Compatible Geometric Matchings and Bent Edges

Elizabeth J. Kupin

1 Existence and counting problems for straight line matchings

One question to check, before beginning work on this area, is to see if there are any sets of points that are automatically bad, and cannot support two disjoint, straight line matchings. Fortunately, this is false. It leads into some interesting small results on sets of points in general position, and how many straight line matchings they can contain.

Theorem 1. *Every set of 2n points in the plane (in general position) can support a perfect matching made up of straight, non-crossing lines.*

Proof. This is actually a common exercise in an algorithms course. Start with an arbitrary matching of the points, made up of straight lines. It is likely that there are many crossings among the lines. Arbitrarily pick two lines that are crossing, and re-connect those 4 points with two line segments that do not cross each other. This may create new crossings, so continue the process while crossings are left in the graph. At every step the sum of the lengths of all the line segments decreases, therefore the process must terminate. □

From this result, we can generalize to a case of particular interest to us:

Theorem 2. *Every set of 2n points in the plane (in general position) can support two disjoint, perfect matchings made up of straight, (mutually) non-crossing lines.*

Proof. As before, start with two arbitrary straight-line matchings. This gives us a 2-regular graph made up of even cycles, and we will think of the edges in each matching as red or blue. We will eliminate the crosses one by one, and the overall sum of the lengths of the line segments used will decrease. The only thing to be shown is that at the end of each step we will still have two disjoint perfect matchings.

If we have a place where two red edges cross each other (or two blue edges), we are free to uncross them as in the proof above. Each point will still be incident to exactly one red edge and exactly one blue edge, but the sum of the lengths of the edges will have decreased. Moreover, if we have a place where a red and a blue edge cross, but they are from disjoint cycles, we can flip the colors on one of the cycles and handle the crossing as in above.

We need to be slightly more careful when we have a red edge crossing a blue edge from the same cycle. There is no way to preserve the coloring if we uncross them, and we need to show that we are able to uncross them and avoid breaking the cycle into two smaller cycles. If we remove the two edges, there should be two even paths left in the graph. Their 4 endpoints form a quadrilateral in the plane, and it must be that the endpoints of one path cannot be diagonally opposite each other: the original edges were crossing, so they must form the diagonals of the quadrilateral. Since it was one single cycle, no original edge could have connected the endpoints of a path.

We have two options, then, to pick two non-crossing edges to connect the endpoints. We could either connect the endpoints of each path to each other (thus creating two smaller, odd cycles), or we could connect the endpoint of one path to the endpoint of the other path that is adjacent to it as we move around the perimeter of the quadrilateral. We pick the second way, that guarantees that we preserve the size of the cycle. However, because we have rearranged the order of the edges, the coloring of the cycle is no longer alternating red and blue. Our last step in this case is to re-color this cycle so that the edges do alternate red and blue. The sum of all the segment lengths has decreased, and at the end of each uncrossing we are left with two perfect matchings, so the algorithm must terminate with two perfect, non-crossing straight line matchings. □

Note that at every step, the number of disjoint cycles created by the two matchings either stays constant or decreases. If we start with a single cycle, we have as a corollary then that every set of points has a straight line, non-crossing Hamiltonian cycle.

It's also interesting to note that, at least for small cases, it is not possible to have three mutually non-crossing, disjoint perfect matchings. With 4 points there are only 3 ways to match them up, and depending on their embedding into the plane these may or may not be non-crossing.

The above theorem is strictly weaker than our original question. Both result in finding two perfect, disjoint, mutually non-crossing matchings. However, in the original question we are given one perfect non-crossing matching by an adversary, and asked to produce the second. In the theorem above, we need control over the first matching as well as the second.

Another interesting, preliminary question is how many perfect, non-crossing matchings a set of points can support. The number of perfect matchings of $2n$ points in K_{2n} is $(2n-1)!!$, or (roughly) asymptotically n^n. This is a rough upper bound, as it doesn't take into account any information about the embedding. As there are no planar embeddings of K_{2n} for $2n > 4$, our number is strictly less than $(2n-1)!!$ for large values. In fact, if the points are in convex position, the number of non-crossing perfect matchings is given by the nth Catalan number, which grows asymptotically as 4^n. It turns out that this is in some sense the worst case, and that we can always guarantee at least that many non-crossing perfect matchings regardless of how the points are embedded into the plane.

Theorem 3. *Every set of $2n$ points in general position in the plane supports at least C_n perfect, non-crossing matchings, where C_n is the nth Catalan number.*

Proof. Let $P(2n)$ be the minimum number of non-crossing perfect matchings of $2n$ points, over all embeddings of the the points into the plane. We will show that $P(2n) \geq \sum_{i=0}^{n-1} P(2n-2i-2) \cdot P(2i)$. From this, and the fact that $P(0) = P(2) = 1$, we will get the desired result.

