UNIVERSITY OF CALGARY

Equal Area Spherical Subdivision

by

Erika E. Harrison

A THESIS

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE

DEGREE OF MASTERS OF SCIENCE

DEPARTMENT OF COMPUTER SCIENCE

CALGARY, ALBERTA

January, 2012

UNIVERSITY OF
CALGARY

Abstract

Numerous subdivision techqniues currently aim to attain smoothness of shape. Some extend this to achieve interpolating conditions, while others aim to additionally preserve volume. Perimeters, for closed curves, or surface area, for closed surfaces, has had less exploration, yet it is of great importance to geographics and manufacturing industries. Particularly with technological advances, and the ability for increased computing power to store and represent the vast data of the Earth's surface, a transition from traditional visualization and analysis approaches - such as the equal area projection described by Snyder - would greatly benefit the cartographic community. An improved integration of visualization with data storage will increase the data that can be analyzed at any given time.

This work, after an initial investigation of traditional projections, and how they may be be improved for visualization purposes, explores the merging of traditional subdivision surfaces with equal area representations. Global as well as local preservations are initially explored on a curve, and then extended to surfaces. Given the requirement of the cartographic community, the problem is simplified to curves that converge to circles, and surfaces that converge to spheres. A limited discussion on extending the approaches to arbitrary shapes is also presented.

Of particular interest, a collection of polyhedral subdivision approaches are tested against the presented equal area subdivision. The results illustrate a variety of issues inherent within unaltered traditional subdivision approaches, and preferential subdivisions. Statistical analysis illustrates a technique which minimizes global area error as compared with the original surface area, or a maximal face distortion of 16%. Future explorations include the extension to arbitrary curves and surfaces, as well as an association of the subdivision surface its regional data.

Acknowledgements

While the bulk of this work has been constructed by the aforementioned author, it would not be accomplished without a host of resources and extremely helpful individuals. The University of Calgary, the Faculty of Graduate Studies and the Department of Computer Science has been greatly assistive through their sharing of resources. Combined, they have provided financial support through teaching assistanceships, logistical resources through office space and a dedicated workstation, as well as an extensive library of references and helpful librarians to ensure a well researched body of work. I wish to thank the Alberta Graduate Scholarships Programme, the NSERC Industrial Grants and our industrial collaborator PYXIS Innovation for providing financial and inspirational support.

I wish to thank my colleagues within the computer graphics lab for providing inspiration and motivation throughout the research. Of most assistance, my fantastic, and supportive supervisor Dr. Faramarz Samavati, has been invaluable, providing supplementary financial support, facilitation of inspirational discussions, assisting in the direction of the research project, and providing continued motivation. I am greatly thankful for his involvement in the work. Of additional note, I wish to thank Ali Madhavi-Amiri for his extremely helpful discussions and technical support, as well as Shannon Halbert for her motivational discussions, and overall review of this body of work.

At a personal level, I wish to thank my family for their emotional support, and continually supportive attitude towards my scholastic endeavours. Friends from my undergraduate studies, particularly Jenny, have been inspirational in the completion of my goals within this program. Lastly, I wish to thank Deb, Chris, Nathaniel, Hamish, Pat, Lidia, Atlas and Ian for being so understanding during my joys and tribulations throughout this research. Their continuing encouragement has helped in the accomplishment of this body of work, and my appreciation for them is insurmountable.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

The Digital Earth framework, as proposed by former US Vice President, Al Gore [25], aims to improve how the Earth's data is portrayed and analyzed. The framework's objective is to efficiently integrate data, and represent it a 3D spherical representation of the Earth's surface. Consequently, planar satellite images, field surveys and other such two dimensional data must be associated with accurate spherical coordinates. Traditional features such as geographic boundaries, bodies of water, and transportation networks are documented to planar form by cartographers to facilitate discussion of such positionally-based information. Traditionally, the association of planar and spherical data has occurred through the employment of mappings or projections. These projections come in a wide range of transformations, having been developed over thousands of years [49].

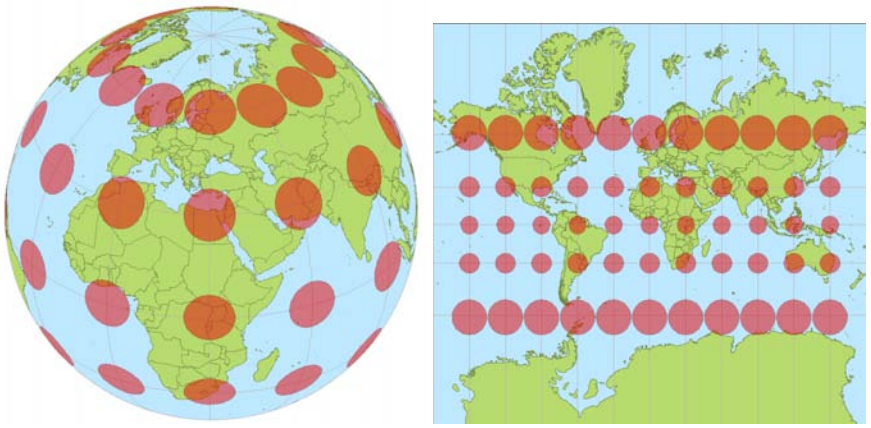

(a) Tissot on World (b) Tissot on Mercator

Figure 1.1: Projection Distortion - Tissot's Indicatrix [53] (Red Circles) are distorted, as they are projected from the world (left) to the map (right) [1]

