, $(0,2k)$, $(0,1)$, for some integer $k > 2$. It can only be decomposed into two $(4k+1)$-gons.

3. Less gravely, I found it quite annoying that some of the figures were drawn with so little precision that one somehow has to guess the intended proportions.

This draft presents some of my own thoughts exploring alternative algorithmic ideas that may eventually lead to an improved solution. As a special feature, this article contains a wrong theorem.

We assume that P, as well as the two parts P_1 and P_2 into which it is to be decomposed, is *regular* in the sense that each set is the closure of its interior. P_1 and P_2 must have no common interior points.

Sometimes we want P_1 and P_2 to have connected interiors. Sometimes we may require that they be simply connected (or simple polygons), i.e., the boundary should be a closed curve without self-intersections.

1 Some results

We concentrate on the case of mirror congruence.

In the plane, a mirror congruence is effected by a reflection and translation (Gleitspiegelung), pure reflection being a special case. That is, a reflection at some line g, followed by a translation parallel to g.

1.1 Mirror congruence — possible transformations

For simplicity, we exclude the case that P is mirror-symmetric. This case is trivial. So from now on we assume that the transformation is not a pure reflection.

First we will discuss how the set of possible transformations can be restricted; then we will see how a particular transformation can be tested, whether it maps a part P_1 of P to its complement.

Let s_P denote the center of gravity of the region P.

Lemma 1 *The reflection line g goes through s_P.*

Proof. s_{P_1} and s_{P_2} lie on different side of g and have the same distance from g. Therefore, $s_P = (s_{P_1} + s_{P_2})/2$ lies on g.

1.1.1 Mirror congruence — possible directions

Lemma 2 *Let g_1 and g_2 be the two supporting lines of P parallel to g. Then g_1 and g_2 lie symmetric with respect to g, see Figure 1*

Proof. Let X_1 be a point of P that lies on g_1 and suppose, w.l.o.g., that $X_1 \in P_1$. By the congruence transformation, this point is mapped to a point X_2 of P_2 that has the same distance from g but lies on the other side. This means that the distance of g_2 from g is at least as big as the distance of g_1 from g. by an analogous argument we can get the reverse inequality.

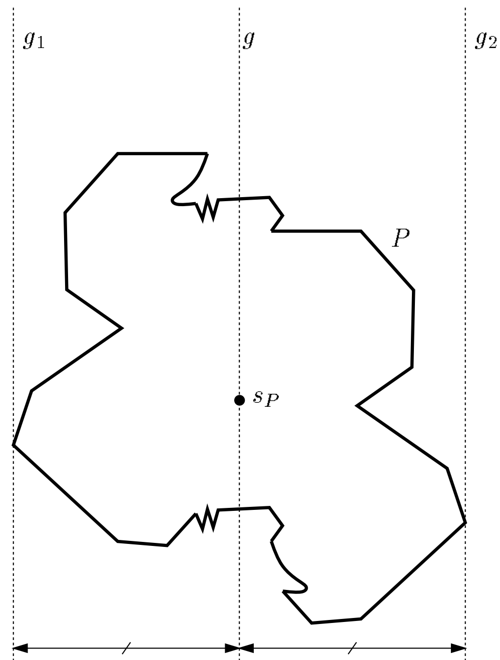

Figure 1: Symmetry about the line g.

Now let t_X denote the reflection about a point X (rotation by π around X).

Lemma 3 *Let $P' := t_{s_P}(P)$. If there is a mirror congruence with reflecting line g, then P and P' have a common supporting line parallel to g.*

Proof. t_{s_P} maps the line g_2 of the previous lemma to g_1.

By this lemma, we can just look at the common tangents of P and P', as shown in Figure 2, and most of the time, we will get a discrete set of possible directions for g. However, it may happen that P and P' have common extreme vertices, and in these cases, this will correspond to an interval of possible directions for g.

Lemma 4 *Let g_1 and g_2 be the two supporting lines of P parallel to g, and suppose that each of them intersects P in only one point X_1 and X_2, respectively. Then there are neighborhoods U_1 and U_2 of X_1 and X_2, respectively, such that $P \cap U_1 \subseteq P_1$ and $P \cap U_2 \subseteq P_2$.*

Proof. Suppose, w.l.o.g., that $X_1 \in P_1$. Hence we must have $X_2 \in P_2$. (We assume that P_1 and P_2 are closed sets.) If X_2 would also belong to P_1, then the image of X_2 would be another point on $g_1 \cap P$, contradicting the assumption. (Recall that we have excluded the case of a pure reflection.) (That case can be handled as well, if one assumes that P_1 and P_2 must be simple polygons.)

