Voronoi Methods for Spatial Selection

Alan Saalfeld

ABSTRACT: We define measures of "being evenly distributed" for any finite set of points on a sphere and show how to choose point subsets of arbitrary fixed size that are as evenly distributed as possible.

KEYWORDS: sampling, evenly-distributed points, point densities, separation, Voronoi regions

Introduction

There are many situations in which one is required to select some fixed-size, spatially representative subset from a given collection of point features. One particular example, the problem which motivated our research, is the following: we wanted to process and analyze data from the approximately 3000 fixed-location continuously operating GPS reference stations throughout the world, but our hardware and software could only handle processing and analysis for about 100 stations. We were looking for "global coverage" results, so we wanted to pick a representative hundred stations on which to do our processing. Intuitively, "representative" meant to us that (1) they spanned the whole region of interest, and that (2) neighboring stations were neither too close nor too far apart, (a property we called "evenly-distributed").

In this paper, we will develop the mathematical principles (definitions, theorems, and their proofs) and the tools (algorithms and data structures) for selecting evenly-distributed spanning point subsets of a pre-specified size. We have to choose from among a given set of points. We use the area of a point's Voronoi region as a measure of the point's separation from its neighbors; and we demonstrate how this measure allows us to adjust and control local point densities of our subsamples efficiently.

Mathematical Preliminaries

Points on a Sphere

Since we think of a sphere as a surface in 3-space or R^3, any finite collection of points on a sphere may also be regarded as vertices of a convex polyhedron, namely the convex hull of the finite collection of points in R^3. The convex hull of the points has edges and facets. If the points are in general position on the sphere, the facets are triangles. Every hull facet is a convex polygon that lies in the plane containing all of the facet's vertices; and that plane is such that all of the other vertices of the other facets of the polyhedron lie on the same side of the plane in R^3. An edge is a straight line segment that is the intersection of two facets; and the two vertices belonging to the edge are called neighbors.

How might we define the notion of "evenly-distributed points on a sphere?" Our first ideas come from the vertices of regular polyhedra (also known as platonic solids) inscribed in a sphere (see Fig. 1). The three conditions that we are looking for in our set of points are:

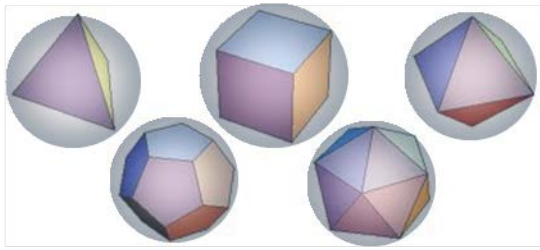

Figure 1. The vertices of platonic solids inscribed in a sphere illustrate "even distributions" of points on a sphere

1. They are not too close together.

2. They are not too far apart.

3. Every other point on the sphere is not too far from one of our chosen set of points.

We call measures of the first property *separation measures*; measures of the second property we call *net-mesh measures*; and measures of the third property we call *coverage measures*.

The symmetry of the regular 3-D figures guarantees that in each of the five regular polyhedra, the edges' lengths in that polyhedron are all identical; hence, all points are equally close to all of their neighboring points. Moreover, on any regular polyhedron, all points have the same number of neighbors Unfortunately, only regular polyhedra have the property that all neighboring pairs of points are the same distance apart; and there are only five such polyhedra, the largest of which has 20 vertices. Hence, since we want to consider polyhedra with more than 20 vertices, we will need to examine situations in which interpoint distances between neighbors vary. How can we describe and classify that variation? What constitutes good bounds on that variation, bounds that for a given size point set give us a reasonably even spacing of points? Let's define some measures of good spacing.

Suppose we are given a set $\mathbf{X}$ consisting of a large number of points on a sphere. We have the following notions of $\mathbf{X}$ being sufficiently separated while at the same time being sufficiently dense everywhere:

For any positive real number $\boldsymbol{\gamma}$, we will say that $\mathbf{X}$ has *gamma-separation* (or $\boldsymbol{\gamma}$-separation) if no two points of $\mathbf{X}$ are closer than $\boldsymbol{\gamma}$, where distance is measured along the minor arc of the great circle passing through the two points.