One of the challenges faced by such projections is the inherent problem of transferring spherical data to a planar form. Figure 1.1, for example, illustrates how the Mercator projec-

tion is distorted from the original Earth's surface. It is impossible to define a projection that preserves both the angles and area of a feature on the Earth [10]. As such, projections are defined and selected according to the task required of them. Desirable characteristics include the preservation of area, preservation of shape, distance and positional accuracy, and an ease of computation [47]. With the advent of computers, this last characteristic - ease of computation - has become less critical to cartographers. Yet within a Digital Earth framework, requiring real-time analysis, computational complexity remains of high importance.

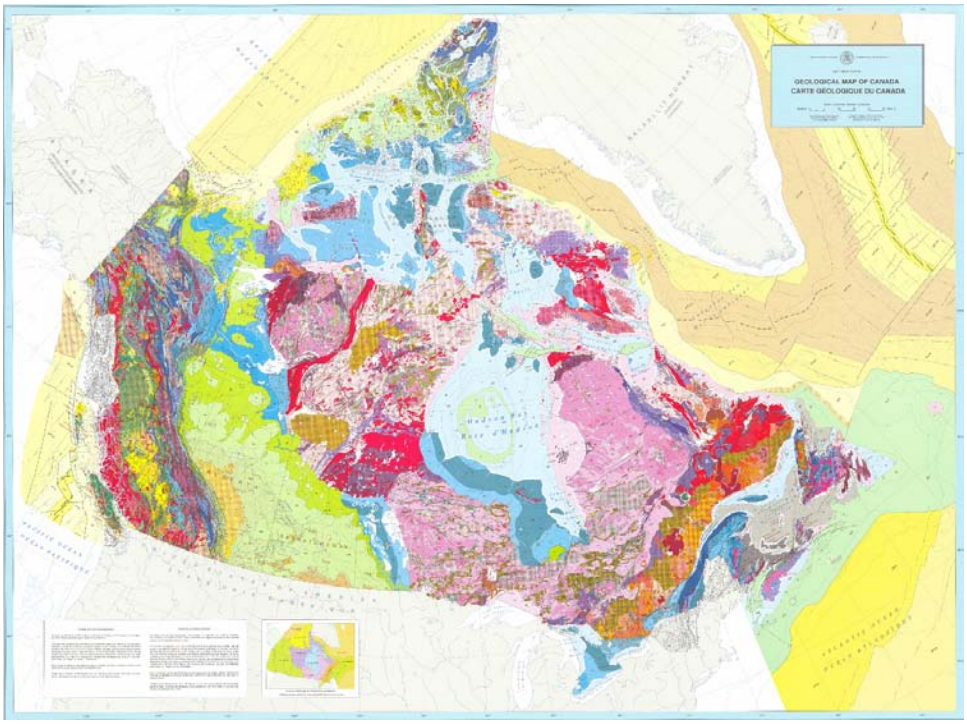

Figure 1.2: Geological Map of Canada (Gov. of Canada, 1996 [61])

The ability to maintain areal equivalence when working between planar and spherical representations of the Earth's surface is imperative for scientific researchers requiring areal or regional analysis. Research topics ranging from regional data relating to biological diversity studies, positional data regarding meteorological phenomena and the geographic spread of medical illnesses all rely on accurate area and regional information. Even business applications such as efficient regional marketing approaches and effective distribution techniques

require such representation. For example, Figure 1.2 (employing a Lambert Conformal Conic Projection) illustrates the geological layout of Canada. If an equal-area projection had been employed, researchers could determine how much space overlaps with endangered species, while decision makers in the oil and gas industry could determine the amount of resources required for a size of a given project. Consequently, when working with data based on a spheroid shape, an approximation of our planet Earth, areal preserving projections become increasingly important.

With the advances of physical computational power, the ability to facilitate increased data, and in particular geoscience information, has vastly improved. Furthermore, improvements within the visualization and computer graphics community offers the potential for supplemental increase in the quantity of data presented, as well as improved integration between the visualization and underlying information. This latter attribute facilitates a highly desirable enhancement to the visualization and analysis for the scientific community. Determining such an association, while maintaining the aforementioned area-preserving characteristic, is the underlying focus of this research.