This lemma helps us to deal with the degenerate case discussed above, where P has two extreme vertices that lie symmetric with respect to s_P: we know that the whole neighborhood of these two vertices must appear mirror-congruent. So, if P and P' have a common extreme vertex V, let h be the direction from V clockwise on P, and let h' be the direction from V counterclockwise on P'. (In case of polygons. Otherwise we must take tangent directions or

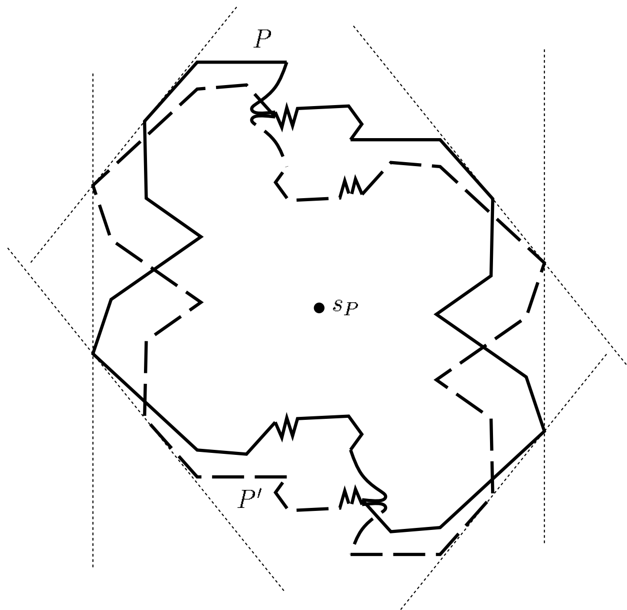

Figure 2: The possible supporting directions for g.

local "outer tangents" or something like that.) Then g must be perpendicular to the angle bisector between directions h and h'. (In addition, we have the two "extreme" possibilities for g, which correspond to the directions where the supporting line at V touches another point of P or P'.)

Conclusion:

Lemma 5 *For a polygon P with n vertices, there is a discrete set of $O(n)$ candidate lines for the mirror line g. This set can be computed in $O(n)$ time.*

1.1.2 Mirror congruence — possible translations

Suppose now that we have fixed a mirror line g, and thereby the translation direction. We still have to find the amount of translation. If g_1 and g_2 touch P at only one point each, we are lucky: we can apply Lemma 4 and this gives us a unique possible transformation.

In the general case, we proceed as follers. Draw g vertically, with g_1 on the left and g_2 on the right. Let X_1 be the highest point on $g_1 \cap P$, and let X_2 be the highest point on $g_2 \cap P$.

Case 1. $X_1 \in P_1$ and $X_2 \in P_2$. This case is easy, and it gives a unique transformation.

Case 2. $X_1 \in P_1$ and $X_2 \in P_1$. We write the reflection at g as $X \mapsto \overline{X}$, and the translation parallel to g by an amount α as $X \mapsto X + \alpha e_g$. e_g is a unit vector parallel to g. So the general transformation that we are looking for has the form

$$X \mapsto \overline{X} + \alpha e_g.$$

We have $X_2' := \overline{X_1} + \alpha e_g \in g_2 \cap P_2$ and $X_1' := \overline{X_2} + \alpha e_g \in g_1 \cap P_2$. X_1' lies below X_1 on g_1, and X_2' lies below X_2 on g_2, see Figure 3. (One pair of points may coincide, but not both pairs: otherwise we would have a pure reflection ($\alpha = 0$.)

Since P_2 is a simple polygon, [**This argument works only for decomposition into simple polygons.**] everything on g_1 which lies below X_1' must belong to P_2. Since X_1' is the highest point of P_2 on g_1, by the definition of X_2, everything on g_1 which lies above X_1' must belong to P_1. The point X_1' may belong to both sets. Similarly, X_2' acts as a dividing point between P_1 and P_2 on g_2.

Suppose that we start with X_1' and X_2' at the highest possible position (where one of them coincides with X_1 or X_2, respectively), and we start translating them downwards by varying α. The part on g_1 above X_1', which belongs to P_1, will increase, whereas its congruent image, the part on g_2 below X_2', will diminish. We can, for example measure the "length" of these

