Annie Laurie Mauhs-Pugh

Hermit Points on an Octahedron

Abstract

Consider a three-dimensional polyhedron. If we define the distance between two points on the polyhedron to be the length of a shortest path between them – restricting ourselves to the surface of the polyhedron – then hermit points are those points that are furthest apart. The question of hermit points appears quite simple. Yet determining the location of hermit points has proved impressively complex. It was not until around 2008 that papers conclusively solving the location of hermit points on a rectangular prism were published [3]; these papers show that the hermit points on a box are not necessarily the diagonally opposed vertices that might seem to most intuitive. In the following pages we will explore the hermit points on a regular octahedron. Consider a point p on a regular octahedron O. By working with the idea of the cut locus p, examining various unfolding maps of O, and looking at voronoi regions on these unfolding maps, it can be seen that the furthest point from p is its antipode. From this conclusion it follows quickly that the hermit points of O are the vertices of O.

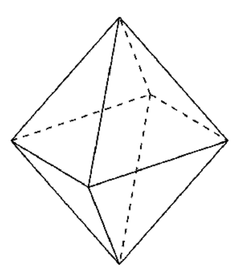

Introduction

Let's say you take two tiny people - we'll call them Alf and Meg - and plop them down on the surface of an octahedron. Looking around Alf will realize that there are infinitely many paths he can use to get to Meg, staying on the octahedron's surface. It follows from a generalization of the Hopf-Rinow Theorem to length spaces [**?**] that one of these paths achieves the infinum length of all these paths . We'll call any such path a trail from Alf to Meg. Now let's suppose Alf and Meg actually cannot stand each other (you chose rather poorly in populating this octahedron). Where can Alf and Meg stand such that the trail between them is a long as possible?

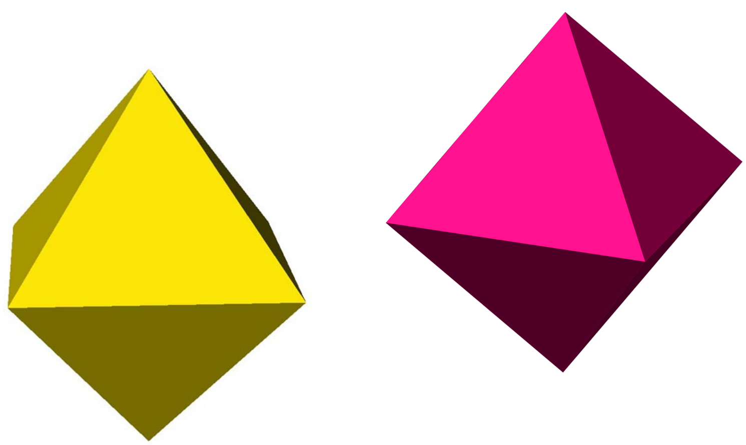

These two points - the points on the octahedron with the maximum trail between them - are what we will call the hermit points of the octahedron . This is the term used by Hess et al. when examining the same question on a rectangular prism [3]. In the following pages we examine the hermit points on a regular octahedron: an octahedron whose faces are all equilateral triangles.

Section 1

We start with an investigation of geodesics on the surface of an octahedron.

Definition For our purposes, a geodesic γ on a body M is a path that is locally distance-minimizing. More formally, a continuous map $\gamma : I \to M$ on an interval I of the reals is a geodesic if for every $t \in U$ there is a δ such that for $t_1 \in (t-\delta, t)$ and $t_2 \in (t, t+\delta)$ γ is the distance minimizing path between $\gamma(t_1)$ and $\gamma(t_2)$.

Note that according to this definition the interval I upon which a geodesic γ is defined can be *any* interval of the reals: closed, half-closed, or open. We also do not require that γ be normalized.

Proposition 1.1. *If $\gamma : [0, \epsilon] \to M$ is a distance minimal path between $\gamma(0)$ and $\gamma(\epsilon)$, then γ is a geodesic.*

Proof. Suppose γ is not a geodesic. Then γ is not locally distance minimizing and so there is a point $t \in [0, \epsilon]$ where γ does not minimize. Hence we can choose $t_1 < t < t_2$ with a path $\sigma : [t_1, t_2] \to M$ such that $\sigma(t_1) = \gamma(t_1)$, $\sigma(t_2) = \gamma(t_2)$, and σ is strictly shorter that γ from $\gamma(t_1)$ to $\gamma(t_2)$. Define $\rho : [0, \epsilon] \to M$ by

$$\rho(t) = \begin{cases} \gamma(t) & \text{if } t \in [0, t_1) \cup (t_2, \epsilon] \\ \sigma(t) & \text{if } t \in [t_1, t_2] \end{cases}$$

Then the distance from $\gamma(0)$ to $\gamma(\epsilon)$ along ρ is strictly shorter than the distance along γ, a contradiction. $\square$

Now let us examine geodesics on an ocathedron. Clearly if we have a geodesic γ restricted to the face of an octahedron then γ must be a straight line.

Proposition 1.2. *Any geodesic containing an edge point of an octahedron must pass through the edge at angle complement to the angle of indicence. That is, suppose $\gamma : (-\epsilon, \epsilon) \to O$ is a geodesic and $\gamma(0)$ is an edge point of O. Choose a point p that lies on the same edge as $\gamma(0)$. Then there must be a $\delta < \gamma$ such that for $t_1 \in (-\delta, 0)$ and $t_2 \in (0, \delta)$ the sum of the angle between rays $\gamma(t_1)\gamma(0)$ and $\gamma(0)p$, with the angle between the rays $\gamma(t_2)\gamma(0)$ and $\gamma(0)p$, must equal π.*

Proof. Suppose to the contrary that there is no such δ. Then for any neighborhood V of 0 we can choose $t_1 < 0 < t_2 \in V$ such that the sum of the angles between the rays $\gamma(t_1)\gamma(0)$, $\gamma(0)p$ and $\gamma(t_2)\gamma(0)$, $\gamma(0)p$ is not π. Choose t_1 and t_2 close enough to 0 so that $\gamma(t_1)$ and $\gamma(t_2)$ lie on the faces on either side of $\gamma(0)$. Let σ be a line between $\gamma(t_1)$ and $\gamma(t_2)$ such that σ passes through the edge containing $\gamma(0)$ at angle complement to its angle of incidence. Let $A \subset O$ denote the area bounded by γ and σ. We can unfold A, mapping it isometrically to a subset of the euclidean plane (where the metric on A is just the length of a minimal geodesic between two points). The image of A under this unfolding is a triangle, and applying the triangle inequality we see that σ is a shorter path from $\gamma(t_1)$ to $\gamma(t_2)$ than γ. Since we can find such t_1, t_2 for any neighborhood of 0, this contradicts the minimality of γ. □

Proposition 1.3. *Consider an octahedron O. No geodesic on O can contain a vertex as an interior point. That is, if a geodesic $\gamma : I \to O$ has a vertex in its image then I must contain at least one of its endpoints, the image of which is a vertex of O.*

Proof. Suppose on the contrary that $\gamma : (-\epsilon, \epsilon) \to O$ is a geodesic and $\gamma(0)$ is a vertex of O. Consider points $a \in (-\epsilon, 0]$ and $b \in [0, \epsilon)$. If $\gamma(a)$ and $\gamma(b)$ lie on the same face, then $\gamma(a)$, $\gamma(0)$, and $\gamma(b)$ form a triangle. Hence the distance from $\gamma(a)$ to $\gamma(b)$ is less than that along γ, a contradiction.

If $\gamma(a)$ and $\gamma(b)$ lie on adjacent faces then the angle sum between the rays $\gamma(a)\gamma(0)$ and $\gamma(b)\gamma(0)$ must be less than π. This follows from the fact the sum of the angles of the four faces meeting at any vertex must be less that 2π and that the angle in each face must be the same. Hence, by the triangle inequality, the distance between $\gamma(a)$ and $\gamma(b)$ must be less than the distance along γ.

Finally, suppose $\gamma(a)$ and $\gamma(b)$ are on opposite faces. Since the sum of the angles of the faces meeting at $\gamma(0)$ must be less that 2π, there must be a direction of travel from $\gamma(a)$ to $\gamma(b)$ such that the rays $\gamma(a)\gamma(0)$ and $\gamma(b)\gamma(0)$ have an angle sum less than π in this direction. Again applying the triangle inequality we see that the distance from $\gamma(a)$ to $\gamma(b)$ must be less than the path traced by γ. In all cases the distance from $\gamma(a)$ to $\gamma(b)$ is strictly less than the distance along γ, contradicting the minimality of γ. □

Definition Given two points p and q on an octahedron O I will call a distance minimal geodesic between them a *trail* between p and q. The existence of such a geodesic follows from an extension of the Hopf-Rinow Theorem to length spaces [**?**].

Define the distance $d(p, q)$ from p to q as the length of a trail γ between p and q. It is easy to check that this is indeed a metric on O.

Definition For any point p on an octahedron O, let F_p denote the set of all points furthest from p. In other words,

$$F_p = \{q \in O | d(r, p) \leq d(q, p) \forall r \in O\}$$

.

Definition Let p be a point on an octahedron O. Then γ_p will denote any geodesic starting at p. That is, γ_p is a geodesic $\gamma_p : [0, \infty) \to O$ or $\gamma_p : [0, \epsilon) \to O$ such that $\gamma_p(0) = p$.

Definition For any geodesic γ_p, γ_p minimizes on a subset of $[0, \infty)$ which, by continuity, is either of the form $[0, t_o]$ or $[0, \infty)$. If the former is true we call $\gamma_p(t_o)$ the *cut point* of p along γ_p.

Definition Let p be a point on a closed surface O. The *cut locus* of p is the collection of all cut points of p.

Proposition 1.4. *Let p be a point on an octahedron O and let C_p denote the cut locus of p. Then $F_p \subset C_p$.*

Proof. I prove that if $q \notin C_p$ then $q \notin F_p$. Suppose a point $q \in O$ is not in C_p. Let $\gamma_p : [0, t_1] \to O$ be a distance-minimal geodesic from p to q with $\gamma(t_1) = q$. As q is not in the cut locus of p, it is not a cut point of γ_p and hence we can find an ϵ such that γ_p is a minimizing path on $[0, t_1 + \epsilon]$. Thus $d(p, \gamma_p(t_1 + \epsilon)) > d(p, q)$, and so $q \notin F_p$. □

Theorem 1.5. *Let p be a point on an octahedron O. Suppose q is in the cut locus of O. Then one of the following must hold:*

1. *q is a vertex of O*
2. *There are two geodesics $\gamma, \rho : [0, 1] \to O$ such that $\gamma(0) = \rho(0) = p$, $\gamma(1) = \rho(1) = q$, and the distances from p to q along γ and ρ are the same.*

Proof. I first show that if v is a vertex of O then v is in the cut locus of p. Let $v_o \in O$ be a vertex. Note than O is path-connected. As such, the set of all paths from p to v_o is non-empty and we can choose a distance-minimal path from p to v_o. This distance-minimal path must be a geodesic (by proposition 1.1), so we can let $\gamma_p : [0, \epsilon] \to O$ be a normalized minimal geodesic with $\gamma(\epsilon) = v_o$. We know that v_o cannot be an interior point in the range of γ and hence $\gamma(\epsilon) = v_o$ must be the cut point of p along γ.

Now suppose that $\gamma_p : [0, \infty) \to O$ is a normalized geodesic and $\gamma_p(t_o)$ (the cut point of p along γ_p) is not a vertex of O. Because $\gamma_p(t_o)$ is not a vertex we can extend γ_p beyond t_o to $t_o + \epsilon$. Construct a sequence of geodesics $\{\sigma_j\}$ as follows:

- take a sequence of real numbers $\{\epsilon_n\}$ such that $\epsilon_n < \epsilon$ for all n and $\epsilon_n \to 0$.
- let σ_j be the normalized minimizing geodesic from p to $\gamma_p(t_o + \epsilon_j)$.

Consider $\{\sigma_j'(0)\}$. At this point the argument varies slightly by whether or not p is a vertex, and edge point, or a face point. We will step through the argument carefully for the case where p is a vertex, and not repeat it in entirety for the other two cases.

If p is a vertex then consider each of its four adjacent faces along with their corresponding $\sigma_{j_i}(0)$ (the σ_{j_i} that leave p through that particular face). In one of these four subsets of $\{\sigma_j\}$ there must be a subsequnce $\{\sigma_{j_n}(0)\}$such that $\{\sigma_{j_n}'(0)\}$ is convergent. Since the plane along this side of the octahedron is compact, $\sigma_{j_n}'(0) \to \sigma'(0)$ for some σ. If p is on an edge or a face we can use similar arguments to find such a σ. By continuity, σ must be a distance minimizing geodesic to $\gamma(t_o)$. I claim that $\sigma \neq \gamma$.