An efficient association of information is of the utmost importance. Consider, for example, a client analyzing bodies of water across Canada. In a rough approximation, taken from a coarse satellite image, there may be a thousand points defining the boundary of a single lake. If there were hundreds of lakes, as can be found in central Manitoba, this translates to over a million points that must be projected through this inversion process. Rivers, park boundaries and roads may also require simultaneous projection to meet the needs for clients such as modern businesses, military planners and scientific researchers. An effective approach must be taken to meet their real-time needs as experienced by geoscience visualization companies [40]. As such, a computationally and memory efficient approach is desirable, yet has not always been of historical importance.

The research herein is transitory in nature. Given established projection techniques, we

employ standard computer algorithms and approaches to improve their calculation time. This industrially-motivated problem lends itself to the consideration of alternative spherical representations. Recognizing a need for areal preservation, and inspired by the optimized projection - Snyder's Inverse Equal Area Polyhedral Projection [48] - we construct a sphere converging subdivision which additionally preserve global surface area. A simplified two-dimensional curve approach is explored prior to that of the spherical.

Given this sphere-converging work, a side exploration and extension is considered. Motivated by the need to represent the non-spherical Earth, a rudimentary construction of an equal area subdivision for arbitrary curves and surfaces is explored. Returning to the original problem at hand, the need for local area preservation is necessary to support the desired equal area quality. As the construction of such a refinement is challenging, a thorough application of the global area preserving, and sphere converging subdivision is applied to existing refinements. They are analyzed for their ability to support a local area preservance, with a final discussion on a recommended approach.

To begin our work, it is important to first understand how the Earth's data has historically been represented. How the connection between planar and spherical data is presently being handled is of additional relevance. Finally, it is of use to identify where modern computer science - in particular computer graphics - algorithms can improve these processes.

1.1 Traditional Cartography

Traditional cartographic projections transform a point on the Earth to a point on a map. This projection can be described through a function:

$$p' = F(p),$$

where p is the point on the sphere, and p' is the resulting point on the map (see Figure 1.3).

Over the years, numerous projections have defined the function F - each with its own underlying motivations and benefits. Some preserve distances along a line, others aim to preserve the shape for aesthetics, and still others preserve area - an important attribute to analysts. The mathematical definitions of these projections are often constructed with this forward projection in mind. The inverse, or $F^{-1}(p)$ in Figure 1.3, defines the projection from the planar map to the spherical Earth. They are often derived directly from the forward projection.

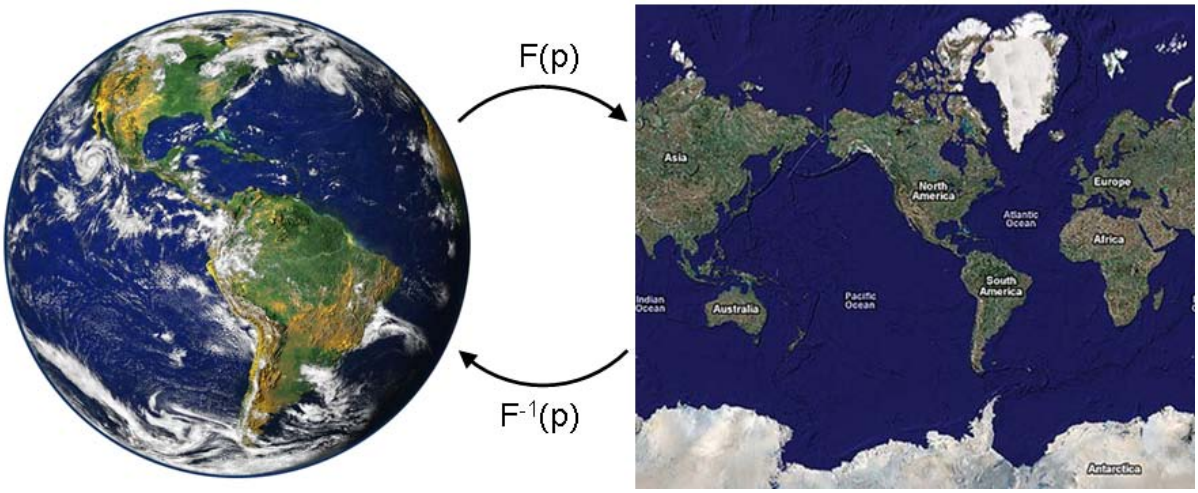

Figure 1.3: Projection from the Earth to Planar Map and Inverse (Right: Blue Marble, NASA [38], Left: Mercator Projection, Google Maps [24])

As modern computers have improved, they have facilitated the ability to analyze geographical data in real time. Much of this analysis, which has otherwise relied on manual map construct, involves data-mining of a regional nature.