For any positive real number $\boldsymbol{\delta}$, we will say that $\mathbf{X}$ has a *delta-net* (or $\boldsymbol{\delta}$-net) on the sphere if there is a triangulation of $\mathbf{X}$ on the sphere by minor arcs of great circles, no arc-length of which is greater than $\boldsymbol{\delta}$. Note that we could define distance between points in terms of the straight line segment distance in $\mathbf{R}^3$, and though distance values would not be the same, distances defined either way would produce the same topology on the sphere.

Suppose we are told that a set **X** has **γ**-separation and also has a **δ**-net; and, moreover, that $\delta = \gamma$. Then since two distances are the same in one metric (the great circle distance) if and only if they are the same in the other metric (the straight-line-segment distance), the set **X** must form the vertices of a regular polyhedron with triangle facets (tetrahedron, octahedron, or icosahedron); and the chords of the great-circle arcs of the **δ**-net are the edges of the polyhedron. Thus the point set **X** must have size 4, 6, or 12.

For any positive real number **ε**, we will say that **X** forms an *epsilon-cover* (or **ε**-cover) of the sphere if every point on the sphere is within a distance **ε** of some point of **X**. In other words, if we place the center of a disk of radius **ε** at each point of **X**, then those disks will cover the sphere (Technically, our "disks" are curved and shaped like bubbles or contact lenses. They consist of all points on the sphere within a distance **ε** of the center, the distance being the geodesic distance, measured along an arc of a great circle of the sphere.).

If **X** forms an **ε**-cover of our sphere, then it has a **δ**-net for all $\delta > 2\varepsilon$.

Conversely, if **X** has a **δ**-net on the sphere, then it forms an **ε**-cover for all $\varepsilon > 3\delta/4$.

Having a **δ**-net or forming an **ε**-cover can also have meaning on subsets of the sphere. We might say, for example that **X** is an **ε**-cover of a subset **A** of the sphere if every point in **A** is within **ε** of some point in **X**. If **A** is, as in our GPS example, all the land masses of the globe, then permanent GPS stations can be denser (have smaller **δ** or **ε**) than on the oceans.

We could have defined our separation and denseness notions on any metric space, but we choose to work with the sphere because it has the following compactness property: Given any non-empty finite set **X**, there exists a smallest $\varepsilon_o > 0$, such that **X** is an **ε**-cover of the sphere for all $\varepsilon > \varepsilon_o$. We will call that smallest value the *ε-cover infimum* for **X** and write $\mathbf{inf}_\varepsilon(\mathbf{X}) = \varepsilon_o$.

Then $\mathbf{inf}_\varepsilon()$ is monotone decreasing in the following way: If $\mathbf{X} \subseteq \mathbf{Y}$, then $\mathbf{inf}_\varepsilon(\mathbf{Y}) \leq \mathbf{inf}_\varepsilon(\mathbf{X})$.

Suppose we have a set **X** of $\mathbf{n > 3}$ points on the sphere that leaves no hemisphere empty. In other words, every great circle not passing through a point of **X** will separate **X** into two non-empty disjoint sets. Equivalently, we could have required that the center of the sphere lie inside the convex hull of the points.

In this case, the sphere may be fully triangulated by minor-arc-of-great-circle edges connecting pairs of points of **X** . Triples of points of **X** will specify the triangulation's spherical triangles. Each triangle lies completely in some hemisphere. No point of the sphere can be more than $\pi/2$-radians from some point of **X**; hence **X** has an **ε**-cover for $\varepsilon = \pi R/2$, and furthermore $\mathbf{inf}_\varepsilon(\mathbf{X}) \leq \pi \mathbf{R}/2$.

Triangulations of the Sphere

A triangulation of a set of points on the sphere is a maximal embedded graph whose vertices are those points and whose edges are non-intersecting minor arcs of great circles.

Counting Edges, Angles, and Triangles

If we have **n>3** points in general position such that no hemisphere is empty, then any triangulation of those points on the sphere by arcs of great circles always has the same number of edges **(3n - 6)**, the same number of triangles **(2n - 4)**, and the same number of dihedral angles **3(2n - 4)** or **(6n - 12)**. Since the dihedral angles at any point always add up to **360°** or **2π** radians, the sum of all interior dihedral angles for all triangles is **2nπ**. Hence, the average size of an interior angle in radians is **2mπ/(6n-12) = π/3 + 2π/(3n-6)**, or, in degrees, we have **60° + 120°/(n-2)**.