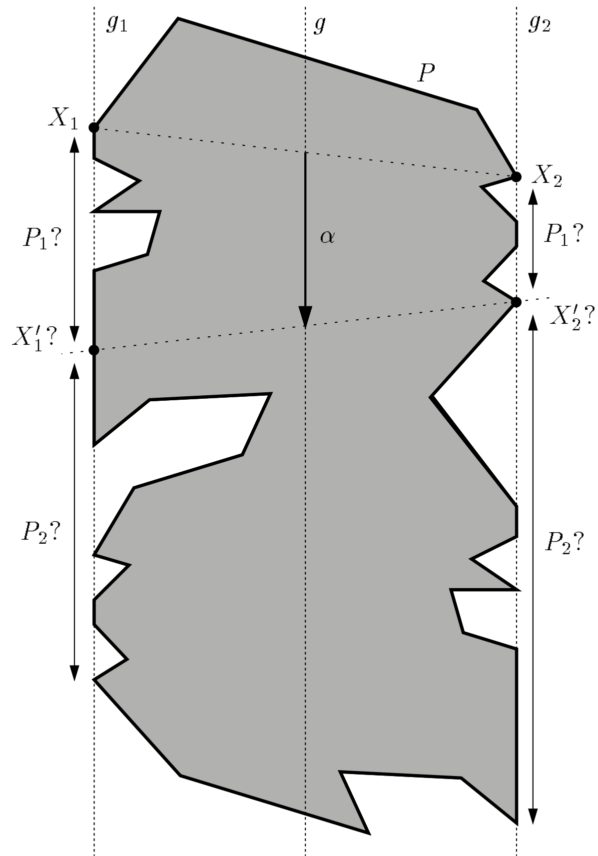

Figure 3: Trying to vary X_1' and X_2'.

parts $P_1 \cap g_1$ and $P_2 \cap g_2$ by their diameter. So there will be at most one value of α where the two lengths coincide. (The length of $P_1 \cap g_1$ does not uniquely depend on α, because X_1' may or may not belong to P_1, but, since the length of $P_2 \cap g_2$ is *strictly* increasing, this does not matter.)

Summary: For a given line g, we have at most two possible translations.

1.2 Testing a transformation in general

Suppose we are given a transformation f, and we want to test if P can be decomposed into two closed sets P_1 and P_2 such that $P_2 = f(P_1)$, $P_1 \cup P_2 = P$, and P_1 and P_2 have no common interior points.

Let f be a transformation. The *orbit* of a point X is the sequence

$$\dots f^{-2}(X), f^{-1}(X), X, f(X), f^2(X), \dots$$

An orbit is either a bi-infinite sequence or a cyclic sequence. By the *orbit inside* P of a point $X \in P$, we mean that part of the sequence that extends in each direction as far as the points belong to P.

Lemma 6 *Let $B = (x_1, x_2, \dots, x_m)$ be an orbit inside P. If no point of B belongs to $P_1 \cap P_2$, then the orbit has even length.*

Proof. (Infinite orbits can be excluded.) Suppose first that the orbit is not cyclic. Since $f^{-1}(x_1) \notin P$, we cannot have $x_1 \in P_2$. Therefore $x_1 \in P_1$ and $x_2 = f(x_1) \in P_2$. By assumption, this implies $x_2 \notin P_1$, and therefore $x_3 \notin P_2$; thus, $x_3 \in P_1$. By continuing this argument we conclude that every x_i with odd i is in P_1 and every x_i with even i is in P_2. If n were odd, we would get $x_n \in P_1$ but $f(x_n) \notin P$, a contradiction. When B is a cycle, we can start with some arbitrary point in it and derive a contradiction similarly.

Another way to look at this is to say that we form the undirected graph from the points of P, by joining every point X to $f(X)$. This graph must then have a complete matching (apart from the exceptions with the boundary points in $P_1 \cap P_2$.

It follows that, if there are no cyclic orbits, for example, if it is a mirror-congruence by not a pure reflection, the decomposition into P_1 and P_2 is unique.

Using the regularity assumption on the sets P_1 and P_2, one can replace the condition of the lemma by the condition that no point of the orbit belongs to the boundary of P.

1.3 Testing a transformation for mirror congruence

In this case there can be no cycles in orbits. (Pure reflection is excluded.)

1.3.1 First try

$P_{\text{init}} := P - f(P)$ is the set of starting points of orbits. These points must belong to P_1. P_{init} had better be contained in $f^{-1}(P)$; otherwise we are doomed.

$P_{\text{fin}} := P \cap f^{-1}(P) - f^{-2}(P)$ is the set of points of P which are mapped to endpoints of orbits.

$P_{\text{direct}} := P_{\text{init}} \cap P_{\text{fin}}$ are the starting points of orbits of length 2.

All these constructions can be carried out geometrically by overlaying polygons and forming Boolean operations on them. So the resulting sets are polygons.