Such area-preserving projections exist but often exhibit severe shape distortion. Snyder's [48] polyhedral projection aims at reducing angular distortion, and has since been recommended for equal area projections [31]. The more recent Slice and Dice projection presented by Leeuwen et al. [57], maintains areal preservation yet distributes angular distortion more uniformly. Transforming information from a planar representation to its spherical locality, or their inversions, is traditionally unimportant. Yet within a computational visualization environment,

this inverse process becomes extremely important since much underlying data is stored in a planar form. Geographic information systems (GIS), for example, employ such tactics to ensure enhanced accuracy of the data presented, enabling researchers with qualitatively improved methods for analysis.

Snyder's equal area approach defines such a forward and inverse projection. His method is unusual in that it projects between a polyhedron, circumscribed within a sphere, and a sphere. Since this polyhedron more closely resembles the sphere upon which the data will be projected, the resulting angular distortion is greatly reduced from the traditional single-plane map projections. Area is maintained by ensuring both global, as well as local preservation. Through a collection of trigonometric equations, the function $F(p)$ is described. The inversion, $F^{-1}(p)$, occurs as a direct reversal of the forward projection. From the trigonometric equations, a nonlinear system evolves. Since neither Snyder nor traditional evaluations are able to construct an analytical or closed form, a numerical technique for finding a solution is employed.

As the recommended numerical approach requires an indeterminate number of iterations for finding a viable solution, it acts as an operational bottleneck for computation time. The closed form of the forward projection avoids this since it does not require such iterations, thereby ensuring a fixed calculation time. However, the iterative process from the solution finding repeats calculations and therefore increases the overall time. Since these iterations are applied for each call of the inversion, which in turn is called potentially millions of times for a single visualization, the process can slow down well below real-time requirements.

There are few other approaches for preserving area, while minimizing distortions as is found with Snyder's polyhedral method. Leeuwen et al.'s Slice and Dice projection offers a valid alternative, by decomposing the equal area projection through a general segmentation approach. Area is equivalence is maintained through the preservation of area in the mapping of infinitely small regions. Unfortunately, the approach does not include an explicit definition of the inversion, let alone a computationally efficient one.

One aim of our research is to optimize the inversion process employed by Snyder. In this way, it may be used to improve environments requiring real-time on-the-fly feedback. A variety of numerical and computational techniques are applied to improve the speed of the inversion. Repeated calculations are identified and removed, efficient polynomial calculations are employed, and - most importantly - a reduction in the iterations of Newton's method is achieved. By evaluating data passed through the iterative process we can reduce $F(p)$ to a one dimensional curve. This improved estimate reduces the time it takes to converge. An additional approach of applying the polynomial approximation directly, without requiring the iterative solution-finding is also explored.

While this improvement to the traditional projection-approach to the association of planar and spherical data is beneficial, it does not take advantage of some of the visualization approaches developed over the last number of years. Integration with computer graphics techniques, including multiresolution representation, offers added benefits to the previously employed representations.

1.2 Technological Developments

Although cartographers have increasingly employed computational tools for their map construction, they have directly used these traditional projections. As such, the approaches are ill designed to effectively take advantage of the fast algorithms and effective hardware solutions available to them. Through improved analysis, we have the potential to perform such geographic tasks in fractions of the time it would have taken without these tools.

Cartographers historically constructed maps manually through paper and ink. Because of the static nature of paper, it is restricted in the amount of information that can be presented. Subsequently, if a meteorologist and a politician required different features highlighted for a similar geographic region, the base boundary features would require manual redrawing. Con-

sequently, different maps would need to be redrawn - by hand - for geological, political, environmental and other boundaries.

Current algorithms are capable of automating much of the cartographic process. Furthermore, for alternative or even dynamic datasets, computers can assist in presenting this information. As processing power of computers have increased, increasingly sophisticated visualizations are becoming capable, which in turn drive the potential and expectation for more improved visualizations. For example, within the last decade, the vectorization of cartographic maps was possible. Duplication and overlaying additional datasets became possible shortly thereafter. As the number of supported layers, alongside improved resolutions, became increasingly capable, so too did the support for the dynamic visualization of datasets. For example, we are now capable of visualizing multiple frames in a changing geographical boundary interactively.

Up until the advent of Google Earth, much information has been generally displayed in a two dimensional planar form or map. As mentioned above, traditional planar projections of a spherical Earth result in distortions - frequently combinations of angular, areal and distance. With a fully 3D spherical representation, these distortions fall away.

Supplementary to this projective distortion, computers offer the ability to transform data. As such, a layering - often of an hierarchical nature - is possible. A coarse, almost zoomed out approach can give a broader picture of a region, as compared with a refined zoomed in approach. Such abilities are desirable within the scientific community, as denoted by their requests for maps at differing resolutions. Simultaneous support of such layering or hierarchical data is possible within our computational environment.