Suppose to the contrary that $\sigma = \gamma$. For the sake of simplicity, look first at the case where p is a vertex point. Since $\sigma = \gamma$, σ and γ must leave p along the same face. Call this face $F \subset O$ and consider the space of tangent vectors at p of curves that leave p through F. Take the exponential map on this space, exp_{F_p}, taking tangent vectors to O. By definition, $\sigma_j(t_i) = \gamma(t_o + \epsilon_j)$ for some $t_i < t_o + \epsilon_j$. Since $\sigma_j' \to \gamma'$, there must be a σ_j in any neighborhood of γ and so for any neighborhood U of $t_o\sigma'(0)$ there must be a j such that $t_i\sigma_j'(0), (t_o + \epsilon_j)\gamma \in U$ and $exp_{F_p}(t_i\sigma_j'(0)) = exp_{F_p}((t_o + \epsilon_j)\gamma$. This means that exp_{F_p} is not singular at $t_o\gamma'(0)$, which, since O is essentially flat, cannot be true.

The cases where p is an edge point or a face point are very similar: in the edge case we look at one of the two faces on either side of p, and in the face case we can actually use the tangent plane at p. In both cases the contradict the singularity of the esponential map, showing that we must have $\sigma = \gamma$ $\square$

Definition Intuitively, an *unfolding* of an ocathedron O is a flat shape obtained by cutting O so that it can lie flat. We formalize this notion by saying that a subset $U \subset \mathbb{R}^2$ is an *unfolding* of an octahedrom O if U is connected, closed, and there is a continuous surjective map $f_U : U \to O$ such that f_U is a local isometry and f_U resricted to the interior of U is an open map. We will call such a map a *folding map* from O to U.

Note that this definition become problematic when we try to examine an unfolding that requires different faces of an octahedron to overlap. A more accurate definition - one that accounts for such overlap - would define an unfolding as a locally isometric immersion of the universal cover of $O\backslash\{v\}$, where $\{v\}$ are the vertices of O. However, this definition is more cumbersome and so for now we will leave our definition as it stands.

This definition allows for a point p on an octahedron O to be represented by multiple points on an unfolding. In fact, much of this paper will be spent examining an unfolding on which several faces of an octahedron are represented twice.

Proposition 1.6. *Let $U \subset \mathbb{R}^2$ be an unfolding of an octahedron O and $f_U : U \to O$ an unfolding map. Then if $\gamma : I \to O$ is a geodesic on O, $f^{-1}(\gamma(I))$ is the disjoint union of straight lines.*

Proof. First suppose $\gamma : A \subset \mathbb{R} \to O$ is a geodesic on O, and take $t \in A$. We then have three cases:

1. t is not an endpoint of A and $f^{-1}(\gamma(t))$ is not in the boundary of U. Then we can choose ϵ such that for all $t_1, t_2 \in (t-\epsilon, t+\epsilon)$, γ is the minimizing geodesic between $\gamma(t_1)$ and $\gamma(t_2)$, and $f^{-1}(\gamma(t_1))$ and $f^{-1}(\gamma(t_2))$ lie in a neighborhood of $f^{-1}(\gamma(t))$ on which f is isometric. Take $t_1 < t < t_2 \in (t-\epsilon, t+\epsilon)$. I claim that that $f^{-1}(\gamma(t))$ lies on the line between $f^{-1}(\gamma(t_1))$ and $f^{-1}(\gamma(t_2))$. For if $f^{-1}(\gamma(t))$ did not lie on the same line as $f^{-1}(\gamma(t_1))$ and $f^{-1}(\gamma(t_2))$ then we would have, by the triangle inequality on $U \subset \mathbb{R}^2$,
$d(\gamma(t_1), \gamma(t_2)) = d(f^{-1}(\gamma(t_1)), f^{-1}(\gamma(t_2)))$
$< d(f^{-1}(\gamma(t_1)), f^{-1}(\gamma(t))) + d(f^{-1}(\gamma(t)), f^{-1}(\gamma(t_2)))$
$= d(\gamma(t_1), \gamma(t)) + d(\gamma(t), \gamma(t_2))$

$= d(\gamma(t_1), \gamma(t_2))$,

a contradiction.

Thus for all $t_1 < t < t_2 \in (t-\epsilon, t+\epsilon)$, $f^{-1}(\gamma(t))$ must lie on the same line as $f^{-1}(\gamma(t_1))$ and $f^{-1}(\gamma(t_2))$. This means that $f^{-1} \circ \gamma$ must be locally straight at t. Since t was arbitrary – under the conditions that it not be an endpoint of A and $f^{-1}(\gamma(t))$ not be in the boundary of U – $f^{-1} \circ \gamma$ must be a straight line across any continuous stretch of all such t. For take a closed interval $[0, \epsilon]$ of such t. We can cover $[0, \epsilon]$ with open intervals which $f^{-1} \circ \gamma$ map to straight segments. By compactness be can take a finite subcover. As this is a cover of open intervals, the intervals must overlap. Patching these intervals together we see that their pre-image under $\gamma^{-1} \circ f$ – the pre-image of $\gamma([0, \epsilon])$ – must be a straight line.

2. $f^{-1}(\gamma(t))$ is in the boundary of U. If we can find an ϵ such that f is isometric on $(f^{-1} \circ \gamma)((t - \epsilon, t + \epsilon))$ then we can apply the above argument.

   Suppose then that we can find no such ϵ. I claim that in this case $f^{-1} \circ \gamma$ is discontinuous at $f^{-1} \circ \gamma(t)$. We know from 1.) that, locally, $f^{-1} \circ \gamma$ approaches $f^{-1} \circ \gamma(t)$ in a straight line. By assumption, we cannot find an ϵ such that $(f^{-1} \circ \gamma)((t - \epsilon, t + \epsilon))$ is a connected straight line on U. Thus either $f^{-1} \circ \gamma$ passes through $(f^{-1} \circ \gamma)(t)$ at an acute angle or is discontinuous at $(f^{-1} \circ \gamma)(t)$. Since the first option would mean that f is not locally isometric at $(f^{-1} \circ \gamma)(t)$, $f^{-1} \circ \gamma$ must be discontinuous at $(f^{-1} \circ \gamma)(t)$.

3. Finally, suppose t is an endpoint of A but $f^{-1}(\gamma(t))$ is not in the boundary of U. By the continuity of γ and the continuity of f^{-1} on the interior of U, $f^{-1}(\gamma(t))$ is either an endpoint of $f^{-1} \circ \gamma$ or at a point of self-intersection of $f^{-1} \circ \gamma$.

The above conclusions, taken together, show that $f^{-1} \circ \gamma(A)$ is a collection of straight lines with any discontinuities occuring at the boundary of U and the image of the the endpoints of A.

□

Lemma 1.7. *If $L \subset U$ is a straight, closed line then $f(L)$ has the same length as L*

Proof. Since L is a geodesic on U we can cover L with open intervals $\{U_i\}_{i \in I}$ on which L is distance minimizing. Furthermore, as f is a local isometry we can cover L with open intervals $\{V_j\}_{j\in J}$ on which f is isometric. Take $\{U_i \cap V_j | i \in I, j \in J\}$. This is an open cover of L by sets on which L is minimal and f is isometric. By compactness we can take a finite subcover $\{W_1, ..., W_n\}$. In order to cover L these intervals must overlap and we can choose points $q_1, ..., q_{n+1}$ from the overlaps such q_1 and q_{n+1} are the endpoints of L, L is distance minimizing from q_i to q_{i+1} for all i, and f is isometric on $\{q_i, q_{i+1}\}$. These properties imply that the length of $f(L)$ is the sum of the distances between q_i and fq_{i+1} for all i, which is exactly the length of L. $\square$

Lemma 1.8. *If $L \subset U$ is a straight line then $f(L)$ is a geodesic.*

Proof. This follows directly from the fact that straight lines are geodesics on any subset of $\mathbb{R}^2$, the definition of distance, and that f is a local isometry.

$\square$

Definition Take a metric space X and a subset $\{p_j\}_{j\in J} \subset X$. Then the *voronoi region* R_j associated with the point p_j, for some $j \in J$, is the set

$$R_j = \{x \in X | d(x, p_j) \leq d(x, p_i) \forall i \neq i\}$$

I will define the *voronoi lines* of $\{p_j\}_{j\in J}$ to be

$$\begin{aligned} L &= \{x \in X | \exists j \neq i \text{ s.t. } x \in R_j \cap R_i\} \\ &= \{x \in X | \exists j \neq i \text{ s.t. } d(x, p_j) = d(x, p_j) \leq d(x, p_k) \forall k\} \end{aligned}$$

. Intuitively, the voronoi lines are the boundaries of the voronoi regions.

Theorem 1.9. *Take a point p on an octahedron O and an unfolding U of O with unfolding map f_U. Let $\{q_j\}$ denote the set of all points in O such that there are geodesics $\gamma \neq \sigma$ that are minimal from p to q_j. Suppose that for every $\gamma_p : [0, \infty) \to O$ with cut point t_o and $\gamma_p(0) = p$ the following hold:*

- *There is a continuous straight line $L \subset U$ such that $f_U(L) = \gamma([0, t_o])$.*
- *f_U is injective on both $f_U^{-1}(\{q_j\})$ and the voronoi lines of $f_U^{-1}(p)$.*

Then the image under f_U of the voronoid lines of $f_U^{-1}(p)$ along with the vertices of O is exactly the cut locus of O.

Proof. Our conditions suppose that each geodesic γ_p from p can be represesented as a straight line on U for as long as γ_p is globally minimal. With this in mind, let C_p be the cut locus of p, $\{v_j\}$ be the vertices of O, and $L_{f_U^{-1}(p)}$ be the voronoi lines of $f_U^{-1}(p)$. We wish to show that $f(L_{f_U^{-1}(p)}) \cup \{v_j\} = C_p$.

By Theorem 1.5, C_p is the the set $\{q_j\} \cup \{v_j\}$. Thus it is sufficient to show $f_U(L_{f^{-1}(p)}) = \{q_j\}$. First take $q_i \in \{q_j\}$. Then we can find $\gamma \neq \sigma : [0, t_o] \to O$, minimal geodesics with $\gamma(t_o) = \sigma(t_o) = q_i$. By supposition there are two straight lines L_1 and L_2 on U such that $f_U(L_1) = \gamma([0, t_o])$ and $f_U(L_2) = \sigma([0, t_o])$

Thus we have two distinct straight lines from $f_U^{-1}(p)$ to $f_U^{-1}(q)$ for each point in these two sets, and so $f_U^{-1}(q) \subset L_{f_U^{-1}(p)}$. Since we supposed that f_U was injective on $\{q_j\}$, this means that there are two distinct, distance minimal lines from $f_U^{-1}(p)$ to $f_U^{-1}(q)$ and so ${}_U f^{-1}(q) \in L_{f_U^{-1}(p)}$. As f_U is continuous on the interior of U, L_1 and L_2 must have endpoints that map to q_i. And since f_U is injective on $f_U^{-1}(\{q_j\})$, these endpoints must be the same. Hence $f_U^{-1}(q_i)$ is in the voronoi lines of $f_U^{-1}(p)$ and so $\{q_j\} \subset f_U(L_{f_U^{-1}(p)})$.

Now suppose $q \in f_U(L_{f^{-1}(p)})$. Then there is a point $x \in U$ such that $f_U(x) = q$ and we can choose $y, z \in f_U^{-1}(p)$ such that there are two distinct minimum length, straight lines from y, z to x. Call these lines L_1 and L_2. We know that $f_U(L_1)$ and $f_U(L_2)$ are globally minimizing geodesics on O, since all globally minimizing geodesics can be represented as straight lines on U. Furthermore, as f_U is both a local isometry and injective at x, there must be a neighborhood about x on which f_U is injective. Thus we must have $f(L_1) \neq f_U(L_2)$. Hence $f_U(x) = q \in \{q_j\}$ and so $f_U(L_{f_U^{-1}(p)}) \subset \{q_j\}$.

□

Section 2

I now examine the cut locus of a regular octahedron, using the above claim to determine F_p for any p on a regular octahedron. In order to understand the cut locus we will spend the next several pages examining the unfolding which is shown on the next page. As this unfolding depends on the point p we will call it U_p and denote its unfolding map by f_{U_p}.

Boundary Lines on an Unfolding of the Octahedron

I start by defining some terminology that will make this examination easier. For the rest of this section we will assume that we have an octahedron O and have chosen a point p of the surface of O.

Definition The *starting face* of p is the face of O on which p resides (if p is on a vertex or an edge it will not matter which of its faces we choose as the starting face).

Definition The *secondary faces* are those faces who share an edge with the starting face.

I will use the terms *starting face* and *secondary face* to refer both to the faces on the octahedron and their pre-image under an unfolding map.

For the rest of this section we will focus on U_p. Intuitively, U_p is formed by cutting the starting face into three pieces (cutting from p to each of its vertices) and unfolding along these lines and the edges of the secondary faces. Each section of the starting face and its secondary face appears twice. A diagram both of O and of U_p are shown on the next page. Note that the arrows on the diagram of U_p represent the two ways a secondary face is portrayed, in a sense rotating the secondary face around an inner vertex.