Every triangulation of **X** has a longest edge. The triangulation with the "best" longest edge of length **δ** determines $\boldsymbol{\delta_0 = \inf_\delta(X)}$**,** the smallest possible **δ** in our definition of **δ**-net. How does $\boldsymbol{\delta_0}$ compare with the longest edge length $\boldsymbol{\delta_{DT}}$ in the Delaunay Triangulation of **X**? An analysis of the counterexample the shows the Delaunay Triangulation in the plane is not the minimum total length triangulation there shows that Delaunay Triangulations on the sphere can have an edge that approaches $1.5\cdot\boldsymbol{\delta_0}$.

Every spherical triangle has a circumcircle that is a lesser circle on the sphere; hence every triangulation with **k** triangles has **k** circumcircles associated with it. The following properties can be easily shown to hold for the Delaunay Triangulation of a set of co-spherical points in general position on the sphere:

The sum of circumcircle areas of all triangles in the Delaunay Triangulation is never bigger than the sum of circumcircle areas of all triangles in any other triangulation (i.e., the Delaunay Triangulation is the triangulation that minimizes the total of all circumcircle areas).

The largest circumcircle of any triangle in the Delaunay Triangulation is never bigger than the largest circumcircle of some triangle in any other triangulation (i.e., the Delaunay Triangulation minimizes the maximum circumcircle).

The largest circumcircle of all triangles in the Delaunay Triangulation corresponds to a triangular facet having no obtuse angles.

If **r** is the radius of the largest circumcircle of any triangle in the Delaunay Triangulation and if R is the radius of the sphere, then $\boldsymbol{\inf_\varepsilon(X)}$ is exactly equal to **R·***arcsin***(r/R)**.

This last result ties the coverage of a set of sites to the Delaunay Triangulation of those sites.

Our new algorithms explore building and modifying Voronoi diagrams and Delaunay triangulations on a sphere by inserting and/or deleting point sites from a given fixed set of many point sites one at a time. Finally we explore some mathematical properties of point sets on spheres and their induced Voronoi diagrams, Delaunay triangulations, and inscribed convex polyhedrons that will allow us to quantify measures of balance of our chosen network.

We choose a subset of **m** point sites based on global coverage and separation. To do this we use an iterative Voronoi approach where each site is ranked by the area of its enclosing spherical Voronoi polygon. From among the eligible **n**, the site of minimum rank is discarded. The rankings on (some of) the remaining **n-1** are updated; and this process is repeated with the remaining **n-1**, then with the remaining **n-2**, and so on, until the desired number **m** of stations is reached. This technique removes one site at a time until it reaches **m**.

Now that we have described what works well, let's provide a little analysis of why it works and how it works (and how it might work even better).

Sub-network selection issues

For all of the applications that we will cover in this exposition, the model of the Earth as a sphere will be sufficiently precise. A ***global geodetic network***, therefore, is a triangulation of the globe (sphere) into spherical triangles (triangles whose edges are minor arcs of great circles). The vertices of the triangulation represent the ground station sites or fixed point locations of the network, sites of known location from which distance measurements may be taken.

The objective of one simplified geometric version of the problem might be stated as follows: find a network on a subset of stations that still affords global coverage and that has network edge lengths (baselines) as small or as nearly equal as possible. Both criteria, small or nearly equal, will get rid of some large baselines. The nature of the distribution of the permanent ground stations, however, is so uneven that not all large baselines can be avoided. An isolated station in the middle of the Pacific Ocean, for example, is necessarily far from all other stations, and, hence, far from its neighbor stations, thus producing one or more long baselines in any possible network. Every site in the triangulation has at least 3 baselines associated with it. Thus truly isolated stations contribute at least 3 long baselines. Nevertheless, leaving out an isolated station would necessarily result in some even longer baselines. Thus isolated stations should belong to the network, if possible. The fact that we want to include isolated stations reinforces the strategy of keeping stations with large Voronoi regions; or at least it strengthens the argument for initially throwing away stations with small Voronoi regions. We will look at some other consequences of the different strategies in later sections.

With the goal of selecting few long baselines, there are several similar measures of optimality that one may try to achieve.

Here are two specific strategies that we considered:

An ***incremental deletion algorithm*** to solve this problem would "throw away" one station at a time until only **m** stations remain.