As such, the resultant question becomes how to computationally support an hierarchical equal-area representation for the data requirements of scientists and those in industry. One approach is to explore the use of a IS3EA data structure. This approach employs the Inverse Snyder Equal Area Polyhedral Projection [45,48]. Through an underlying polyhedron, each face is

subdivided using an established planar division - constructing subsequent sub-faces. Upon subdivision completion, the defining points, or vertices, of the faces are projected through Snyder's projection onto the sphere. While this is viable for hundreds of repeated vertex calculations, even with optimizations, it is costly and computationally intense.

This has inspired the discussion of approaching the visualization process in a more comprehensive manner. Commonly, analysis occurs separately to the visualization process. If we could retain the data within structure which easily facilitates visualization, we may be able to improve both tasks. For example, a researcher may be interested in the data immediately adjacent to a given region. With data stored traditionally, such a search might require extensive searching. However, if the data is stored in a spatially oriented structure, this operation becomes constant.

The work herein endeavoured to explore the visualization of hierarchical data. Parallel explorations within my research group, under the same project, are exploring the indexing and data association from such a hierarchy.

It should be pointed out that the advent of Google Earth has assisted in pushing for the technological development, and incorporation of computer techniques within the realm of geography. Interestingly enough, they employ a traditional latitude and longitude format for data representation and association. This rectangular representation of the Earth's surface, while traditionally simplistic, and popularly employed in such systems as Geographic Information Systems (GIS) and Global Positioning Systems (GPS), it does not offer the analytical robustness of an equal area representation. Although it offers a popular approach, the lack of technically involved representations hinders its ability for ease of use within the scientific and analytic community.

Snyder's equal area projection [48], employed simultaneously a polyhedron as well as a polyhedral globe - or spherical polyhedron[1] - defining the projection between the two repre-

[1] A spherical polyhedron is a polyhedron wherein vertices are positioned a fixed radius from a single point or

sentations. Since these polyhedrons - the platonic solids as well as the truncated icosahedron - more closely resemble a sphere, they reduce much of the angular distortion. From this suggestion, he opened the gates for developing equal area projections, and further exploration on polyhedral globes.

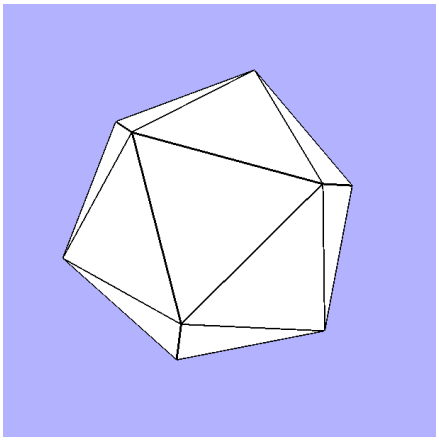

(a) Icosahedron

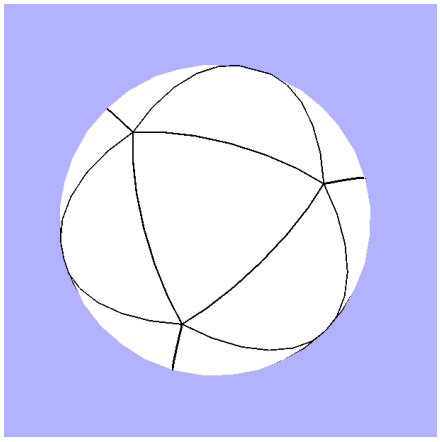

(b) Spherical Icosahedron

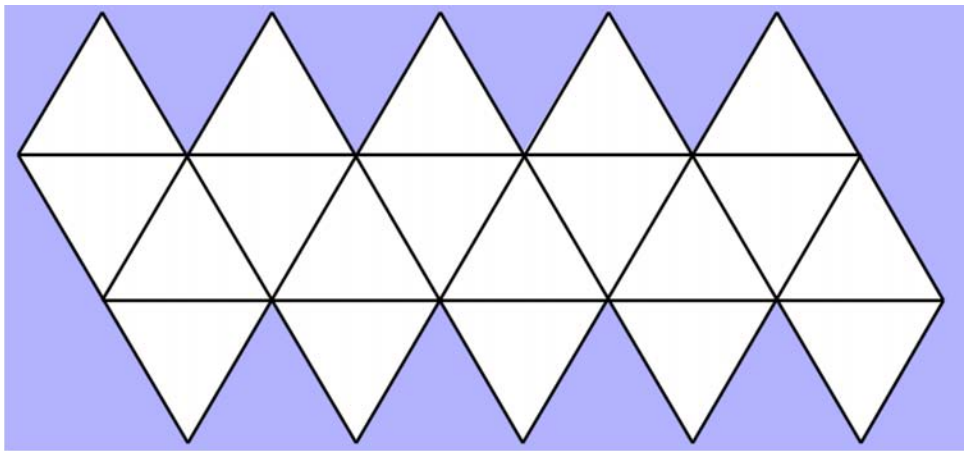

(c) Unfolded Icosahedron

Figure 1.4: Various Visualizations of the Icosahedron

The use of a polyhedral mesh as the replacement structure influences the resultant angular distortion. During any projection, only a limited number of properties may be retained. For example, if distances are kept consistent with the original spherical surface, the resulting map

origin. The linear edges of the polyhedron are replaced with great circle arcs, again residing on same sphere of the given radius and origin.