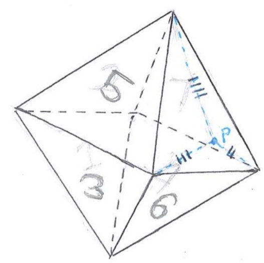

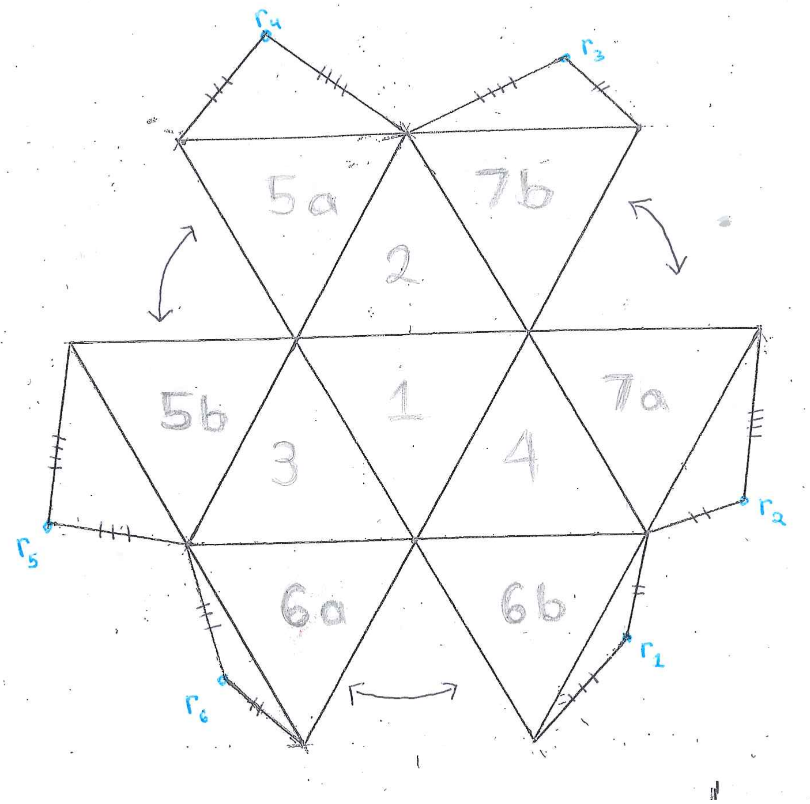

I have labelled the faces of the diagram above. Faces 5a through 7b are secondary faces. The starting face is cut into sections but will be called face 8. Face 1 is what I will call the *antipodal face* of p. On U_p I will use the term *outer vertex of a secondary face* to refer to the vertex of a secondary face that does not intersect any of faces 2, 3, or 4. The *inner vertices* of that face are the other two vertices.

Definition Note that on U_p, $f^{-1}(x)$ consists of six points. These points will be called *point representations* and labelled r_1 through r_6.

In order to examine the cut locus of the ocathedron O we will look at the voronoi lines of $\{r_1, ..., r_6\}$. Note that these voronoi lines can be determined from the perpendicular bisectors of all possible pairings of $r_1, ..., r_6$. A point z is in the voronoi lines of $f^{-1}(p)$ if and only if it lies at the first intersection of a geodesic from a point-prepresentation and one of its perpendicular bisectors.

Definition For linguistic simplicity I will call a perpendicular bisector between r_i and r_j the *boundary line* between r_i and r_j, written $r_iBr_j = r_jBr_i$.

Note that in our chosen unfolding the six point-representations lie in a roughly cirlular shape. I will say two point-representations are *one apart* if there is one point-representation separating between them (clockwise or counterclockwise), *two apart* if there are two point-representations between them, etc. In the following exploration I break the boundary lines into four categories:

- A *1st category boundary line* is a boundary line r_jBr_i such that r_i and r_j are next to each other and f is not injective on the sections of starting face that r_i and r_j lie on. In other words, r_i and r_j lie on different representations of the same section of starting face.

- A *2nd category boundary line* is a boundary line r_jBr_i such that r_i and r_j are next to each other and f is injective on the sections of starting face r_i and r_j lie on.

- A *3rd category boundary line* is a boundary line r_jBr_i such that r_i and r_j are one apart.

- A *4th category boundary line* is a boundary line r_jBr_i such that r_i and r_j are two apart.

1st Category Boundary Lines

For these argument refer to Figure 1 below. Let r_i be one of the six point-representations on U_p. r_i corresponds to one secondary face (in this case 6b), and we will refer to the vertex of this secondary face nearest the center of the unfolding as r_i's *interior point* v . Let r_j denote the point-representation gotten by rotating r_1's secondary face around v by $\frac{2}{3}\pi$. Note that $d(r_1, v) = d(r_j, v)$. Thus r_iBr_j passes through v. To determine the angle at which r_iBr_j

FIG. 1.

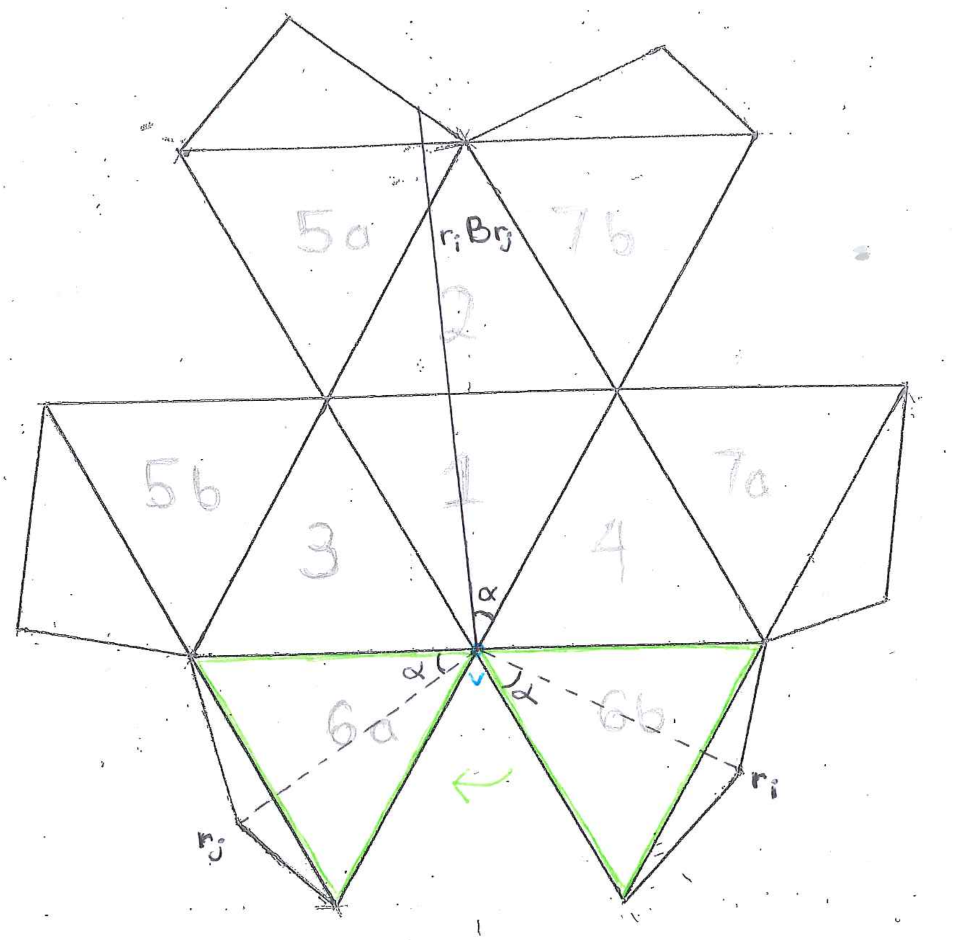

passes through v, note that this angle must be the same as the angle α formed by the line from r_i to v with the edge of the secondary face.

With this in mind we look at a partial unfolding of O that maps a triangle to the starting face and the three secondary faces – faces 5, 6, 7, and 8, shown in Figure 2. Each of the lines from the pre-image of p to a corner of the larger triangle represents a line from a point-representation r_k to its interior vertex . Thus, applying the previous logic to each r_iBr_xj

where r_{xi} and r_{xj} share the same section of starting face, we determine the relationship shown on Figures 2 and 3 below. Each such $r_{xi}Br_{xj}$ starts at one of the inner vertices of the unfolding. The angle they make with the edge of the inner face is a reflection over the angle bisector of the angles shown in Figure 2

FIG. 2.

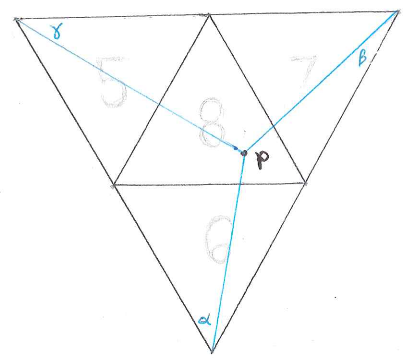

FIG. 3.

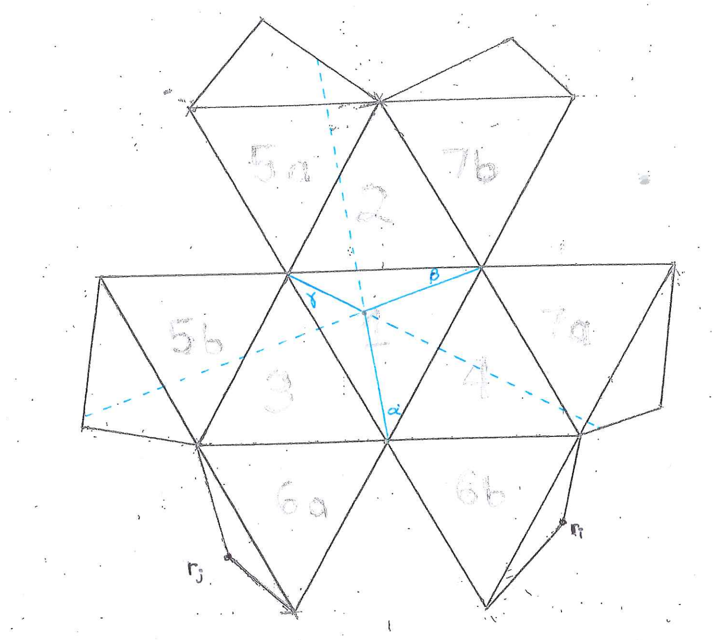

Proposition 2.1. *The three 1st categoryboundary lines intersect at one point.*

Proof. Take the triangle fin Figure 2 formed by the starting face and the three secondary faces (faces 5, 6, 7, and 8). Draw lines from the pre-image of p to each of the vertices.

Now compare that to the inner face of U with the 1st category boundary lines drawn on (on Figure 3). Both of these are triangles with lines passing through each vertex. The lines in the second case are gained from the first by reflecting them over their respective angle bisectors.

Hence we have two triangles such that the first triangle has a point in the interior and lines connecting this point to each of its three vertices. The second triangle also has lines passing through its vertices, lines which are formed by reflecting the lines in the first triangle over their perpendicular bisectors. It is known that in this situation the lines in the second triangle meet at one point. I will sketch the proof of this fact below, as I will later need to generalize it beyond triangles. The idea of this proof is credited to Doyle [**?**].

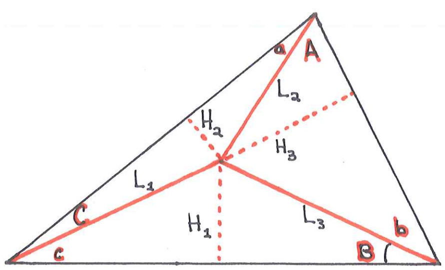

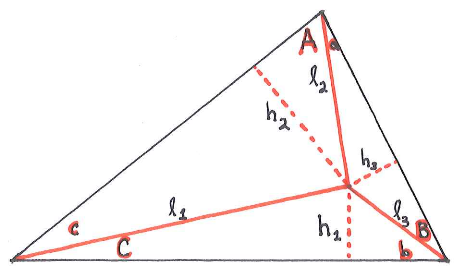

Take the first triangle (an example is shown above). The three lines in this triangle form three sub-triangles: draw a perpendicular bisector of each of these (labeled H_1 through H_3 on the diagram below). Note that $H_1 = sin(a)L_1$, $L_1 = \frac{H_2}{sin(b)}$, etc. Combining these equalities, starting at H_1 and moving clockwise, we get

$$H_1 = \frac{sin(a)}{sin(b)}\frac{sin(c)}{sin(d)}\frac{sin(e)}{sin(f)}H_1$$

$$\text{Hence } \frac{sin(a)}{sin(b)}\frac{sin(c)}{sin(d)}\frac{sin(e)}{sin(f)} = 1.$$

Now look at the second triangle and the lines gained by flipping the lines of the first triangle over their perpendicular bisectors. These lines, along with the edges of the triangle, form three triangles. The question is whether or not we can fit these three triangles together to form the larger triangle without overlapping them. Note that if start with h_1 we can shrink or stretch the subtriangle clockwise to fit it. This would make $l_1 = \frac{h_1}{sin(b)}$ and $h_2 = l_1 sin(1) = h)1\frac{sin(a)}{sin(b)}$. We can then scale the third subtriangle to fit this triangle. The question is whether or not at this point the third triangle will fit the first subtriangle. Let l_3 denote the original length of the counter-clockwise side of the first subtriangle. Let l_3'

denote the clockwise side of the third subtraingle after the above process. Then, combining trigonometric equalities,

$$l_3' = \frac{sin(c)}{sin(f)} \frac{sin(a)}{sin(d)} \frac{sin(e)}{sin(b)} l_3 = \frac{sin(a)}{sin(b)} \frac{sin(c)}{sin(d)} \frac{sin(e)}{sin(f)} l_3$$

But given what we know from above, this means $l_3' = l_3$. Thus the three subtriangles fit together.