$\tilde{P}_{\text{init}} := P_{\text{init}} - P_{\text{direct}}$ are the starting points of orbits of length 4, 6, ..., and $\tilde{P}_{\text{fin}} := P_{\text{fin}} - P_{\text{direct}}$ are their endpoints. The set $\tilde{P}_{\text{init}}$ must be connected to the set $\tilde{P}_{\text{fin}}$ via the points

$$P_{\text{rest}} := P \cap f^{-1}(P) - (P_{\text{init}} \cup P_{\text{fin}}).$$

This means:

$$x_0 \xrightarrow{f} x_1 \xrightarrow{f} x_2 \xrightarrow{f} \cdots \xrightarrow{f} x_m \xrightarrow{f} x_{m+1},$$

with $x_0 \in \tilde{P}_{\text{init}}$, $x_1, \ldots, x_{m-1} \in P_{\text{rest}}$, and $x_m \in \tilde{P}_{\text{fin}}$, $m \geq 2$ and even.

Now we reduce it to the case of pure translation by taking two steps of f at a time. Set $\hat{P}_{\text{init}} := \tilde{P}_{\text{init}} \cup f(\tilde{P}_{\text{init}})$ (first two elements of each orbit of length more than three) and $\hat{P}_{\text{fin}} := \tilde{P}_{\text{fin}} \cup f(\tilde{P}_{\text{fin}})$ (last two elements of such orbits). These two sets must be connected by repeated applications of the function f^2 via the set P_{rest}. This means, for every $x_0 \in \hat{P}_{\text{init}}$ we must have

$$x_0 \xrightarrow{f^2} x_1 \xrightarrow{f^2} x_2 \xrightarrow{f^2} \cdots \xrightarrow{f^2} x_m,$$

with $m \geq 1$, $x_1, \ldots, x_{m-1} \in P_{\text{rest}}$, and $x_m \in \hat{P}_{\text{fin}}$. the same must be true for every $x(= x_m) \in \hat{P}_{\text{fin}}$. (If one checks beforehand that P_{init} and P_{fin} have the same area, only one of these conditions is sufficient.)

Do we have to say that every $x \in P_{\text{rest}}$ must belong to such a chain? No. By definition of P_{rest}, if $x \in P_{\text{rest}}$ then $f^{-2}(x) \in P$ and $f^2(x) \in P$. So the point $f^{-2}(x)$ is either again in P_{rest}, or it is in $\hat{P}_{\text{fin}}$.

Lemma 7 *If $x \in P_{\text{rest}}$ then $f^2(x) \in P_{\text{rest}} \cup \hat{P}_{\text{fin}}$ and $f^{-2}(x) \in P_{\text{rest}} \cup \hat{P}_{\text{init}}$.*

Proof. By definition we have

$$\begin{aligned}
x \in P_{\text{init}} &\iff f^{-1}(x) \notin P \text{ and } x \in P \text{ and } [\ f(x) \in P\ ] \\
x \in P_{\text{fin}} &\iff x \in P \text{ and } f(x) \in P \text{ and } f^2(x) \notin P \\
x \in P_{\text{direct}} &\iff f^{-1}(x) \notin P \text{ and } x \in P \text{ and } f(x) \in P \text{ and } f^2(x) \notin P \\
x \in \tilde{P}_{\text{init}} &\iff f^{-1}(x) \notin P \text{ and } x \in P \text{ and } f(x) \in P \text{ and } f^2(x) \in P \\
x \in \tilde{P}_{\text{fin}} &\iff f^{-1}(x) \in P \text{ and } x \in P \text{ and } f(x) \in P \text{ and } f^2(x) \notin P \\
x \in P_{\text{rest}} &\iff f^{-1}(x) \in P \text{ and } x \in P \text{ and } f(x) \in P \text{ and } f^2(x) \in P \\
x \in \hat{P}_{\text{init}} &\iff f^{-1}(x) \notin P \text{ and } x \in P \text{ and } f(x) \in P \text{ and } f^2(x) \in P \\
&\qquad \text{or } \left(f^{-2}(x) \notin P \text{ and } f^{-1}(x) \in P \text{ and } x \in P \text{ and } f(x) \in P\right) \\
x \in \hat{P}_{\text{fin}} &\iff f^{-1}(x) \in P \text{ and } x \in P \text{ and } f(x) \in P \text{ and } f^2(x) \notin P \\
&\qquad \text{or } \left(f^{-2}(x) \in P \text{ and } f^{-1}(x) \in P \text{ and } x \in P \text{ and } f(x) \notin P\right)
\end{aligned}$$

From this the lemma follows.

So every $x \in P_{\text{rest}}$ is automatically part of an f^2-orbit. We only have to check that an f^2-orbit starts in every point of $\hat{P}_{\text{init}}$ and an orbit ends in every point of $\hat{P}_{\text{fin}}$.