would be required to also be spherical in form. Consequently, projections are designed for the specifications required - such as distance preservation along a single line of longitude - keeping in mind the inherent limitations. Simultaneous retention of areal and angular (or shape) properties is not possible. As we aim to preserve area for the analytical benefit of researchers, attempting to minimize the angular distortion is also desirable.

Since the polyhedron more closely approximates the sphere, as compared with a traditional planar map, it offers an improvement in minimizing angular distortion. In fact, Snyder's projection results in less than 3.75^o angular deformation when the truncated icosahedron is used. Consequently, employment of this non-planar underlying structure not only benefits the community for computer graphics applications, but also for reductions in angular deformations.

It should be noted that the use of a polyhedron is indeed untraditional. The resultant shape is often presented in three dimensions, which is non-standard for viewing maps. Should an individual prefer direct manipulation of entities on the polyhedron it is possible to unfold the shape onto a flat plane (See Figure 1.4). While this will exhibit several interruptions between adjacent locations, it is generally not intended for planar evaluation.

To improve the association between the data and the visualization, we adopt Snyder's polyhedral tactic, and employ a mesh as our base data structure. Through the use of an advanced computer graphic technique known as multiresolution subdivision citezorin98, we can refine our coarse polyhedral representation of the Earth's surface, to a more refined spherical approximation. The objective then becomes the establishment of a subdivision process which preserves area throughout the iterations.

1.3 Subdivision Approach

One of the many facets of computer graphics is the subdivision algorithms developed therein (See Figure 1.5). Designed to enhance continuity (or smoothness) of form, these resultant

surfaces often exhibit self-similar decomposition in their divided levels. If we could retain areal equivalence for a spherical entity, throughout this process, we would assist in meeting the needs of the geographic community.

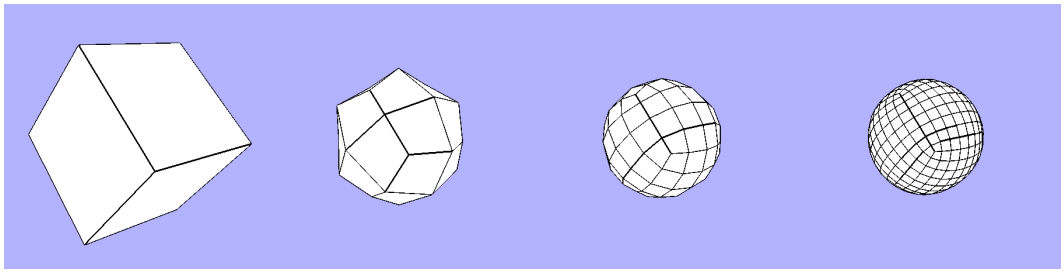

Figure 1.5: Catmull-Clark Subdivision [12] Applied to a Cube

By defining a multiresolution structure or subdivision which preserves the area through the refinement, we improve the facilitation of data association. Such a preservation of area requires a local maintenance. By extension, this ends up supporting a global preservation of area. The global requires the net surface area to remain constant throughout the subdivision process, while the local preservation requires areal equality across all faces at a given resolution. The former is simpler to ensure, particularly with the employment of a scaling factor. The latter requires a sophisticated evaluation of face generation, and while maintaining traditional characteristics of a subdivision process, such as smoothness and convergence. Since the problem in question relies solely the data representation of a spherical surface, this second local area preservation can be simplified by ensuring the subdivision surfaces generated aim to converge to a spherical polyhedron.

To initially develop such a scheme, testing within a two dimensional environment offers the ability to establish a more fundamental approach. Starting with a closed polygon inscribed within a circle, it is important to develop a subdivision which preserves the length, or perimeter, while improving its circular approximation. Because the subdivision aims to better approximate a circle, it inherently becomes smooth at the limit curve.

Upon the establishment of a two dimensional subdivision, the next stage is the development

of a similar subdivision for a closed polyhedral surface. In this case, the closed polyhedron will be circumscribed within a sphere, and the subdivision will aim to preserve the surface area while improving its spherical approximation. The global preservation of surface area becomes more disjoint with the local preservation due to the additional dimension, and resulting flexibility in subdivision development.