□

2nd Category Boundary Lines.

Figure 4 below will be helpful in understanding this argument. Take a point-representation r_i. Let r_j denote the point-representation neighboring r_i that does not share a secondary face with r_i. The secondary faces corresponding to r_i and r_j share a common vertex, v. Note that $d(r_i, v) = d(r_j, v)$, since on the octahedron these sections of the starting face meet at p. Thus r_iBr_j must pass through v. To determine the angle at which r_iBr_j passes through v note that this angle must be the same as the angle α formed by the line between r_i and v and the edge of the starting face, as shown below.

With this in mind we now examine the starting face. Note that all angles on U_p between

FIG. 4.

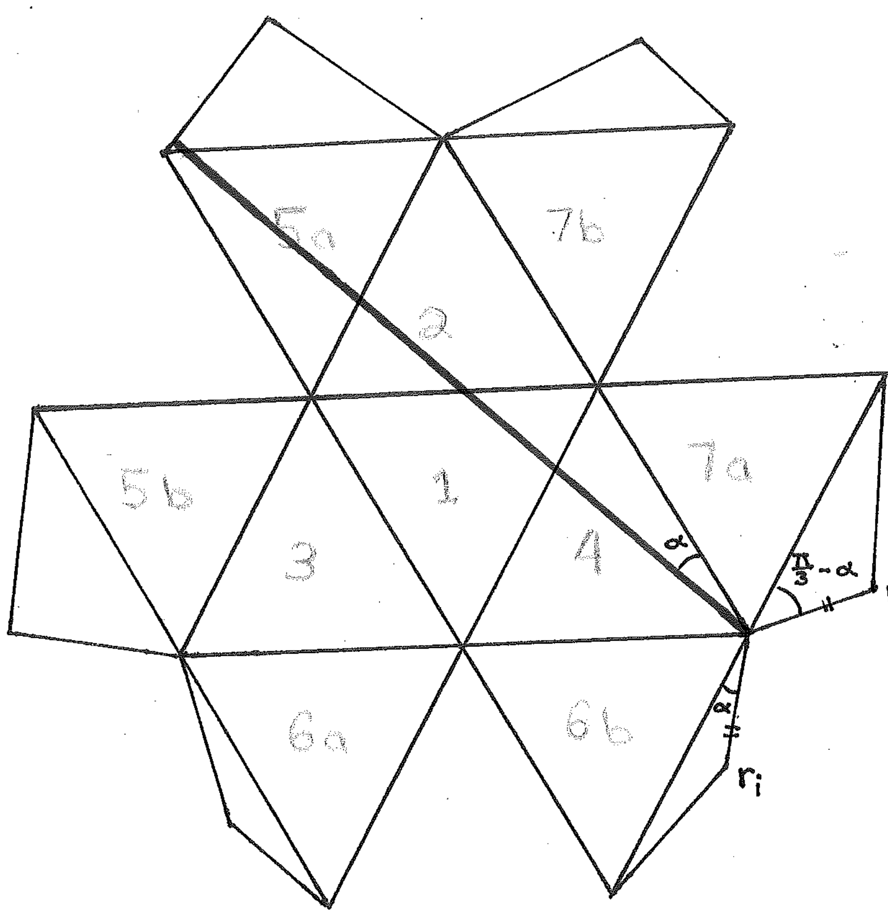

point-representations and the vertices of their starting face sections can be represented on the starting face by drawing lines from p to the three vertices of the starting face. Applying the same logic that we used to determine r_iBr_j we find the relationship shown in Figures 4 and 5: take the triangle in U_p composed of faces 1, 2, 3, and 4. Each 2nd category boundary line, r_iBr_j, passes through a vertex of this triangle. The angle each r_iBr_j forms with this triangle is a reflection over the angle bisector of the angle p makes with r_i's section of the starting face.

FIG. 5.

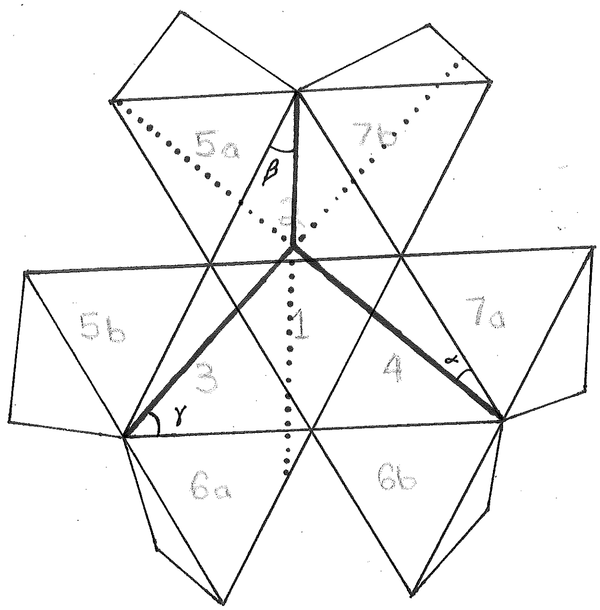

FIG. 6.

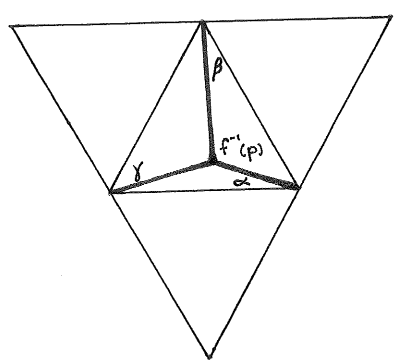

Proposition 2.2. *The 2nd category boundary lines intersect at one point.*

Proof. This proof is the same as for the 1st category boundary lines, concluding that the boundary lines intersect at what is called thier isogonal conjugate. ☐

3rd Category Boundary Lines

Once again, let r_i be a point-representation on our unfolding. Let r_j denote a point representation one away from r_i. To determine the location of r_iBr_j it will not matter if r_j is clockwise or counter clockwise from r_i.

FIG. 7.

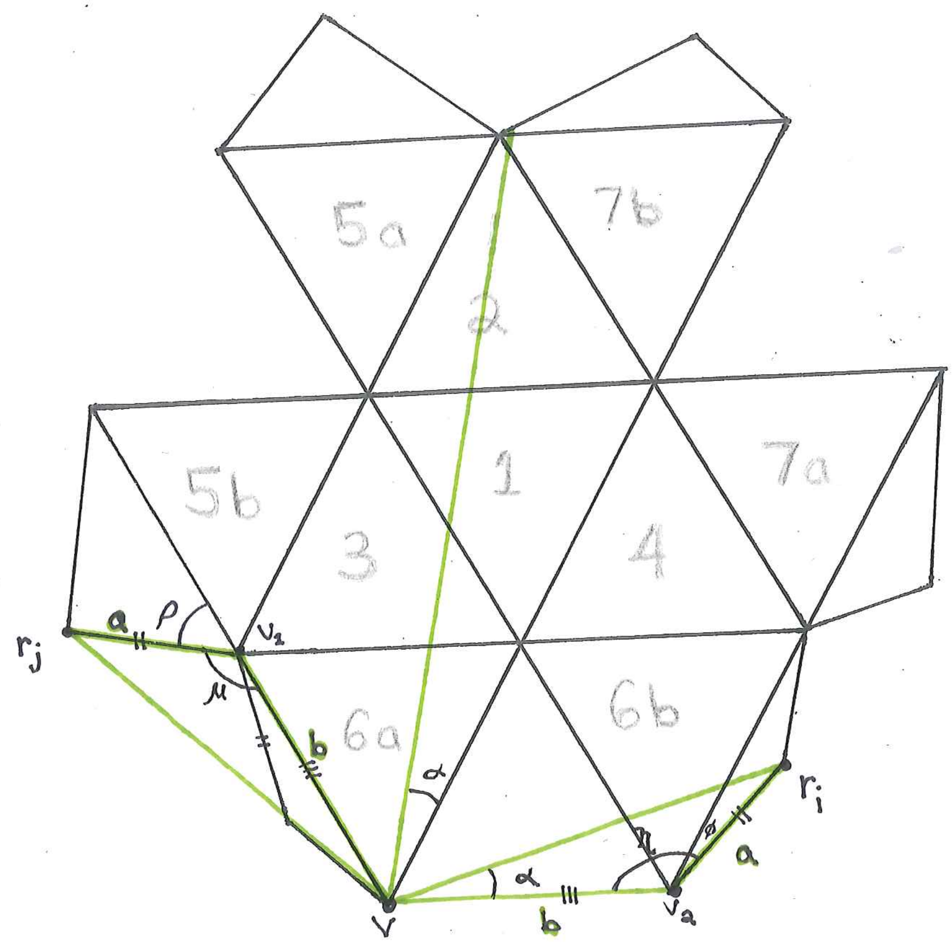

Lemma 2.3. *Let v be the outer vertex of the secondary face between r_i and r_j. Then $d(v, r_i) = d(v, r_j)$.*

Proof. To follow this proof it will be very helpful to look Figure 7 above. Draw a line between r_i and v, and a line between r_j and v. Consider the two triangles formed by these lines and the pre-images v_a and v_2 of the vertex that r_i's and r_j's secondary faces share (face 6 on the diagram). I claim that these two triangles are equivalent. We know that two of their sides are of equal length (labelled a and b in the diagram). What remains to show is that the two angles, labeled η and μ, are the same. To see this, note first that the angles labelled ϕ and ρ must sum to $\frac{\pi}{3}$. Then

$$\eta = \frac{\pi}{3} + \frac{\pi}{3} + \phi = \frac{\pi}{3} + \frac{\pi}{3} + (\frac{\pi}{3} - \rho) = \pi - \rho = \mu$$

Thus the two triangles are identical. Their longest sides, which have lengths $d(r_i, v)$ and $d(r_j, v)$, must then be of equal length. $\square$

Hence r_iBr_j must pass through v. To determine the angle at which r_iBr_j passes through v note that it must pass through v at the same angle as the angle nearest v in the triangles just examined (this angle is labeled α). α is determined by the location of p: we see that α is the angle formed by drawing a line from p to the face two away from the starting face, to the vertex of that face furthest from the starting face.

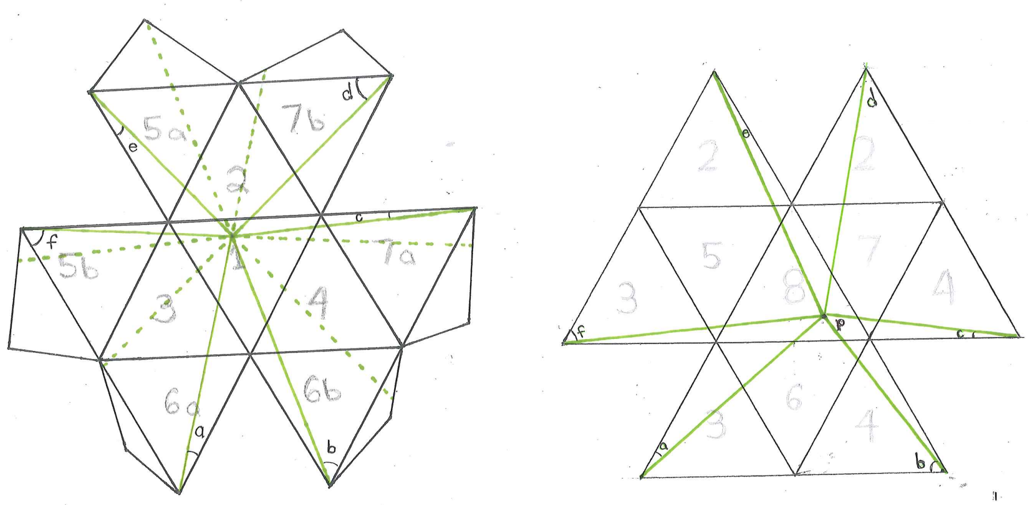

To understand α better we turn to a partial unfolding U_o of O, shown directly above. U_o is similar to U_p but has face 8 at its center. It also does not include face 1. Draw a line from $f_{U_o}^{-1}(p)$ to each of the vertices on U_o that are furthest from the starting face. We label the angles that these lines creates at the vertices a though f. These are exactly the angles formed by drawing a line from p to the face two away from the starting face, to the vertex of that face furthest from the starting face. Applying he argument from the previous paragraph to each of these angles we are able to determine each 3rd category boundary line r_iBr_j. The relationship between the two unfoldings is shown below. Each $\mathrm{t}r_iBr_j$ passes through an outer vertex of U_p at an angle that is $\frac{\pi}{3}$ less than one of the angles in U_o. Note that the angles are arranged in the same clockwise order.