Fix a point x on the convex hull of the set of $2n$ points in the plane. We will consider only perfect, non-crossing matchings where the edge containing x splits the set of points, that is, no other edge in the perfect matching crosses over the line determined by the segment containing x. If all the points are on the convex hull of the set, then this is always true and we aren't discarding any matchings.

To be part of a valid perfect matching, the segment containing x must split the remaining $2n-2$ points into two even groups. There are exactly n ways of doing this. To see this, consider all $2n-1$ line segments that could contain x. These line segments are contained in an angle of no more than π degrees, and so have a natural order from the smaller angle to the larger. We know that the first and last ones are valid, as they separate out no other points. As we move through, every new line segment adds one point to one side, and removes it from the other. Therefore, out of the $2n-1$ possibilities, there are exactly n valid ones.

Once we have the first split, we will form a perfect matching by matching the points in each piece separately. Any matching of one side, combined with any matching of the other side, will give us a unique perfect matching. There are at least $P(|G_1|) \cdot P(|G_2|)$ ways of doing this for a given split. Adding them up over all possible splits that x can be involved in gives us the formula for the nth Catalan number. □

The reverse question is what is the largest number of non-crossing perfect matchings of $2n$ points that can be achieved, and how can we characterize the embeddings that achieve this maximum? This is related to studying the crossing number of K_{2n}, since an embedding with very few crossings is a good candidate for an embedding with many perfect, non-crossing matchings. In fact, using a simplified version of the currently optimal construction for the crossing number of K_{2n}, we can give a construction with on the order of 8^n perfect matchings.

Suppose $n/3$ is an integer, and consider the following embedding of K_{2n}: Group the points into three sets of size $\frac{2n}{3}$. Each of these will be equally spaced along the perimeter of an oval of width 1 and height ϵ for some $\epsilon > 0$, and the three ovals will be arranged on the lines with slope 0, slope $\frac{2\pi}{3}$, and slope $\frac{4\pi}{3}$. This means that $\frac{n}{3}$ points will be on each side of each oval, and so "visible" to the points in the next oval on that side (the top and bottom points are visible to both sides). If we only match up the points between ovals to "visible" points, then none of the line segments in the matching will cross any of the ovals. The remaining points in each oval will therefore be in convex position and we are free to use any matching of them (nothing is being blocked).

A lower bound is given by the following, where the fist summation counts perfect matchings where one pair of ovals has more than $\frac{n}{3}$ edges between them (counted by ℓ), and the second summation counts matchings with no more than $\frac{n}{3}$ edges between each pair of the ovals:

$$3 \sum_{\ell=\frac{n}{3}+1}^{\frac{2n}{3}} \sum_{\substack{r,b=0 \\ \text{same sign as } \ell}}^{\frac{2n}{3}-\ell} \binom{\frac{2n}{3}-\ell}{r}^2 \binom{\frac{2n}{3}-\ell}{b}^2 \left(\frac{2n}{3}-\ell\right)^2 C_{(\frac{n}{3}-\frac{\ell+r}{2})} C_{(\frac{n}{3}-\frac{\ell+b}{2})} C_{(\frac{n}{3}-\frac{r+b}{2})}$$

$$+ \sum_{\substack{r,\ell,b=0 \\ \text{all with same sign}}}^{\frac{n}{3}} \binom{\frac{n}{3}}{r}^2 \binom{\frac{n}{3}}{\ell}^2 \binom{\frac{n}{3}}{b}^2 C_{(\frac{n}{3}-\frac{\ell+r}{2})} C_{(\frac{n}{3}-\frac{\ell+b}{2})} C_{(\frac{n}{3}-\frac{r+b}{2})}$$

To show that we get the order of growth at least 8^n, we will focus just on one term of the second sum, when r, ℓ and b are as close to $\frac{n}{6}$ as possible, and all three have the same sign. This yields

$$\binom{\frac{n}{3}}{r}^2 \binom{\frac{n}{3}}{\ell}^2 \binom{\frac{n}{3}}{b}^2 C_{(\frac{n}{3}-\frac{\ell+r}{2})} C_{(\frac{n}{3}-\frac{\ell+b}{2})} C_{(\frac{n}{3}-\frac{r+b}{2})} \sim \binom{\frac{n}{3}}{\frac{n}{6}}^6 \left(C_{(\frac{n}{3}-\frac{n}{6})}\right)^3$$

$$\sim \left(\frac{2^{\frac{n}{3}}}{\left(\frac{n}{6}\right)^{\frac{1}{2}}\sqrt{\pi}}\right)^6 \cdot \left(\frac{2^{\frac{n}{3}}}{\left(\frac{n}{6}\right)^{\frac{3}{2}}\sqrt{\pi}}\right)^3 \sim \frac{2^{(2n+n)} \cdot 6^{\frac{15}{2}}}{n^{\frac{15}{2}} \cdot \pi^{\frac{9}{2}}} \sim \left(3,971.55 \cdot n^{-7.5}\right) \cdot 8^n$$