An ***incremental insertion algorithm*** would be an algorithm that adds one station at a time (starting from **0**, or starting from some preselected set of **j** locations, for some **j** < **m**) until the set has size **m**.

Any algorithm to select a set of **m** stations would then need to be followed by a network-building algorithm to add edges to the network so that the resulting network could be assessed for quality (in terms of the lengths of the ``baseline" edges). One simple and reasonably effective way to recover edges once the sites have been selected is to build the spherical Delaunay Triangulation of the site set. Although the Delaunay Triangulation is not necessarily the minimum total length triangulation in the plane or on the sphere, the Delaunay Triangulation does make its component triangles as equiangular as possible; and this fact, at least locally, does tend to make within-triangle-edge-length-ratios closer to one.

Incremental Deletion

We opted for an incremental deletion algorithm that successively discarded the site with the smallest Voronoi region among all remaining sites. After a site is thrown away, we redistributed the site's Voronoi area by giving some to each of its Delaunay neighbors, the equivalent of recomputing the Voronoi regions for the set with one fewer vertex or site. Here is pseudo-code for a simple Voronoi-based sub-optimal incremental deletion algorithm for choosing m well-distributed point sites from n sites:

Input: A set of n points in R^3 at $\{(x_1,y_1,z_1), (x_2,y_2,z_2), (x_3,y_3,z_3),..., (x_n,y_n,z_n)\}$

Output: A subset of size m $\{(x_{i1},y_{i1},z_{i1}), (x_{i2},y_{i2},z_{i2}), (x_{i3},y_{i3},z_{i3}),..., (x_{im},y_{im},z_{im})\}$

```
Let S = {(x1,y1,z1), (x2,y2,z2), (x3,y3,z3),..., (xn,yn,zn)}

Compute the Voronoi Diagram on the sphere for the locations in S and
the areas of the Voronoi regions of the points in S;

    For i = 1 to n - m

        Remove from the set S the location whose Voronoi region has
        smallest area

        Update the Voronoi diagram and Voronoi region areas by
        reallocating the area that belonged to the removed site;
```

Although we know that Delaunay networks are not shortest total length networks in general, why should we choose them if we are trying to avoid long baselines, and what makes us think they might help solve the problem?

Why try Voronoi?

The Voronoi region of a site consists of all points that are closer to that site than to any other site. The Voronoi region is sometimes called the site's *region of influence*. A site with a large Voronoi region has more responsibility for keeping track of activity in a region less sampled by other sites or less accessible to multiple observations.

Insights from the 1-D analogue to the problem

The circle is the 1-D analogue to the sphere, and Voronoi regions on the circle are just simple arcs.

Suppose that we are given some collection of point sites on a circle. Which point should we delete to minimize the new adjacent point distances? The only new distance that we create by a single point removal is the distance between the two points that straddled (were adjacent on either side of) the removed point. So if we remove the point with the shortest straddling distance, we will add the smallest possible new adjacent link on the circle. But, in terms of Voronoi region size (in this case, arc length), we see that this strategy equates to removing the point with the smallest Voronoi region:

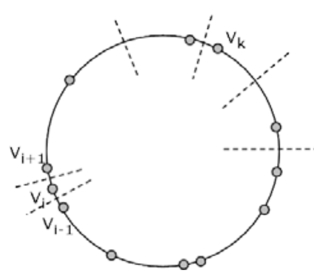

Fig. 2. Points on a circle with some of their Voronoi regions separated by dashed lines.

Lemma. Suppose the points $\{\mathbf{v_1}, \mathbf{v_2}, \ldots, \mathbf{v_k}\}$ all lie in order on a circle. Then the point $\mathbf{v_i}$ has the smallest Voronoi region if and only if the straddling pair $\{\mathbf{v_{\{i-1\}}}, \mathbf{v_{\{i+1\}}}\}$ subtend the smallest arc of any straddling pair $\{\mathbf{v_{\{j-1\}}}, \mathbf{v_{\{j+1\}}}\}$.

Proof: From Figure 2, it is clear that the length of Voronoi region arc for every vertex $\mathbf{v_j}$ equals exactly 1/2 the length of the arc connecting $\mathbf{v_{\{j-1\}}}$ and $\mathbf{v_{\{j+1\}}}$.