Theorem 1 *A decomposition into two congruent parts P_1 and $P_2 = f(P_1)$ is possible if and only if $f(P_{\text{init}}) \subseteq P$,*

$$f^2(\hat{P}_{\text{init}}) \subseteq P_{\text{rest}} \cup \hat{P}_{\text{fin}},$$

and

$$f^{-2}(\hat{P}_{\text{fin}}) \subseteq P_{\text{rest}} \cup \hat{P}_{\text{init}}.$$

Proof. These conditions are clearly necessary. By the lemma, they are sufficient to get the orbit going from each point in $\hat{P}_{\text{init}}$ and $\hat{P}_{\text{fin}}$.

This theorem is very surprising. **In fact, the theorem is wrong!** [Hint: FIND THE ERROR!] It should be possible to apply it to the case of rotations as well.

All sets involved can be constructed by overlaying a constant number of different congruent copies of P and performing Boolean operations. If P is an n-gon, the resulting sets might be regions with $O(n^2)$ vertices and edges, and they can be constructed within $O(n^2)$ time. These bounds hold in general. Maybe, by special considerations, one can show that in the case that interests us, i.e., when the condition of the theorem is fulfilled, this high complexity cannot arise. Some approach by a sweepline-algorithm seems promising.

Above we have reduced the problem to pure translation. This is perhaps more intuitive to deal with, but it might be possible to bypass this step. But I think no. How are we to treat the parity constraint for the length of the orbits?

1.3.2 Second try

Let us look at the orbit of a point X, and write a x for each point lying in P and a ∘ for each point outside P. We get an infinite x-∘-sequence like

∘ ∘ ∘ ∘|x x|x x|∘|x x|∘ ∘ ∘ ∘|x x|x x|x x|x x|∘ ∘ ∘ ∘,

which we call the *characteristic sequence* of X. Such a sequence is called a *legal* sequence if each consecutive block of x's has even length. In this case, the x's can naturally be grouped into pairs, as indicated in the above sequence. We call the left x of each pair the *odd* x and the other x the *even* x.

By Lemma 6, legality is a necessary and sufficient condition for decomposability, if this holds for every point X for which no point of the orbit lies on the boundary of P.

In principle, we can construct the sequence for each orbit and check its legality. It is sufficient to take one representative for each orbit, for example we can restrict X to some fundamental domain D of the transformation f, like a half-infinite half-strip of width 2α bounded by a segment of g of length 2α and two rays perpendicular to g. We can take all iterated images $f^i(P)$ of P for $i \in \mathbb{Z}$ and overlay them over D. This will split D into finitely many regions. If we vary the point X inside one such region, the associated x-∘-sequence remains constant because no point of the orbit crosses the boundary of P. Therefore it would be sufficient to select one representative point from each region and check legality. Unfortunately, the number of these regions cannot be bounded in terms of n, since an edge of P can intersect an arbitrary number of copies of another edge inside D.

Let us see how the characteristic sequence changes as X moves a little and some points of the orbit cross the boundary of P.

∘ ∘ ∘ ∘|x x|x x|∘|x x|∘ ∘ ∘ ∘|x x|x x|x x|x x|∘ ∘ ∘ ∘ ∘|x x|∘ ∘ ∘ ∘ ∘

∘ ∘|x x|x x|x x|∘|x x|x x|∘ ∘ ∘ ∘ ∘|x x|x x|x x|∘ ∘ ∘|x x|x x|∘ ∘ ∘ ∘

Each change from x to ∘ or from ∘ to x corresponds to some point $f^i(X)$ in the orbit which crosses an edge of P, and vice versa. We refer to these positions in the sequence as the *swap positions*. We refer to these edges as *the edges at X*.

Since the total number of x's remains even, there must be an even number of such crossings, and we match them pairwise. We see that some x's change from even to odd or from odd to even. A block of contiguous x's that change in this way is bounded by a swap position on each side. Such a block can either grow or shrink in both directions, and then the two swap positions change in the same sense, or it can sort-of "shift" to the left or right, and then the two swap positions change in the opposite senses. In any case, the two swap positions cannot be adjacent. We group the two swap positions into a pair, which we call a *skew pair*. The skew pairs of swap positions are indicated in the example by arrows. On the other hand, an odd-even pair of x's on one side can be exchanged for two adjacent ∘'s on the other side; in this case we group the two swap positions into a *straight pair*.

Lemma 8 *Each swap position belongs to a unique pair. The straight and skew pairs exhaust all swap positions.*

The characteristic vector changes only if the point X crosses an edge in the arrangement of all overlaid polygons inside D. In this overlaid arrangement, (part of) some edges from different copies of P will coincide and form a single edge of the arrangement. It is possible that, as X moves, it crosses several different edges of the arrangement (such as when X moves across a vertex of the arrangement). However, we will only be interested in the changes as X crosses single edges: This is sufficient because X can go from any region to any other region by crossing only single edges at a time.