Such a scheme exploits the circular (curve) or spherical (surface) nature of the system to retain net-length property while further maintaining smoothness. Currently it is applied only to circular or spherical converging polygons or polyhedrons.

Due to the nature of the global preservation, it becomes apparent that as supplementary subdivision may be developed. Since it is reliant on circular or spherical convergence, the extension to an arbitrary curve or mesh might be readily applicable. Notice that the tangents of a circle define all possible tangents along a curve. If one could develop a curve such that its tangents represent a continuous transition of tangents along a circle, then the curve would demonstrate a smooth quality. Such a subdivision scheme is presented herein along with an affiliation with global length preservation.

1.4 Overall Outline

The research is presented systematically, following fundamental developments, preliminary applications, experimentation and finally results. An initial discussion surrounding past and present approaches in the representation of geographical data, with particular attention towards areal preservation, is presented in chapter 2. Upon completion, a more in-depth discussion surrounding modern approaches to geographic visualization is explored in chapter 3. Chapter 4 looks at the improvement of the traditional projections, with particular focus on Snyder's Equal Area Polyhedral projection.

Upon the establishment that such improvements are still disjoint with a comprehensive so-

lution, the main body of work is presented chapter 5. Here, the global area preservation is presented for both curves and surfaces. A local preservation is trivially discussed for curves, and its more sophisticated surface application is explored. A minor digression on the application of such approaches to arbitrary curves and surfaces is discussed and briefly evaluated.

Chapter 6 employs an experimental approach is for local preservation of the spherical subdivision process. A collection of possible candidates are presented and evaluated for their potential to naturally be employed as an equal area subdivision. An analysis of their success and failures is presented. A final discussion on potential extension is presented in chapter 7, and the conclusion is presented in chapter 8.

Chapter 2

Related Work

Cartography has a rich and diverse history. Dating over two thousand years, the visualization of the Earth's surface on planar maps has long been an important task. Traditional projections are used to mathematically define the transition from a point on the sphere to a point on such a planar representation. These can be done in a variety of ways, while maintaining a collection of properties. The preservation of area, especially during such a projection, is highly important for information analysis. Examples of regional-based analysis include population densities, determination of overlaying geographic regions, and optimal distribution of goods. Understanding and supplying such information is important within the biological sciences, within socio-economic evaluations, and within businesses through their determination of effective marketing strategies. Unfortunately, due to the curvature inherent in a sphere, it is impossible to define a projection that simultaneously conserves area and shape, or angles. Practically and mathematically, flattening a region of the Earth will result in an uneven stretching, therefore resulting in an inability support both conformality (shape) and area [10]. The resulting goal is to define a projection such that area is preserved, and shape distortion is minimized.

Much of this research is motivated by the last few decades of technological improvement, and the consequent ability to support increased data storage and analysis. To support global data, one must either maintain a separate data storage location, and project the results to the spherical representation of the Earth, or seek out an alternative approach which more closely knits the data with the visualization. A promising approach is to incorporate the computer graphics concept of multiresolution subdivision. Such a tactic will enable a well founded hierarchical structure for the data therein, while closely associating the data visualization with its storage. From this close association, global information will be more positionally aligned,

and can therefore support an improved neighbourhood awareness - a property that is highly desirable during data mining tasks.

While a projection-based approach for data visualization is possible - and currently employed by numerous companies - there are a number of inherent detriments. For members of the research community who work with the resultant maps, the distortions inherent in the projection, such as areal, positional, or conformal, can often wreak havoc on the data analysis. Furthermore, storage of a projection-based approach makes for ineffective retrieval of information. Since data ends up being indexed based, which is not well associated with positional information, due to the projection, it becomes challenging to quickly find data that is geographically nearby. While this may seem solved due to the increasingly faster computer systems available, for the terabytes of data potentially streamed every second - for example from satellite communications - this projection approach rapidly becomes a limited solution for real-time visualization and analysis.

The basis for an alternative, subdivision-based approach was inspired by the polyhedral projections from the geographic community. These are generally limited to a low number of faces, or limited polygons, however, constructing a multiresolution polyhedron which preserves area through the duration of the subdivision may better facilitate data analysis.

Traditional projections, as well as promising computer graphics concepts are explored in the following subsections. Through this discussion of related works, a more robust definition of the overall problem is presented. Furthermore, existing approaches, including their benefits and detriments are described and discussed.

2.1 Cartography and Equal Area Projections

Maps of the earth have been around since we first developed pictorial forms of communication. With the advent of the printing press, these maps were distributed in more expansive forms,

and enabled the proliferation of knowledgeable and therefore safer travel.

As previously discussed, maps of the Earth's surface are well known dating back to around 200 B.C. Ptolemy's *Geographica* manuscripts, in the second century, explore a variety of approaches to mapping [49]. Over the years, different mappings have been designed to address particular situations, or desirable presentations of the underlying geographic information. Commonly incorporated attributes include the preservation of shape, distance, area, and sphericality. Depending upon the objectives of the user of the map, different mappings are selected to offer the best visualization for the problem at hand. For our research, the preservation of area has been deemed highly important, due to its necessity within the scientific, military, and business community.