Proposition 2.4. *All 3rd category boundary lines intersect at one point.*

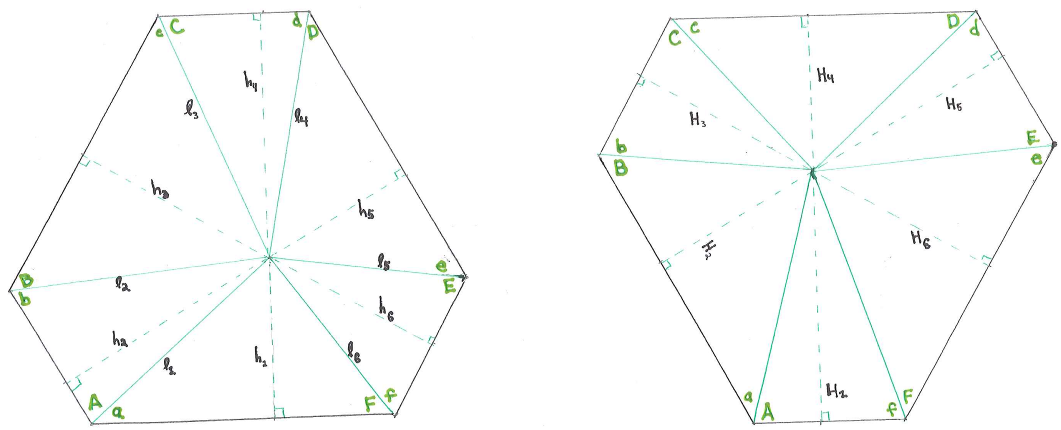

Proof. Consider the two hexagons illustrated above. Each has six lines, one starting at each vertex. The interior angles are the same in the second hexagon as in the first. A line starting at a vertex v in the first hexagon corresponds to a line in the second hexagon that has been reflected over the angle bisector at v. Looking at U_p and U_o we see that this is exactly the relationship described above in the 3rd category boundary lines: the hexagons can be gained from the unfoldings by drawing lines between each of the relevant vertices. To prove the proposition it will be sufficient to show that if the lines in the first hexagon intersect at one point then so do those in the second, since the lines in U_o intersect at $f_{U_O}^{-1}(p)$

Suppose that the lines in the first hexagon intersect in one point. Consider the triangles formed by the intersections, as shown above. We label each line segment $l_1, ..., l_6$ and draw in lines representing the height of each triangle, labeled $h_1, ..., h_6$.

We see that

$$l_1 = \frac{h_1}{sin(a)}$$
$$h_2 = l_1 f(A)$$
$$l_2 = \frac{h_2}{sin(b)}$$
$$h_3 = l_2 sin(B)$$

etc.

Substituting in we get

$$h_1 = \frac{sin(F)}{sin(f)}\frac{sin(E)}{sin(e)}\frac{sin(D)}{sin(d)}\frac{sin(C)}{sin(c)}\frac{sin(B)}{sin(b)}\frac{sin(A)}{sin(a)}h_1$$

or

$$1 = \frac{sin(F)}{sin(f)}\frac{sin(E)}{sin(e)}\frac{sin(D)}{sin(d)}\frac{sin(C)}{sin(c)}\frac{sin(B)}{sin(b)}\frac{sin(A)}{sin(a)}$$

With this is mind, turn to the second hexagon. Consider the triangles formed by the reflected lines (for example, a triangle with interior angles A and f). Since the sum of the unlabeled angles of each of these triangles must be 2π, we know that we can size these triangles such that they all meet at one vertex. The question is whether or not we can size the triangles such that their unlabeled angles meet at one vertex and their sides align such that they form a hexagon.

Start by selecting one triangle, T_1. Size the triangle that is clockwise, T_2, so that their bordering sides are the same length. Repreat this for all the triangles, moving clockwise around the hexagon. We will show that once you have completed this process, the size that T_6 requires T_1 to be is ecaxtly the sixe that T_1 originally was. In order to see this, label the heights of the traingles $H_1, ..., H_6$ and the triangle sides $l_1, ..., l_6$. Note that the triangles will fit together if and only if

$$\begin{aligned}\frac{H_1}{sin(A)} &= \frac{H_2}{sin(a)}\\ \frac{H_2}{sin(B)} &= \frac{H_3}{sin(b)}\\ &\dots\\ \frac{H_6}{sin(F)} &= \frac{H_1}{sin(f)}\end{aligned}$$

We can combine these inequalities, eliminating every H_i except for H_1. Thus the triangles will fit together if and only if

$$\frac{H_1}{sin(A)}\frac{sin(a)}{sin(B)}\frac{sin(b)}{sin(C)}\frac{sin(c)}{sin(D)}\frac{sin(d)}{sin(E)}\frac{sin(e)}{sin(F)} = \frac{H_1}{sin(f)}$$

or

$$H_1\frac{sin(a)}{sin(A)}\frac{sin(b)}{sin(B)}\frac{sin(c)}{sin(C)}\frac{sin(d)}{sin(D)}\frac{sin(e)}{sin(F)}\frac{sin(f)}{sin(F)} = H_1$$

But this must be true since we showed above that

$$\frac{sin(a)}{sin(A)}\frac{sin(b)}{sin(B)}\frac{sin(c)}{sin(C)}\frac{sin(d)}{sin(D)}\frac{sin(e)}{sin(F)}\frac{sin(f)}{sin(F)} = (\frac{sin(F)}{sin(f)}\frac{sin(E)}{sin(e)}\frac{sin(D)}{sin(d)}\frac{sin(C)}{sin(c)}\frac{sin(B)}{sin(b)}\frac{sin(A)}{sin(a)})^{-1} = 1^{-1} = 1.$$

$\square$

4th Category Boundary Lines

Consider a point-representation r_i on U_p. Let r_j denote the point-representation that two away from r_i.

Proposition 2.5. *r_iBr_j passes through the center face of tU_p (face 1), at an angle that is parallel to the perpendicular bisector of the two edges of the face nearest r_i and r_j.*

FIG. 8.

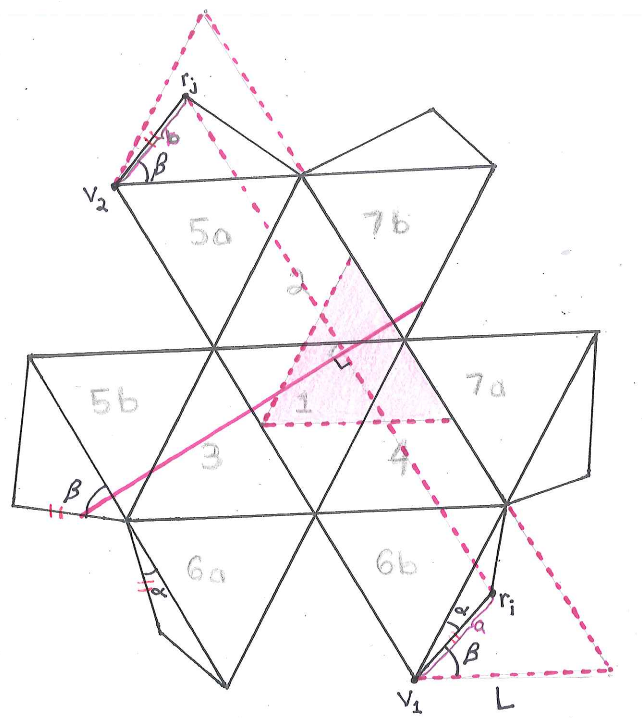

Proof. We first show that r_iBr_j is parallel to the perpendicular bisector of the edges nearest r_i and r_j. To see this we will label the two outer vertices of r_i's and r_j's secondary faces v_1 and v_2, respectively ($f_{U_p}(v_1) = f_{U_p}(v_2)$. Note that the anges labeled α and β on the unfolding below must sum to $\frac{\pi}{3}$. In addition, the lengths of the line segments labeled a and b must be the same. With this in mind, draw a line L through v_1 that is parallel to the edge of the secondary face nearest r_j. Since $\alpha + \beta = \frac{\pi}{3}$, the angle γ between L and the line from v_1 to r_i must be equal to β. Thus r_j is offset from v_1 exactly as r_i is offset from v_2. Thus the line between r_i and r_j is parallel to the line between v_1 and v_2, and so r_iBr_j is perpendicular to the line between v_1 and v_2. Thus r_iBr_j is parallel to the the perpendicular bisector of the two other edges of face 1.

We now show that r_iBr_j must intersect face 1. To do this, note that the intersection of r_iBr_j and and face 1 depends on the location of the point half-way between r_i and r_j. For the sake of notation, call this point p. Let us start when r_i is the vertex point v_2. In this case $r_j = v_1$ and p lies at the center of the line between v_1 and v_2. Since we know that a movement in r_i implies the exact same movement in r_j, this also means that a movement in r_i implies the same movement in p. Thus all possible locations of p form a triangle equal in size, shape, and orientation to the faces containing r_i and r_j, but shifted along the line between the two. Hence we can see that p is furthest from the center face when r_i is either of the vertices not v_2. In the case when r_i is the vertex furthest from the center face, r_iBr_j passes through the vertex of face 1 closest to v_2, and when r_i is the other vertex r_iBr_j passes through the vertex nearest v_1. Thus r_iBr_j must always pass through face 1. $\square$

We will stop our investigation of 4th category boundary lines at this point. The claim just proved will be sufficient for the following arguments.

Section 3

Obtaining and Examining the Cut Locus Using Boundary Lines

Theorem 3.1. *U_p is sufficient to determine the cut locus for any point p on an octahedron.*

Proof. Let C_p be the cut locus of p on O and and let $L_{f_{U_p}^{-1}(p)}$ be the voronoi lines of $f_{U_p}^{-1}(p)$ in U_p. Let $\{v_n\}$ be the set of all vertices of O and $\{q_j\}$ denote the set of all points in O such that there are geodesics $\gamma \neq \sigma$ minimal from p to q_j. Recall that according to Theorem 1.9, $f_{U_p}(L_{f^{-1}(p)}) \cup \{v_n\} = C_p$ *if*

1. for every globally minimizing geodesic $\gamma_p([0,t])$ there is a continuous straight line L such that $f_{U_p}(L) = \gamma_p([0,t])$.

2. f_{U_p} is injective on V_p and the pre-image of $\{q_j\}$.

I claim that both of these conditions hold for U_p. Start with the first condition. I will show that if t_o is the cut point of γ_p then there is a straight line whose image is $\gamma_p([0,t_o])$. This is sufficient to prove (1).

To do this, first note that is a geodesic γ_p neither passes through two distinct secondary faces (or pass through the same secondary face twice) nor passes through a vertex, then γ_p can be represented as a straight line on U_p. As we know that no geodesics can pass through a vertex, this means that if γ_p is a geodesic and $\gamma_p([0,t])$ does not pass through two secondary faces then there is a straight line $L \subset U_p$ such that $f_{U_p}(L) = \gamma_p([0,t])$.

I claim that if t_o is the cut point of γ_p then $\gamma_p([0,t_o])$ cannot pass through two secondary faces. Suppose by way of contradiction that $\gamma_p([0,t_o])$ does pass through two secondary faces of O. Consider $f_{U_p}^{-1}(\gamma_p([0,t_o]))$. γ_p passes through two secondary faces if and only if $f_{U_p} \circ \gamma_p$ does. $f_{U_p}^{-1} \circ \gamma_p$ must leave a point representation r_i in a straight line. Given our unfolding, $f_{U_p}^{-1} \circ \gamma_p$ will leave from *two* point representations, passing immediately through two different representations of the same secondary face. One of these lines will continue straight into faces 2, 3, or 4, while the other line will, in a sense, "join it there". The only way for $f_{U_p}^{-1} \circ \gamma_p$ to pass through a second secondary face of U is to continue through the triangle formed by faces 1, 2, 3, and 4 (since the only other way to reach the border of U is to go through a vertex and geodesics cannot continue through vertices).

However, in order for $f_{U_p}^{-1} \circ \gamma_p([0,t_o])$ to pass into another secondary face it must necessarily cross a 1st or 2nd category boundary line. This holds since – from what we know of 1st and 2nd category boundary lines – 1st and 2nd category boundary lines cannot lie entirely outside of the triangle formed by faces 1, 2, 3, and 4.

Thus if $\gamma_p([0, t_o])$ passes through two secondary faces then there must be a straight line L on U from a point representation r_j to $(f_{U_p}^{-1} \circ \gamma_p)(t_o)$ such that L is shorter than $(f_{U_p}^{-1} \circ \gamma_p)([0, t_o])$. Since $f_{U_p}(L)$ is a geodesic on O and the length of $f_{U_p}(L)$ is the same as the length of L, we have a geodesic from p to $\gamma_p(t_o)$ that is shorter than γ_p). This contradicts the minimality of $\gamma_p([0, t_o])$. Thus $\gamma_p([0, t_o])$ cannot pass through two secondary faces. Hence, by what was shown two paragraphs earlier, there is a straight line $L \subset U_p$ such that $f_{U_p}(L) = \gamma_p([0, t_o])$. THence (1) is proved.

It remains to show that f_{U_p} is injective on the voronoi lines of $f_{U_p}^{-1}(p)$ and on $f_{U_p}^{-1}(\{q_i\})$. In terms of the voronoi lines, recall what was shown in section 2. From this we can see that the voronoi lines must lie entirely withing faces 1, 2, 3, and 4. As f_{U_p} restricted to these faces in injective, f_{U_p} is injective on the voronoi lines.