We have seen above that swap positions come at least in pairs. Therefore, when X crosses an edge, there are (parts of) other edges that the orbit must necessarily cross. Above we have matched swap positions into skew and straight pairs. We can extend this matching to the corresponding edges. (It may be necessary to partition edges of P into parts that correspond to the single edges of the arrangement and that can then be matched.)

Definition. A sub-edge is a maximal segment e in the boundary of P such that every segment in the orbit of e is either completely contained in the boundary of P or has no subsegment in common with the boundary of P.

Sub-edges are the units which the algorithm treats. (It would be possible to combine adjacent sub-edges into one path, if they correspond to the same pieces in the orbit. (by extending the above definition to polygonal paths instead of segments.)

The edges of the arrangement can be found in a preprocessing step. For each edge of P, we construct the line through it and find all other edges that are contained in an iterated image or pre-image of this line. At all endpoints of these edges (more precisely, at their appropriate images or pre-images), each of these edges must be subdivided. This gives the subdivision of each edge into sub-edges, and also, for each edge of the arrangement, the class of sub-edges of P that belong to it. (Here, a segment is considered as a single edge of the arrangement even if it is crossed by other edges of the arrangement. It is only required that the set of sub-edges of P corresponding to it is fixed.

Lemma 9 *If u, v is a skew pair of sub-edges, with $v = f^i(u)$, $(i \geq 2)$, then none of the images $f(u)$, $f^2(u)$, …, $f^{i-1}(u)$ may cross any other edge of P.*

Proof. Suppose $f^j(u)$ if the first edge in the above sequence that crosses some other edge w. Suppose X lies close to u. As $f^j(X)$ crosses w, the symbol in the j-th position changes between x and ∘. As X moves from one side of u to the other, the $(j-1)$-st position changes between an even x and an odd x (or between an odd x and a ∘, in case $j = 1$). In one of these positions, the resulting sequence is therefore illegal when the symbol in the j-th position is ∘.

So we proceed as follows. For each sub-edge of P, we place one representative point X somewhere in the interior of that edge arbitrarily, and we see what happens as X crosses this edge, by constructing the characteristic sequences on both sides.

Theorem 2 *If the characteristic sequence is legal in the vicinity of the (arbitrarily chosen) representative point of each sub-edge, and if the condition of Lemma 9 is fulfilled for every skew pair of sub-edges, then there is a partition.*

Proof. We show that we can cross an arbitrary number of sub-edges and always maintain a legal sequence. For each sub-edge e, there is one point X where we have already tested that it is possible to cross the edge there. Suppose we cross e at a different point X'. We will perform precisely the same swaps between x and ∘ as at X; but it is conceivable that now they lead from a legal sequence to an illegal sequence. In particular this happens when x's change between even and odd status. Note, however, that this happens only in case of a skew swap.

Suppose that, as we go along e from X to X', a skew swap happens as we cross another edge w, between positions i and j in the sequence. By Lemma 9 applied to w, none of the sub-edges of P that are involved in the swaps for e crosses between positions i and j, i.e., the skew swap between positions i and j does not affect correctness of the swaps of e since these swaps do not involve any of the positions from i to j.

Now suppose that, as we go along e from X to X', a straight swap happens as we cross another edge. Then we know that even and odd x's remain so.

Consider a skew swap happening at X, between positions i and j in the sequence. Again by Lemma 9, but this time applied to e, nothing happens between i and j as X moves to X'. Therefore the swap is still legal. Any straight swap happening at X must remain legal, since in the vicinity of the swap, x's cannot have changed between even and odd.

[After reading your thesis, I trust that you would be able to express this "independence" of swaps more elegantly, based on general principles.

Example: The following must be excluded:

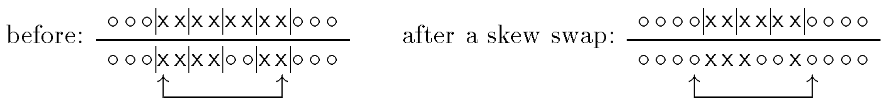

The right swap is now illegal.

It is also conceivable that another edge does not cross e, but it approaches e and makes a turn precisely when it hits e, possibly splitting into other sub-edges that continue partially on the same side of e as before and partially on the other side. The above argument can be extended to this case. The only dangerous case is when another edge hits e and (partly) joins e in one direction. But in this case the sub-edge e would have been split into smaller sub-edges at this point.