2 Bent line matchings

A slightly easier question to answer is whether or not we can find a second disjoint, non-crossing matching if we are allowed to consider not simply line segments but piecewise linear functions. We will call these “bent line” matchings. I will show that we can always find a second such matching.

We first break down the convex hull of the original matching into smaller regions, each with a property analogous to being convex hull connected, as follows: Begin with an edge that has neither endpoint on the convex hull, and extend out in a straight line in each direction until either it exits the convex hull, or it intersects a new edge in the matching. If a new edge is intersected and it does not have an endpoint on the convex hull, we continue from the intersection along the part of the edge that has a positive component in our initial direction. If the new edge does have an endpoint on the convex hull, we move towards that endpoint and exit the region. In this way, we have made a single jagged line, splitting the edges into two pieces. We assign every edge that has one endpoint on the split line and one endpoint off to belong to the side that contains the off endpoint. Finally, we assign the one edge with both endpoints on the split line to the side with an odd number of edges.

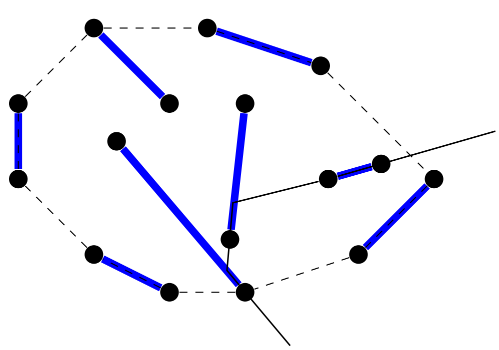

An example of a matching that requires only one split to form perimeter-connected regions.

We can repeat this process on each half of the original, until there are no edges that do not have one endpoint either on the original convex hull, or on a split line. The resulting pieces are not necessarily convex, but all the edges are tied down to the boundary, consisting of either the original convex hull or the split line. We will call this property "perimeter connected". Every edge that has one endpoint on the split line and another endpoint within a region will still be assigned to that region. Moreover, every region has an even number of edges. We will now modify the original proof for a convex-hull connected matching, so that we can match up the vertices in each piece with bent line segments. Then we will show that if we can do this, we will be able to then put these line segments together to form a non-crossing bent-line matching.

Following the proof of the convex-hull connected result, we will take alternate gaps around the perimeter. These are certainly piecewise linear, although since each region is not convex, they may not be straight line segments. We then open up a wedge around each segment that extends into the region, and apply the polygon lemma to match up the remaining points. It remains only to be shown that we can move these alternate gaps slightly, while still keeping them piecewise linear, so that they do not overlap with any edges in the original matching, or with each other.

For every pair of points x and y in our original matching, consider the finite set of perpendicular distances between the line that x and y determine, and the other points in S. Since this is a finite set it has some minimum, and moreover there is some δ that is less than the minimum for all pairs x and y. We will take ϵ to be less than $\frac{\delta}{2}$. For any region R, we can surround its perimeter in with balls of radius ϵ, without containing any points that were not already on the perimeter of R. From our original bent line segment, we will create a new segment as follows: replace each intermediate point with a new point, so that the following two conditions hold:

- The new point is on the correct side of any edges of the original matching that run along the perimeter and were assigned to the region (this can only happen if the perimeter runs along an edge for a while and then turns away - we want to move the turning point in the direction that the path turns). In this case, we will move in at most $\epsilon/2$.

- The new point is moved into the interior of the region if the line segment on either side of that point contains a point on the perimeter of R that was assigned to the other region. In this case we will move in by at least $\epsilon/2$, but less than ϵ.

We see that we can successfully move the edges in our matching so that they don't intersect with any of the edges in the original matching. We also need to verify that we will have avoided any intersections between these bent edges. Since when we move to avoid points assigned to another region we step in, and the regions are connected and planar, we will never by moving in create a problem. We may have to edge into another region when we are avoiding an edge that was assigned to our region, but we will only edge out by at most $\epsilon/2$, while the other path that may run along that perimeter will move in at least $\epsilon/2$ (but at most ϵ), so that we will again be able to avoid intersection.