In terms of $\{q_i\}$, suppose f_{U_p} was not injective on $f_{U_p}^{-1}(\{q_i\})$. Given the definition of U_p, this would require that there be a point $q_j \in \{q_i\}$ such that q_j lies within a secondary face. But then there would have to be two distinct minimal geodesics from p to a point on a secondary face. By what was just shown, each of these geodesics must pass only through the starting face and the secondary face on which q_i lies. Clearly there are not two distinct minimal geodesics which do this.

Thus $C_p = f_{U_p}(L_{f_{U_p}^{-1}(p)}) \cup \{v_j\}$. Hence we can determine C_p from the voronoi lines of $f_{U_p}^{-1}(p)$ and pre-images of vertices of O on our unfolding.

$\square$

Having seen that U_p is sufficient to determine the cut locus of O, I turn to F_p. Note, before we begin,that on a regular octahedron F_p cannot lie on a secondary face of p. Recall also that F_p must be a subset of the cut locus. Hence there must be a subset of $L_{f_{U_p}^{-1}(p)} \cup f_{U_p}^{-1}(\{v_j\}$ that maps to F_p. Hence to determine the subset of U_p that maps to F_p it will be sufficient to look at the voronoi lines and vertices that lie on faces 1, 2, 3, and 4. On these faces the pre-image of vertices of O are contained in $L_{f_{U_p}^{-1}(p)}$, as can be seen from our investigations in Setion I, so we need only examine the voronoi lines $L_{f_{U_p}^{-1}(p)}$. Furthermore, f_{U_p} restricted to faces 1, 2, 3, and 4 is an injection. Hence for the rest of this section we will act as though f_{U_p} is injective.

In order to determine the subset of $L_{f_{U_p}^{-1}(p)}$ that maps to F_p, introduce a function D_p. Recall that $U_o \subset U_p$ denotes the union of the faces 1, 2, 3, and 4 on the unfolding U_p. Then $D_p : U_o \to \mathbb{R}$ is defined by

$$D_p(y) = d(y, r_i),$$

where r_i is the point-representation in whose voronoi region y lies. Intuitively, D maps a point of faces 1, 2, 3, or 4 to its distance from the nearest point representation. D is well defined, since if y lies in the voronoi regions of multiple point-representations the distance to y must be the same for all the point-representations.

Proposition 3.2. *For $q \in f_{U_o}(U_o)$, $D_p(f_{U_p}^{-1}(q)) = d(q,p)$.*

Proof. Take $q \in f(U_o)$ and suppose $f_{U_p}^{-1}(q)$ lies in the voronoi region of r_i. Let L_1 denote the line from r_i to $f_{U_p}^{-1}(q)$. By Lemma 1.8, $f_{U_p}(L_1)$ must be a geodesic from p to q. Furthermore, Lemma 1.7 tells us that L_1 and $f_{U_p}(L_1)$ have the same length. Hence $D_p(f_{U_p}^{-1}(q))$ is the length of $f_{U_p}(L_1)$, and it will suffice to show that $f_{U_p}(L_1)$ is the minimal geodesic from p to q.

Suppose not. Then there must be a geodesic $\sigma : [0,t] \to O$ from p to q such that the length of $\sigma([0,t])$ is strictly less than the length of $f_{U_p}(L_1)$. Since σ is minimal from p to q, Theorem 3.1 tells us that there must be a straight line $L_2 \subset U_p$ such that $f_{U_p}(L_2) = \sigma([0,t])$. Since both $f_{U_p}(L_1)$ and σ contain q and f_{U_p} is injective on q, L_1 and L_2 must both intersect at q (their endpoints). As L_1 and L_2 are distinct lines, they must start from different point-representations; let r_j be the point representation that L_2 starts at. Then L_2 is a line from r_j to $f_{U_p}^{-1}(q)$ which has the same length as $\sigma([0,t])$, which is strictly less that the length of L_1. This contradicts the fact that $f_{U_p} f^{-1}(q)$ is in the voronoi region of r_i.

□

Lemma 3.3. *For any $p \in O$, $F_p = f_{U_p}(D^{-1}(max(D(U_o))))$. Intuitively, this means that $p \in F_p$ if and only if there is a point $x \in U_p$ such that $f(x) = p$ and x is as far from its nearest point representation as possible.*

Proof. Suppose $q \in F_p$. Then $D(f_{U_p}^{-1}(q) = d(q,p) = max\{d(r,p) | r \in O\} = max(D(U_o))$. Thus $f_{U_p}^{-1}(q) \in D^{-1}(max(D(U_o)))$ and so $q \in f_{U_p}(D^{-1}(max(D(U_o))))$.

Now suppose $q \in f_{U_p}(D^{-1}(max(D(U_o))))$. Since f_{U_p} is injective on U_o and D is only defined on U_o, this means $f_{U_p}^{-1}(q) \in D^{-1}(max(D(U_o)))$. Then $D(f_{U_p}^{-1}(q)) = max(D(U_o)) =$

$max\{d(r,p)|r \in f_{U_p}(U_o)\}$. Recalling that this maximum must be achieved at U_o, we see that $D(f^{-1}(q) = max\{d(r,p)|r \in O\}$. Hence $q \in F_p$. □

Thus in order to find F_p it will be sufficient to find the point in U_p that is furthest from its nearest point-representation. Since we showed earlier that $F_p \subset f_{U_p}(L_{f^{-1}(p)})$, we need only look for the point in $L_{f_{U_p}^{-1}(p)}$ which is the furthest from the point-representation in whose voronoi region it is contained. I will eventually show that this point is the point at which the 3rd category boundary regions intersect. From here on this point will be denoted by y_o.

In order to prove this claim we will need a better understanding of the voronoi lines and regions of $f_{U_p}^{-1}(p)$. To this end we prove the following six lemmas about the voronoi lines.

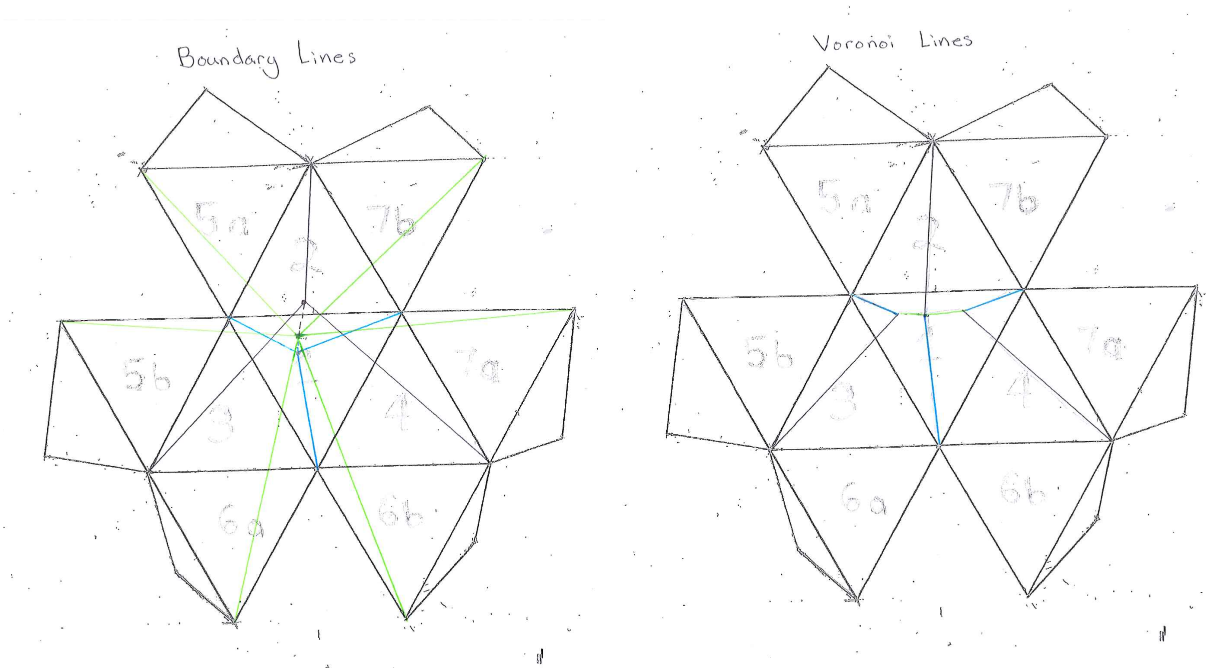

Lemma 3.4. *The union of the voronoi regions of $f_{U_p}^{-1}(p)$ must cover the entire unfolding.*

Proof. Suppose there is a point z on U such that z is contained in none of the voronoi regions of $f^{-1}(p)$. Consider a line L_1 from a point representation r_i to z. As z is not in the voronoi region corresponding to r_i, L_1 must cross a boundary line r_iBr_j. So now consider the line L_2 from r_j to z, which we know must be shorter than L_1. Since z is not in the boundary region corresponding to r_j, is must cross a boundary line r_jBr_k. Consider the line L_3 from r_k to z, which must be shorter than L_2.

Continue repeating this process, choosing lines $L_1 < L_2 < L_3 < ... < L_n$. We know that there are only finitely many point-representations, so we will eventually have to use a point-representation r_n twice. But this means that the line L_m from r_n to z is strictly short than itself, a contradiction. $\square$

Lemma 3.5. *The voronoi region of a point-representation $r_i \in f_{U_p}^{-1}(p)$ is entirely determined by its 1st, 2nd, and 3rd category boundary lines.*

Proof. Suppose otherwise. Then there is a point-representation r_i whose voronoi region is partially determined by its 4th category boundary line. That is, there is a line L leaving r_i who intersects the 4th category boundary line corresponding to r_i before it intersects any of r_i's other boundary lines. This 4th category boundary line must thus pass within the region determined by the two 3rd category boundary lines corresponding to r_i. We know from our previous investigation that the 4th category boundary line is parallel to the edge of face 1 that is closest to r_i. Thus, given the vertices that the 3rd category boundary lines must leave from, this 4th category boundary line must pass closer to r_i than y_o, in a sense "cutting y_o off from r_i" (see Figure 10). Now travel along this 4th category boundary line until it intersects another boundary line corresponding to r_i, say r_iBr_j. Consider the area A formed by the 4th category boundary line, r_iBr_j, the 3rd category boundary line that passes through r_i's secondary face, and the other 3rd category line closest to r_j. Note that this area can be contained in none of the voronoi regions of $f_{U_p}^{-1}(p)$, contradicting Lemma 3.4.

FIG. 9.

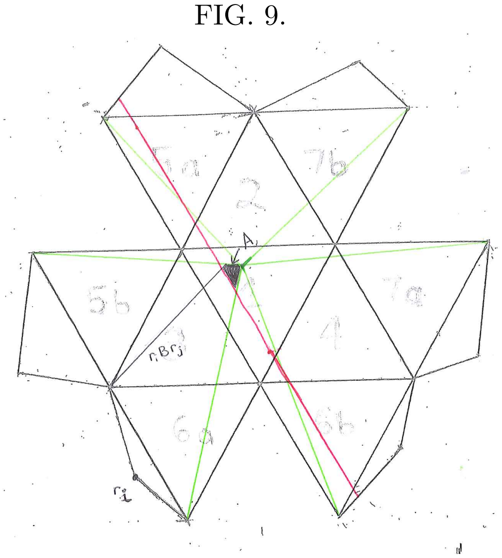

$\square$

Lemma 3.6. *Take a voronoi region or a point-representation $r_i \in f_{U_p}^{-1}(p)$. Then the voronoi region R_i corresponding to r_i must be partially determined by both r_i's 1st and 2nd category boundary lines. That is, there are lines L_1 and L_2 leaving r_i whose first intersections with a boundary line of r_i is with a 1st and 2nd category boundary line, respectively.*

Proof. This follows directly from the fact that the 1st and second category boundary lines of r_i both start at an inner vertex of r_i's secondary face, from which they continue inside the triangle formed by faces 1, 2, 3, and 4. Note that the 3rd category boundary lines corresponding to r_i start at the outer vertices of the secondary faces on either side of r_i and must stay contained within these faces. Hence the vertices from which the 1st and 2nd category lines start must always be in the border of the voronoi region of r_i. □

Lemma 3.7. *Consider a point representation r_i. If the voronoi region R_i that corresponds to r_i does not contain y_o, then both of the voronoi regions either side of R_i must.*

FIG. 10.