Let us discuss how this check can be carried out algorithmically. The check of the characteristic sequence at particular point X can be done by a sweep along a line through X parallel to g simultaneously with a reflected line on the other side of g. We maintain an interval of length 2α as a representation of D. We sweep down these lines, maintaining the status of the part of the sequence thus far constructed, for each point of the interval. We have to do some work for each edge of P that we cross on either of the two lines. Quadratic time suffices clearly, and by using the fact that one can stop if an illegal sequence is detected, it can be improved to $O(n \log n)$ time. The total number of sub-edges (in the arrangement) (this is the number of representative points X) is $O(n)$: Each endpoint of a sub-edge comes from some endpoint of an original edge. As we process each point X, we can also identify the skew and straight pairs of edges.

The test of Lemma 9 can be done in $O(n)$ for each skew pair u, v of sub-edges: We simply check every other edge. Since they are line segments, it can be determined in constant time, which of a series of periodic translated copies of u they intersect.

How many skew pairs can there be? Certainly at most $O(n^2)$, but probably $O(n)$?

This gives an overall bound of $O(n^3)$ or maybe $O(n^2)$ for testing a *single* transformation f.

There is ample space for improvement: (a) algorithmic ideas, i.e., one plane-sweep with a pair of parallel line (symmetric about g) that does it all; (b) structural insights: How can the approach be simplified?

The resulting partition polygons P_1, P_2 are not necessarily connected. They can be arbitrary regions.

It should be possible to apply these ideas to the case of rotations as well.

1.4 A property of the boundary of a simple polygon in case of mirror congruence

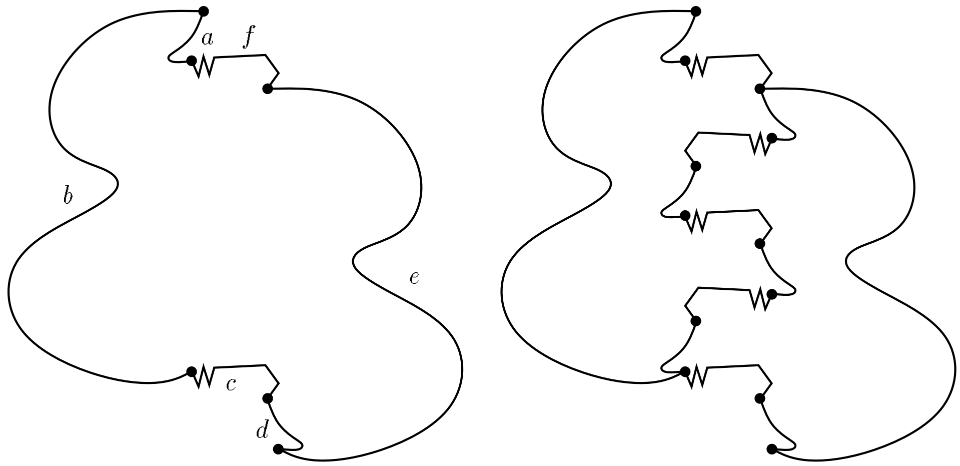

Figure 4: The structure of the boundary in case of mirror symmetry.

Theorem 3 *Suppose P is a* simple polygon *(or rather: simply connected region) that can be partitioned into P_1 and P_2, where P_1 is congruent to P_2 by a mirror congruence with reflecting line g.*

Then the boundary of P can be partitioned into six pieces a, b, c, d, e, f (in this order) with the following properties, see Figure 4: (Any of the pieces can be missing.)

1. *c is a copy of f translated parallel to g.*
2. *The two endpoints of c (and of f) have the same distance from g and lie on different sides of g.*
3. *a and d are mirror-congruent by a reflection along g and a translation parallel to g.*
4. *b and e are mirror-congruent by a reflection along g and a translation parallel to g.*

This theorem has not been thoroughly proved. Something can be said about the amount of translation, too.

1.5 Proper congruence

Theorem 4 *The center of rotation C lies on the furthest-point Voronoi diagram of P.*

Proof. Let X_1 be a point of P that is furthest from C, and suppose, w.l.o.g., that $X_1 \in P_1$. Since $X_1 \neq C$, this point is mapped to a different point X_2 of P_2. So X_1 and X_2 are two furthest points from C in P, and C lies on the furthest-point Voronoi diagram of P.

Lemma 10 *Suppose that the center of rotation C lies in the interior of P. Then the angle of rotation is a fraction of π. Moreover, if we want P_1 and P_2 to have connected interiors, the angle of rotation is π, and P is symmetric about C.*

So, if P is not symmetric about a point, and if we are only interested in parts with connected interiors, We can exclude the interior of P from consideration.