2.1 Two Bends

The piecewise linear condition in the above result is not very limiting, since any curve can be approximated with a piece wise linear path. With some more detailed analysis, however, we will show we can find a second matching where there are no more than two bends in any line segment. Consider the portion of the jagged perimeter between two "adjacent" points on the perimeter of a region. Some number of the bends in that path will be reflex points, and some will not. We can represent the path between these two points p_1 and p_2 with the list of intermediate points. If they are reflex points, we will denote them r_i, and if not we will denote them x_i.

Lemma 1. *There can be at most two reflex points between adjacent vertices p_1 and p_2 on the perimeter of a region.*

Proof. Every time we have a reflex point, it must be from a time when the jagged split line intersected an edge of the original matching, and then turned away from our region, making one of the endpoints of that edge contained in the interior of our region. If that is the case, then both endpoints of that edge are assigned to our region. It follows that on one of the straight line segments that form the sides of the angle at that reflex, there is one point that has been assigned to our region. From this it follows that in the path between two adjacent points on the perimeter there can be at most two reflex points. □

We have no control over the number of bends in the perimeter that are not reflex points, but we will attempt to replace the jagged path along the perimeter given by $p_1, x_1, x_2, \ldots x_k, r_1, x_{k+1}, \ldots, x_{k+m}, r_2, x_{k+m+1}, \ldots, p_2$, with the shorter path p_1, r_1, r_2, p_2. Since we only skip over points that are not reflex points, this new path does not cross over this portion of the perimeter.

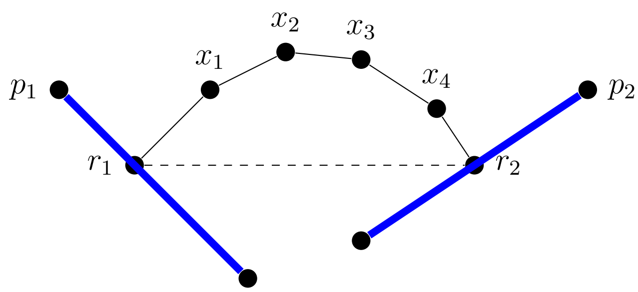

Two reflex points.

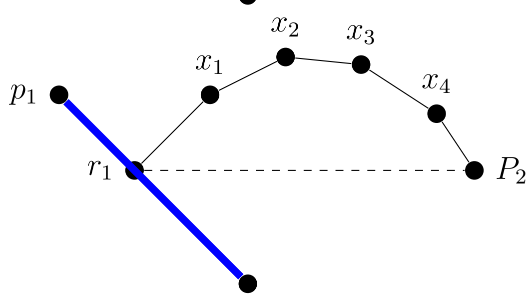

One reflex point.

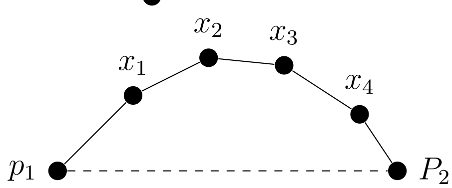

Zero reflex points.

If we can replace the perimeter with this shorter path, without intersecting any edges within our own region, we have successfully reduced the number of bends needed in that edge to at most two. However it's possible that if we do this replacement we will cross internal edges, and possibly even cross over some other part of the perimeter. If we do cross the perimeter, however, there must have been another reflex point, and we must have crossed over the corresponding edge. Therefore it suffices to avoid crossing over any edges. Consider X, the set of points in the area between the original jagged perimeter and the proposed shorter segment. Take the convex hull of $X \cup \{p_i, p_{i+1}\}$, where p_i and p_{i+1} are the endpoints of this proposed segment. The convex hull will include the segment $p_i, p - i + 1$. We will replace that segment in our path with the other portion of the convex hull. This is again made of straight line segments, and if we take that as our perimeter we will not have crossed any internal edges.

We will, however, have created (potentially many) splitting edges: edges that have both endpoints on the perimeter of the region, but not at adjacent points on the perimeter. Since every edge had at least one endpoint on the perimeter, any point we add to the perimeter will now have both endpoints on the perimeter, and the cannot be adjacent. We will break our region into smaller regions based on these splitters, so that at every step we maintain an even number of edges in each region. This process will certainly terminate, since we must have at least four edges to have this situation occur, and every time we split a region we reduce the number of edges in each sub-region.

We can continue this algorithm until every portion of the perimeter is made up of piecewise linear segments with at most 3 pieces (i.e. two bends). We then follow the result above and take alternate gaps, shifting slightly as necessary so that they don't overlap each other or the edges of the original matching. We then apply the polygon lemma to match the remaining points, and we have successfully found a bent-line matching where every edge has at most two bends.