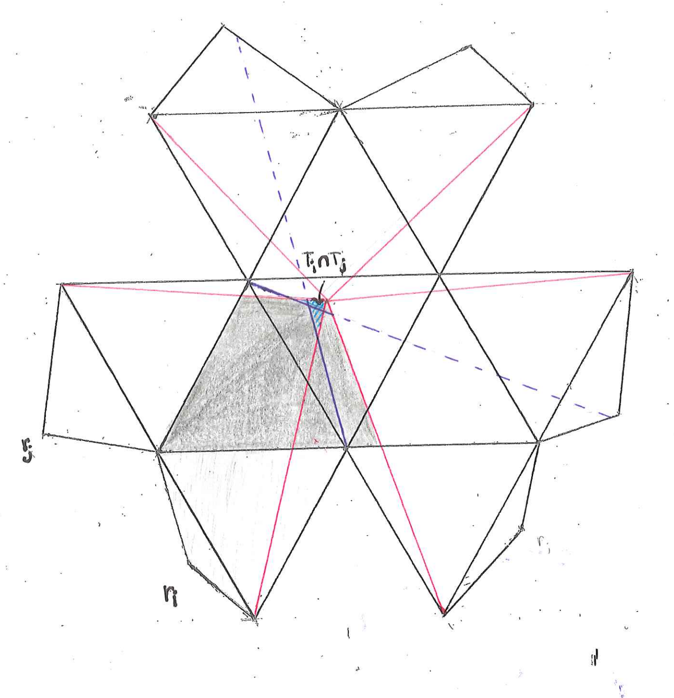

Proof. Suppose that R_i does not contain the 3rd order point. Note that R_i must be a subset of the region determined by the two 3rd category boundary lines corresponding to r_i (see figure 10 above). Thus if R_i does not contain y_o there must be a 1st or 2nd category boundary line of r_i that cuts across the 3rd order boundary lines. This boundary line, together with the two 3rd order boundary lines, forms a triangle T_i that is not contained in R_i.

Now suppose that one of the boundary regions bordering R_i, say R_j, also does not contain the 3rd order point. Then there is a triangle T_j formed by the two 3rd category boundary lines of r_j and another boundary line, such that $T_j \cap R_j = \varnothing$.

We see that $T_i \cap T_j \neq \varnothing$. Hence there is a point $z \in T_i \cup T_j$ such that $y \notin R_i \cup R_j$. However, since each boundary region is a subset of the region formed by its two 3rd category boundary lines, there is no other boundary region that can contain z. Thus we have produced a point that is contained in none of the voronoi regions, which we showed was imposible. ☐

Lemma 3.8. *Take a voronoi region or a point-representation $r_i \in f_{U_p}^{-1}(p)$. If the 1st and 2nd category boundary lines corresponding to r_i intersect within the region determined by r_i's two 3rd category boundary lines then they must intersect on the 3rd category boundary lines corresponding to the point-representations on either side of r_i.*

FIG. 11.

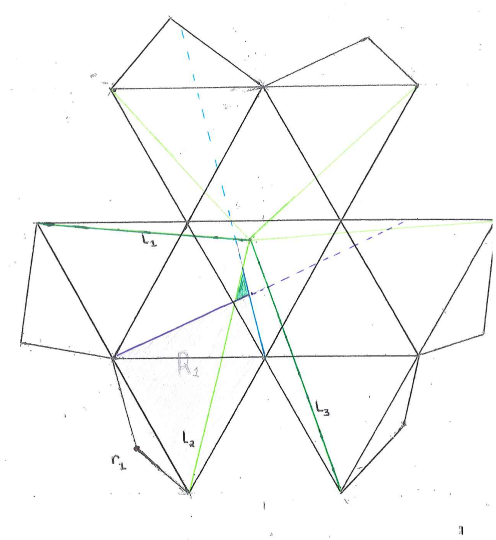

Proof. For this argument it will be helpful to refer to Figure 11 above, on which an example is given. On this diagram the three 3rd order boundary lines in question are labeled L_1, L_2, and L_3. Suppose by way of contradiction that the two second order lines intersect within the region formed by the two 3rd order boundary lines, but not on the third one, L_2. Now, the 1st category boundary line starting nearest L_1 cannot cross L_1 and the 2nd order boundary line starting nearest L_3, for otherwise we would have two neighboring boundary regions that do not contain the 3rd order point. Thus both of the 2nd order boundary lines must cross L_2. Consider the triangle formed by L_2, the two 2nd order boundary lines, and their point of intersection (in green). This triangle is neither in r_i's voronoi region nor the voronoi region on whose side of L_2 it lies. Hence there is no boundary region which contains this triangle, which we have shown is not possible. □

Lemma 3.9. *Consider a point-representation $r_i \in f_{U_p}^{-1}(p)$. Suppose the 1st order boundary line corresponding to r_i passes out of the region formed by r_i's two 3rd category lines by intersecting one of these lines, r_iBr_j, at a point other than y_o. Then r_i's 2nd category boundary line cannot do the same.*

Proof. Suppose the 1st category boundary line, r_iBr_j, leaves the region by intersecting a 3rd category boundary line r_iBr_m, and the 2nd category boundary line does the same. We break this into two cases, both of which are illustrated on the next page.

1. r_iBr_m is the third category line nearer to r_j. Since we know from Lemma 3.8 that r_iBr_j and the 2nd category line cannot intersect within the region formed by the 3rd category boundary lines, this means that the second category line must pass closer to y_o than r_iBr_j. Consider the triangle T_i formed by the 2nd category boundary line, r_iBr_m, and the 3rd category line corresponding to the point-representations on either side of r_i. Neither of the voronoi regions on either side of r_i can contain T_i, so T_i is not in the union of the voronoi regions, a contradiction.

2. r_iBr_m is the third category line further from r_j. In this case Lemma 3.8 tells us that this means r_iBr_j must pass through region deermined by the 3rd category lines closer to y_o than the second category boundary line. Consider the triangle T_j formed by r_iBr_j, r_iBr_m and the 3rd category line corresponding to the point-representations on either side of r_i. By the same reasoning as above T_j cannot be contained in any of the voronoi regions, a contradiction.

□

FIG. 12.

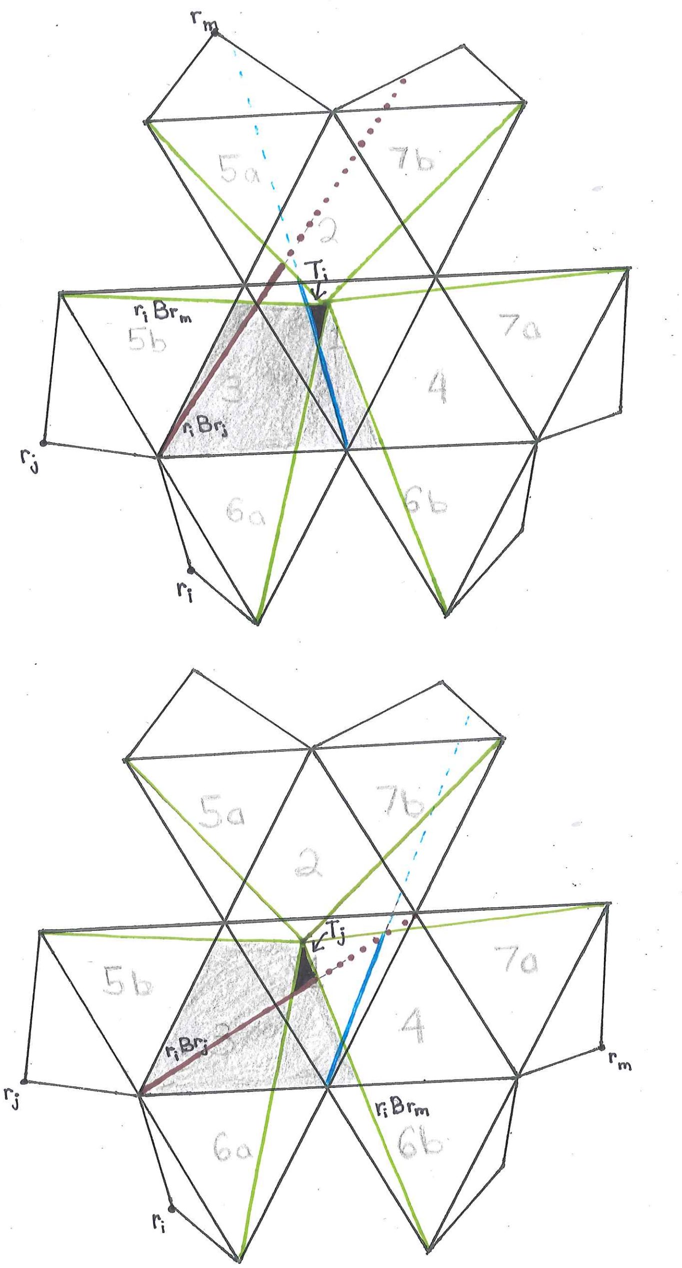

Lemma 3.10. *Take a point representation $r_i \in f_{U_p}^{-1}(p)$. Let r_iBr_j be a 1st, 2nd, or 3rd category boundary line. Then for any point $q \in r_iBr_j \cap U_o$, as you move along r_iBr_j towards the point-representation two away from r_i your distance from r_i increases.*

Proof. Note that the point closest to r_i on r_iBr_j will lie directly between r_i and r_j. As you move away from this point your distance from r_i will increase. Thus to prove the claim it will be sufficient to show that if r_j is next to r_i or one away from r_i then the midpoint between r_i and r_j must lie outside of U_o. (We consider only U_o rather than U because we already know that the furthest point from p cannot lie on a secondary or starting face.) This is obviously the case when r_j is next to r_i. To see that this holds for the case where r_j is one away from r_j note that in this case the midpoint is as close as it can be to U_o when p is a vertex of O, and the midpoint lies on the edge of U_o (which still means our distance from r_i will increase as we move away along r_iBr_j). □

Theorem 3.11. *I now finally show that $F_x = \{f_{U_p}(y_o)\}$.*

Proof. I will show that for any $z \in L_{f_{U_p}^{-1}(p)}$, $D(z) \leq D(y_o)$. Applying Lemma 3.3 and the fact that $f_{U_p}^{-1}(F_p) \in L_{f^{-1}(p)}$, this means that $F_x = \{f_{U_p}(y_o)\}$.

Let $z \in L_{f_{U_p}^{-1}(p)}$. Let r_i be a point representation in whose vonoroi region z lies. Then we know from Lemma 3.7 that z lies on a 1st, 2nd, or 3rd category boundary line corresponding to r_i. Call this boundary line r_iBr_j (choosing one line if z is at an intersection). I examine 4 cases:

- **Case 1:** z lies on a 3rd category line corresponding to r_i:

  Since part of the border of r_i's voronoi region is determined by its 3rd category boundary line, claims (re intersection of 1st and 2nd) tell us that this voronoi region must contain y_o. Hence we can follow r_iBr_j to y_o. Lemma 3.10 tells us our distance from r_i will increase constantly along this path, and so $D(y_o) > D(z)$.

- **Case 2:** r_iBr_j is a 1st or 2nd category boundary line which continues through y_o.

  In this case follow r_iBr_j to y_o. Lemma 3.10 again tells us that the distance from r_i increases along this path, so $D(y_o) > D(z)$.

- **Case 3:** r_iBr_j is a 1st or 2nd category boundary line which does not intersect one of r_i's 3rd category lines.

  In this case r_i's 1st and 2nd category lines must intersect within the region determined by its 3rd category lines. Lemma 3.8 tells us that these lines must then intersect on the 3rd category boundary line which corresponds to the two point-representations next to r_i. Call them r_j and r_k. We can continue from z along r_iBr_j until we intersect r_kBr_j at a point z_o. Our distance from r_i will increase along this path and so $D(z_o) > D(z)$.

When we reach z_o there will be no direction within r_i's voronoi region that will increase our distance from r_i. However, z_o lies in the voronoi region of r_j, on a 3rd category line of r_j. Thus we can apply the arguments from Case 1 to z_o, travelling a path of increasing distance until we reach y_o and $D(y_o) > D(z_o) > D(z)$.

- **Case 4:** r_iBr_j is a 1st or 2nd category boundary line which does intersect one of r_i's 3rd category lines.

  In this case we can follow r_iBr_j away from z until we intersect the 3rd category line at a point z_o. We can then follow this boundary line to y_o. Lemma 3.10 again tells us that $D(y_o) > D(z_o) > D(z)$.

The only case not examined above is the case where $z = y_o$, which is trivial. The combination of these cases show that for any $z \in L_{f_{U_p}^{-1}(p)}$, $D(y_o) \geq D(z)$. Furthermore, our arguments also showed that $D(y_o) = D(z)$ only when $y_o = z$. Hence $y_o = D^{-1}(max(D(U_o)))$ and so $F_p = \{f_{U_p}(y_o)\}$. □

Theorem 3.12. *$f_{U_p}(y_o)$ is the antipode of p.*

Proof. To see this let us return to the 3rd category boundary lines. Diagrams are provided on the next page. Recall how we determined the 3rd category lines on U_p: on U_o we drew lines from $f_{U_o}^{-1}(p)$ to the outer vertices of faces 2, 3, and 4. Then, with U_p oriented upside-down from U_o, we drew corresponding lines on U_p starting at the outer vertices of the secondary faces, so that the angle a given line made with the edges of its face was flipped from the corresponding lines in U_o. For example, from one of the representations of face 3 on U_o we have a line going towards $f_{U_o}^{-1}(p)$ at an angle of a degrees. This corresponds to a line leaving the vertex of face 6a in U_p at an angle of $\frac{\pi}{3} - a$. We also showed that the 3rd category boundary lines intersect at y_o.