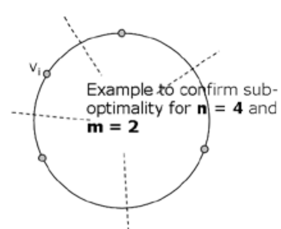

Fig. 3. The optimal solution for m=2 cannot be reached by an incremental deletion algorithm.

Note: This incremental approach can be easily seen in the example in Figure 3 to fail to be globally optimal on the 1-D analogue (i.e., for points on a circle), but the 1-D version of the algorithm can also be shown to be within a small constant multiple of optimality. The only best distribution for m=2 sites consists of v_i and the vertex opposite it. However, the incremental algorithm throws v_i away at the first step.

Persistence of Certain Sites

The strategy of always throwing away first the site with the currently smallest Voronoi region has some very desirable outcomes.

Here is one nice property of the m-subset selected by incremental deletion: If a point starts out with a relatively large Voronoi region, then it is not discarded by the incremental deletion algorithm.

Proposition. If a point site starts out with a Voronoi region bigger than $[1/m]^{th}$ the size of the sphere area, then that point is not removed by this iterative point-removal procedure which stops when m points remain.

Proof: The smallest area of i remaining regions, $m < i \leq n$, is always less than or equal to $[1/i]^{th}$ the size of the sphere, which in turn is less than $[1/m]^{th}$ the size of the total sphere area. When a Voronoi region is removed, its area gets redistributed to (some of) the remaining regions. Hence, after site removal, the area of each of the remaining regions stays the same or increases. Thus, any site that starts out with more than $[1/m]^{th}$ the total sphere area continues to have more than $[1/m]^{th}$ the total sphere area, and will never be chosen as the smallest remaining site.

The persistence result in the proposition guarantees that the incremental deletion algorithm will leave some sites untouched. We might use those *always-in* sites as an initial set for an incremental insertion algorithm. For example, in a specific situation of having to choose, say, 100 sites from 300, there may be thirty or forty of those 300 sites, each of which has more than $[1/100]^{th}$ of the sphere's total area in its Voronoi region. While an insertion strategy (and we have not yet even suggested how one might be implemented) might not give the same 100 sites as our site deletion strategy, the thirty or forty sites that represent large areas should probably be in both sets.

Here are some of our as yet undeveloped ideas about how to design an algorithm that incrementally adds sites. One could try to build a network of sites by incrementally adding the site which claims the largest Voronoi region possible. At issue is how to efficiently determine what that site is and thus implement the algorithm. An algorithm that checks each and every as-yet-un-inserted point at each insertion step would be very inefficient. Perhaps we can find the size (or update the sizes, or bound the sizes) of Voronoi regions resulting from all possible one-point additions to a growing network (Note: adding a point to a set will make the other points' Voronoi regions smaller or the same, but never larger). We have thought to look for a site within the largest still empty circle because Voronoi vertices are centers of circumcircles of Delaunay triangles. The lower bound that we computed for the Voronoi region of such an added point depends on where in the empty circle the site-to-be-added was located (closer to the center was better).

We also noted that an incremental insertion algorithm needs to be *primed*. There should probably be at least 4 well-distributed points discovered in an initialization procedure that is different from the incremental procedure. Until some 4 points have been chosen, we do not have a Delaunay triangulation or a polyhedron to work with.

Once we have a working incremental insertion algorithm, we will be able to build a pair (or any number) of disjoint sub-networks by "choosing up teams" using the algorithm and taking turns in the selection process. These disjoint sub-networks could provide some statistical information about precision of measurements generated by the different independent sub-networks using the same forms of analysis.

This exclusive alternating rule could be modified to allow designated shared sites by starting each team with the same set, then alternating selections of the maximal-Voronoi-claiming site from the two growing sets.

Conclusions

There is still much to learn about building better GPS ground station networks, and many interesting open problems about how to select more equitably distributed subsets from sets of fixed point locations. We have observed that incremental methods do not always build optimal solutions, and often we do not even have tools to verify that something is an optimal solution. We saw via theory and examples that sites with larger Voronoi regions becomes more important to keep, especially since not keeping them leaves a big hole and a longer distance to span to reconnect the network. Selecting sites with larger Voronoi regions has produced some new ideas and insights about network stability. We found, for example, that if a very remote site had two GPS stations, then each would have a large Voronoi region, and thus, by our criteria, both would be selected by our incremental deletion procedure. At first glance, we thought that two stations on one remote island would be redundant, and that choosing just one would have been a better choice. It turned out that having two sites in almost the same isolated place was more useful that we ever imagined because of the way the measurements sometimes reinforced each other and at other times served to focus our attention on discrepancies.