Lemma 11 *If the center of rotation C lies in the exterior of P (not on the boundary), then it must lie on the (closest-point) Voronoi diagram of P.*

Proof. Similar to the proof of Theorem 4.

Lemma 12 *Suppose P_1 and P_2 have connected interiors. Consider the intersection of P_1 and P_2 with a circle k around C. These two intersections can be "separated" on k. In other words, the sets $(P_1 - P_2) \cap k$ and $(P_2 - P_1) \cap k$ are contained in two circular arcs with disjoint interior.*

This means that there are two curves that start at C and move outwards to infinity (the distance from C is weakly increasing) which together separate P_1 from P_2.

We can try to find the rotation angle by looking at the intersection of P with any circle k centered at C. The following lemma is handy in this situation.

Lemma 13 *Let P' be a disjoint union of arcs of a circle k, and suppose that P' can be decomposed into two parts P'_1 and P'_2, which are separated on k (in the sense explained above) and which are congruent by a rotation by α, where $\alpha \neq \pi$. Then this decomposition is unique.*

Proof. Let's parameterize the circle by the angle with the positive x-axis, so that it is parameterized by $[0 \,..\, 2\pi]$. W.l.o.g., we assume that that P'_1 extends over the arc from 0 to ϕ, and P'_2 extends over the arc from α to $\alpha + \phi$, where $0 < \phi \leq \alpha < \pi$. If we try to separate P into two parts in a different way, we must separate it at corresponding points $\psi \in (0 \,..\, \phi)$ and $\psi + \alpha$, because only in this way it is ensured that both parts have the same total length of intervals. But any such division does not cut P into two congruent parts, since one part contains a gap between $\alpha + \phi$ and 2π, whereas the other part contains a gap between ϕ and α in the corresponding position (or no gap at all, if $\phi = \alpha$). These two gaps are of different size, and therefore the two parts cannot be congruent.

---

2 An example where Kimmo's algorithm fails

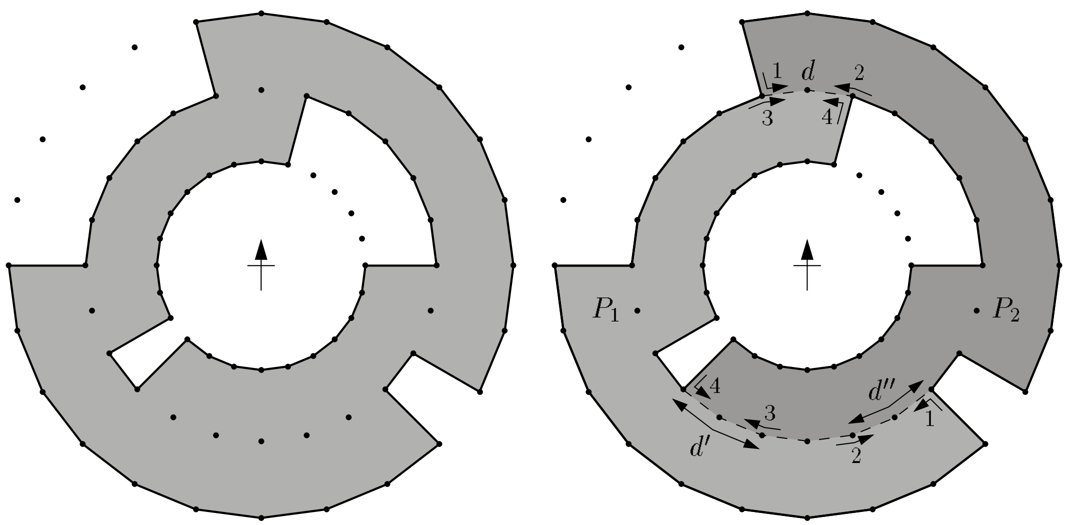

Figure 5: A counter-example to the algorithm.

This ringlike shape P can be decomposed uniquely into two congruent parts P_1 and P_2 as shown in the figure. If the algorithm is to find this division, one of the two paths that try to trace out the boundaries of P_1 and P_2 must at some time leave the boundary of P and enter the small dividing segment d in the upper part of P. There are four possibilities to do this, which marked by little arrows. Supposing that the second path at the same time correctly traces out the boundary of the other part, this second path should proceed as shown by the arrows in the lower part of the figure. In each case, it can be checked that the second path will see no reason to take this course, because it is either following the boundary of P and will happily continue to do so, or it is in the interior and will continue to follow the movements of the first path.

(One might say that the second path might follow the correct course because it hits some part of a path that was drawn before. But that part (d' or d'') could never have been drawn as long as d is not drawn.)