Consider U_p with the 3rd category lines drawn on, and choose two 3rd category lines that start at opposing vertices. Consider the triangle formed by these two lines and the line that joins their starting vertices. There are three such triangles on U_p and they determine the location of y_o. Furthermore, we can see these same triangles on U_o, where they determine the location of $f_{U_p}^{-1}(p)$. Observe that the position of these triangles of face 8 is exactly opposite their position on face 1. Hence y_o is the antipode of $f_{U_p}^{-1}(p)$.

□

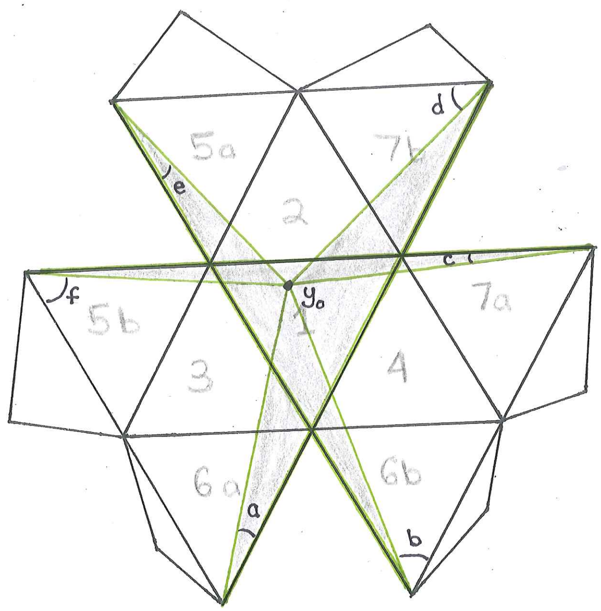

5a
7b
d
e
2
c
f
y_0
7a
5b
1
3
4
6a
6b
a
b

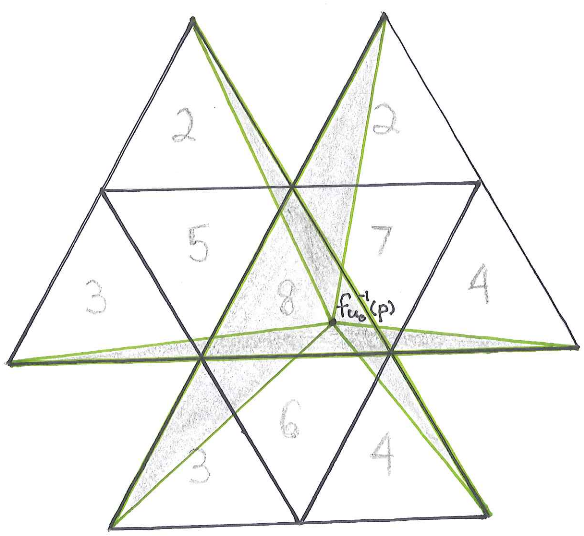

2
2
5
7
3
8
$f_{u_0}^{-1}(p)$
4
6
3
4

Section 4

In this section I use the fact that $\{F_p\}$ is exactly the antipode of p to calculate the distance from p to its furthest point, finally maximizing this as p moves over a face. For any point $p \in O$ let p' denote the antipode of p. Clearly p' must lie on the face opposite the starting face of p. Let us consider the possible face-paths that we can take to get from p to its opposite face. The minimum number of faces that we can use is four. If we travel over four faces and do not pass through the opposite face we will necessarily have to pass through two secondary faces (or the same secondary face twice). We know from the previous section that if any line segment passes through two or more secondary faces, there is a shorter line segment between the same two points. Thus to take the shortest path from p to p' we must take a four-face path.

FIG. 13.

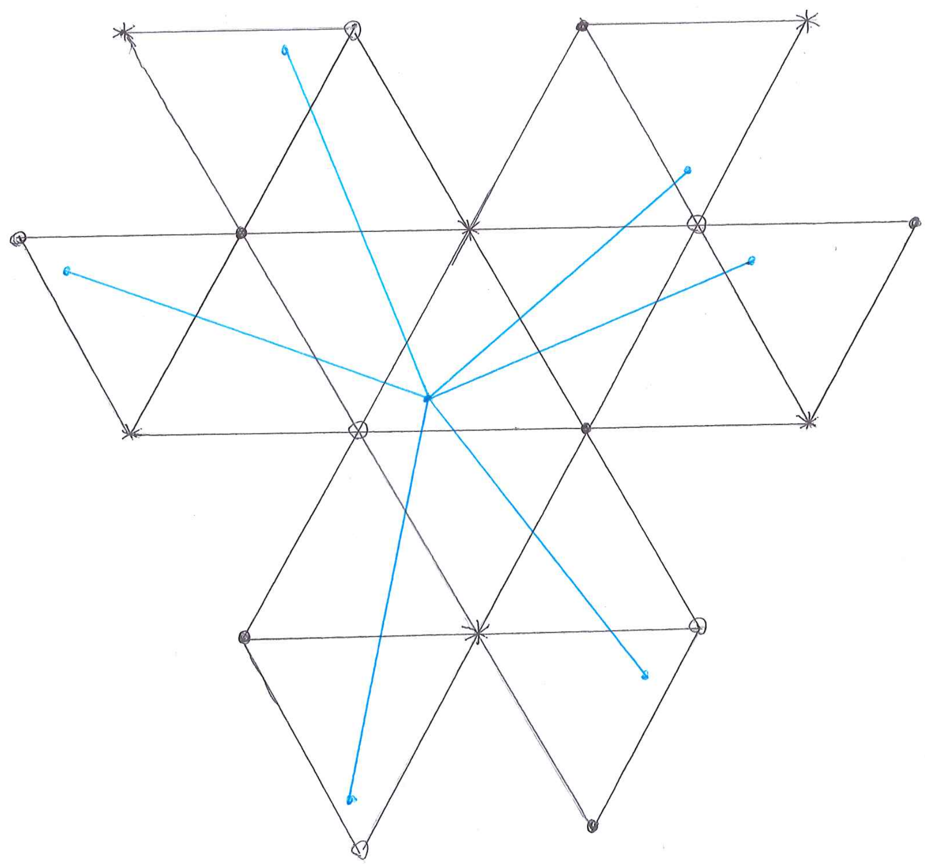

Let us then consider all possible four-face paths from the starting face to the face opposite it. There are six such paths, shown on figure 13 above. Each of these paths corresponds to a different line segment from x to x'. However, we can see that there is a symmetry such that half of the line segments are essentially equivalent to the other half. To see this, take a four face path. Flip the four faces over and translate them up until they lie ontop of the four-face path across from its original location. The two paths are identical. Thus we need only consider three paths from x to x', shown on Figure 14.

Denote the lengths of the three paths by A, B, and C. If we place the three face-paths on a coordinate grid with the lower left vertex the starting face at the origin, we can use trigonometry and simple coordinate distances to determine the length of each path.

FIG. 14.

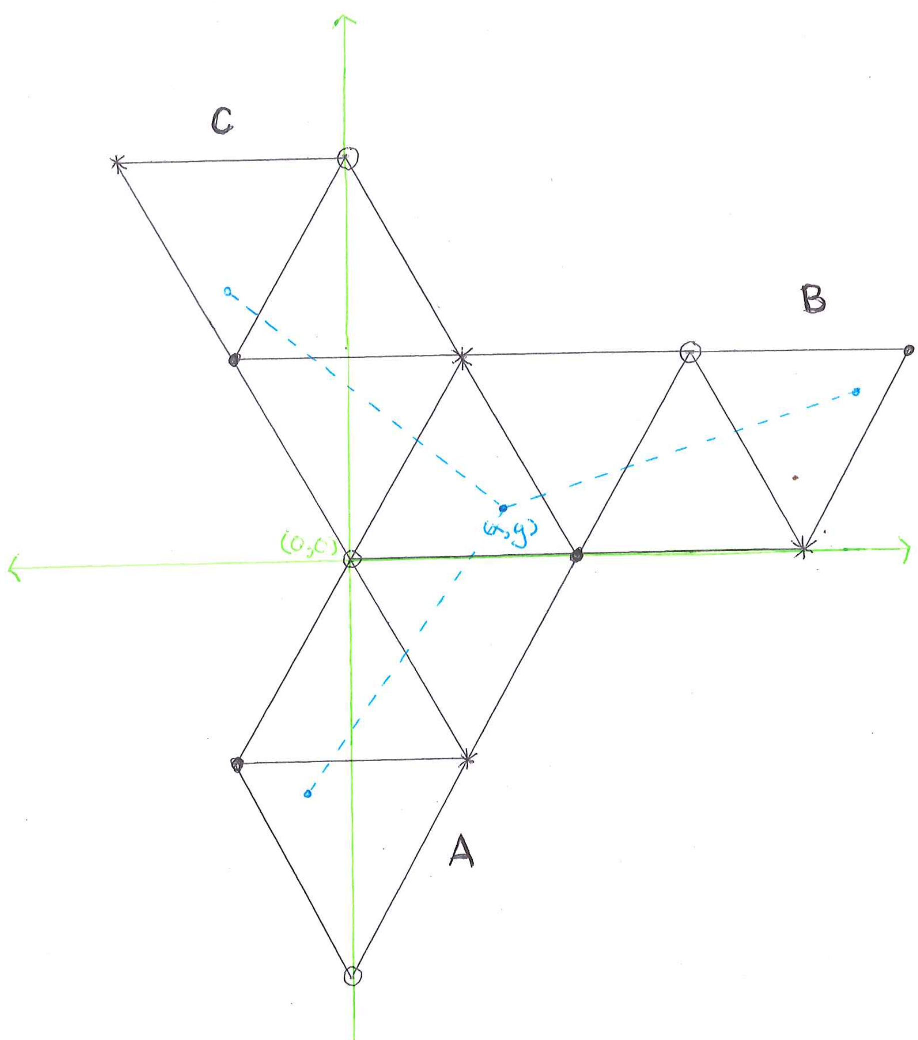

If we set the edge of a face to be equal to two units on a coordinate grid, we find the following relationships. When the coordinates of a point p on the starting face are (x, y)

then the coordinates of p' are:

- **Along path A:** $(\frac{y\sqrt{3}-x}{2}, -2\sqrt{3} + \frac{x\sqrt{3}+y}{2})$
- **Along path B:** $(3+x, \sqrt{3}-y)$
- **Along path C:** $(\frac{-y\sqrt{3}-x}{2}, 2\sqrt{3} + \frac{-x\sqrt{3}+y}{2})$

Calculating the distance from $p = (x, y)$ to each of these points, we get the following lengths of the three different paths:

- **Length of A:** $3x^2 - 2\sqrt{3}xy - 6x + y^2 + 2\sqrt{3}y + 12$
- **Length of B:** $4y^2 - 4\sqrt{3}y + 12$
- **Length of C:** $3x^2 + 2\sqrt{3}xy - 6x + y^2 - 2\sqrt{3}y + 12$

The distance from p to p' is the minimum of these lengths. Note that while the specific locations and path lengths depend on the scale we chose, the relative lengths of the paths do not. To determine where on the starting face paths A, B, and C are minimal we set their lengths equal to eachother.

- The lengths of A and C are equal when $x = 1$.
- The lengths of A and B are equal when $y = \frac{x}{\sqrt{3}}$.
- The lengths of C and B are equal when $y = \frac{2-x}{\sqrt{3}}$

Note that these are exactly the perpendicular bisectors of the starting face. Using these inequalities we can see (and the reader can check) that A minimizes when p is in the lower right third of the starting face, B minimizes when p is in the top third of the starting face, and C minimizes when p is in the lower left third of the starting face. Consider the functions f_P that takes a point p on the starting face to the length of path P (where P is one of A, B, or C) from p to p'. Define $f(p)$ to be the minimum of $\{f_A(p), f_B(p), f_C(p)\}$ (diagram ???). Then f gives the distance of a point p on a face to its antipode.

Given that we can rotate the entire diagram to turn one path into the position of another, we know that f_A, f_B, and f_C act the same on the thirds of the starting face on which they minimize. Hence we need only consider one of these functions, restricted to the third on which it minimizes, to determine which points p are furthest from their antipodes. It is clear that f_B reaches its maximum value at the vertex of the starting face. Thus the distance from p to its antipode is maximized when p is a vertes, and so the hermit points of a regular octahedron are exactly all the vertices of this octahedron.

[1] I. Adlestein *Closed Geodesics via Critical Points of Energy*, Doctoral Dissertation in progress, Dartmouth College, 2014.

[2] M. DoCarmo, *Riemannian Geometry* (transl. by Francis Flaherty), Math. Theory Appl., Birkhauser Boston Inc., Boston, MA, 1992.

[3] P. Doyle,

[4] M. Gromov *Metric Structures for Riemannian and Non-Riemannian Spaces* (transl. by Sean Bates) Progress in Math., Birkhauser Boston Inc., Boston, MA, 1999.

[5] R. Hess, C. Grinstead, M. Grinstead, D. Bergstrand, *Hermit Points on a Box*, The College Math. Journal (2012), 12-23.

[6] U. Lang, *Length Spaces*, Lecture Notes for a Course on Metric Geometry, ETH Zurich, 2004.

[7] Costin Vilcu, *On Two Conjectures of Steinhaus*, Geometraie Dedicata **79** (2000), 267 - 275.