We have added some well-known and some not-so-well-known results in spherical geometry and convex polyhedrons that we think may be useful for computing areas of spherical polygons. One result which we prove at the end of this section allows easy computations of Voronoi region areas from the facet angles of a convex polyhedron on those same site points.

Spherical Excess

The average angle in a spherical triangle is greater than 60°; and, in fact, the amount by which the triangle's angles exceed 180° (or, more correctly, exceed π radians) is equal to the triangle's area in radians squared. This so-called spherical excess result is also true for spherical polygons in general: a plane polygon on n vertices always has (n-2)180° as the total of its interior angles. The amount by which a spherical polygon's interior dihedral angles exceed (n-2)180° or $(n-2)\pi$ is a measure in radians squared of the area of the polygon.

Delaunay/Voronoi/Polyhedron Theory

Every general-position set of site points on the sphere corresponds to (1) an associated spherical Delaunay Triangulation (whose edges consist of great-circle arcs), (2) an associated spherical Voronoi diagram (whose region boundaries are made up of great-circle arcs, and (3) a convex polyhedron (whose edges are straight-line-segments and chords of the sphere), as illustrated in Figure 4. There is further a one-to-one correspondence among the edge sets that preserves adjacency.

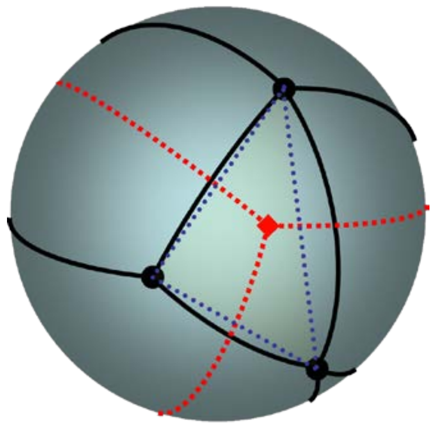

Fig. 4. The Delauney great-circle arc edges (black), the Voronoi great-circle arc edges (dashed red), and the polyhedron straight-line-segment edges (dotted blue), all for the same (black) point sites.

Delaunay Triangulations

There is a nice relationship between Delaunay Triangulations in the plane and Delaunay Triangulations on a sphere that permits computation to be done in whichever domain is preferred.

Lemma. The Delaunay triangulation of points on the sphere is combinatorially equivalent to the Delaunay triangulation of any stereographic projection of the points to the plane (except for edges that would lie outside the convex hull of the plane projected point set).

Proof: The stereographic projection and its inverse preserve all circles and all of their containment relations on the two surfaces. Circles free of vertices in their interiors are precisely the constraints that are preserved and also that define Delaunay triangles/triangulations.

Lemma. The Delaunay triangulation of vertices on a sphere is combinatorially equivalent to the convex hull of the polyhedron on the same set of vertices in 3D; and the great circle arcs that make up the Delaunay edges on the sphere have their chords as the corresponding edges of the polyhedron under the combinatorial equivalence.

The following lemma shows how some checks and computations can be done using 3D Cartesian coordinates and dot products and cross products:

Lemma. Suppose we build the Delaunay Triangulation of a set of **n** points (vertices) on the sphere. Then the following are equivalent:

1. A spherical triangle $\mathbf{\Delta v_1 v_2 v_3}$ belongs to the Delaunay triangulation of vertices on a sphere.

2. A triangle facet $\mathbf{\Delta v_1 v_2 v_3}$ belongs to the convex hull of the polyhedron in 3D generated by the vertices on a sphere.

3. All of the other (n-3) vertices lie on one side of the plane of $\mathbf{\{v_1,v_2,v_3\}}$.

4. The dot product of any vertex with μ, the outward normal to the plane of $\mathbf{\{v_1,v_2,v_3\}}$, attains a maximum value at $\mathbf{v_1 = v_2 = v_3}$.

5. If $\mathbf{x = (v_1-v_2)\times(v_3-v_2)}$, then for any vertex **w** other than $\mathbf{\{v_1,v_2,v_3\}}$, and for i = 1, 2, 3, the **3(n-3)** real numbers $\mathbf{\{(v_i-w)\cdot x\}}$ all have the same sign. If $\mathbf{\{v_1,v_2,v_3\}}$ are in clockwise order, then that sign is positive; and if they are in counter-clockwise order, then that sign is negative.

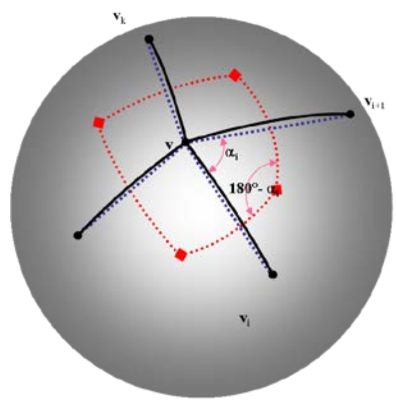

Fig. 5. The spherical excess of v's Voronoi region (dashed red) equals the angle defect at the polyhedron vertex v

Voronoi/Polyhedron Correspondence

Lemma. On the sphere, the arc-of-great-circle edges of the Delaunay Triangulation that form a co-cycle around a vertex **v** are in 1-to-1 correspondence with the arc-of-great-circle perpendicular bisector edges (shown in dashed red) that form a cycle of Voronoi edges bounding the Voronoi cell associated with the site v. Moreover, the angle between the chords $\mathbf{vv_i}$ and $\mathbf{vv_{\{i+1\}}}$ (shown as dotted blue lines in Figure 5) is the supplement (180° - α) of the dihedral angle α of the arcs of the Voronoi edge perpendicular bisectors of the Delaunay arcs $\mathbf{vv_i}$ and $\mathbf{vv_{\{i+1\}}}$.

Proof: Every Voronoi edge corresponds to exactly one Delaunay edge and thus to one chord of the convex polyhedron that shares endpoints with the Delaunay edge. Suppose that two consecutive chords of the polyhedron meet at angle α. The planes of the two Voronoi edge arcs are each perpendicular to their corresponding polyhedron edges because as perpendicular bisectors to the Delaunay arcs, they are also perpendicular bisectors of the chords of those arcs.

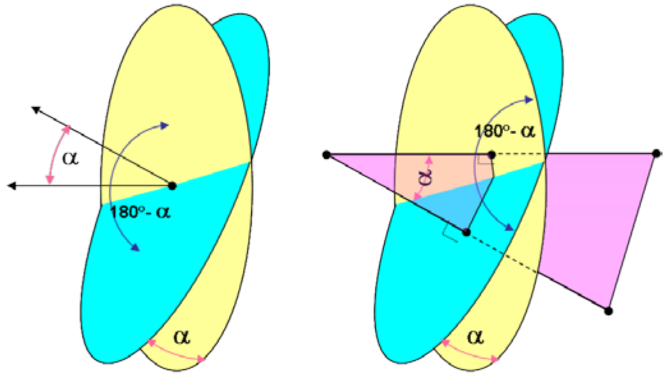

Fig. 6. The polyhedron facet angle and the corresponding Voronoi dihedral angle are supplementary.

As can be seen in Figure 6, the dihedral angle of the Voronoi edge intersection must equal (180° - α).

That lemma takes us easily to a rather surprising and useful theorem:

Angle defect/spherical excess Theorem. The angle defect at any vertex of the convex polyhedron on the set of site points on a sphere is equal to the spherical excess (and hence, the area) of the Voronoi region of the vertex regarded as a site in the Voronoi diagram on the sphere.

Proof: If $\{\alpha_1, \alpha_2, ..., \alpha_s\}$ are the polyhedron angles at the site in question, then the angle defect is $360° - \sum_{\{i=1\}}^{s} \alpha_i$. The spherical excess is the amount that the sum of the s dihedral angles of the Voronoi region exceed **(s-2)180°**:

$$\sum_{\{i=1\}}^{s} (180° - \alpha_i) - (s-2)180° = [s - (s-2)]180° - \sum_{\{i=1\}}^{s} \alpha_i.$$

This last result also gives the Voronoi region's area, also equal to the polyhedron's angle defect in radians squared. So for any set of points on the sphere, we can compute the area of each point's Voronoi region on the sphere without building the Voronoi Diagram or the Delauany Triangulation, but merely by measuring angles on the convex polyhedron of the point set.