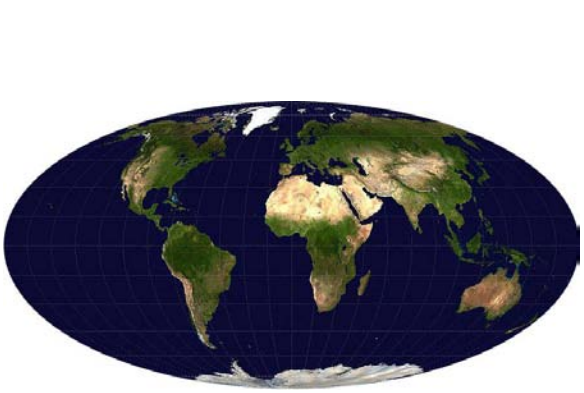

(a) Mollweide

(b) Werner

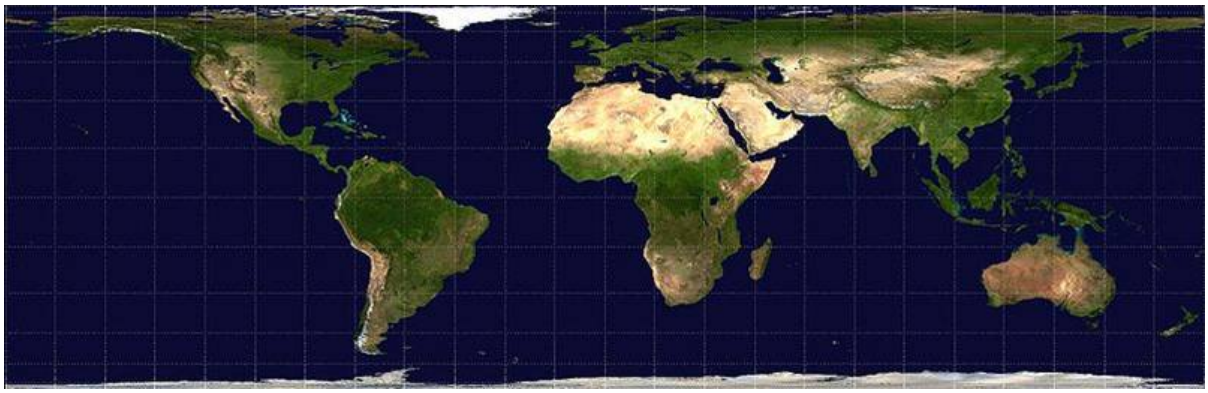

(c) Lambert Cylindrical Equal Area

Figure 2.1: Area Preserving Projections (Blue Marble Series, NASA [38])

The mathematical definitions of equal area mappings, or projections, have long been studied. Werner's projection [49, 60] of the 16th century (Figure 2.1(b)), for example, generates a heart-shaped map. This projection simultaneously preserves distances along each parallel and a central meridian. This means that someone using the map would have an accurate measurement of distance - with some scaling effect - along this region.

The Lambert Equal-Area projection was established in 1772 by Johann Lambert [50], and is one of the few equal area projections still in popular use [13, 55] (Figure 2.1(c)). They project the surface of the Earth to a more regular rectangular space. Consequently, the amount of distortion is reduced, as compared with Werner's. However the additional distance preservation characteristic is lost.

The Mollweide projection [49] of the 19th century (Figure 2.1(a)) retains areal equivalence while further reducing angular distortion due to its elliptical form. Unfortunately, the angular distortion is still high along the boundaries and non-uniform across the planar mapping.

Numerous others have been described over the years [50]. Both Lambert's (1772) and Behrmann's (1910) Cylindrical Equal Area projections result in vast stretching in the poles. The Sinusoidal (16th century) and variants, resembles a spinning top, severely distorting around the boundaries of the map. The Eckert IV (1906), which is popularly used in textbooks, has a recti-oval shape, reducing the distortions of the Sinusoidal, but still exhibiting them around the boundaries and throughout the poles. Albers' (1805) and Lambert's (1772) Conic Equal-Area projections result in half-circle maps, exhibiting extreme distortion around the boundaries. As a consequence of their planar and continuous maps, severe angular distortion occurs.

If we accept the notion of a non-smooth boundary for our contiguous mapping, we can explore alternative approaches to the projection process. The McBryde S3, Goode Homolosine (Figure 2.2) and others [50] have experimented with equal-area projections involving interruptions in the map. However, these approaches tend not to be well suited for conversion to a 3D visualization as they are not easily converted to a consistent 3D form. The involved mathemat-

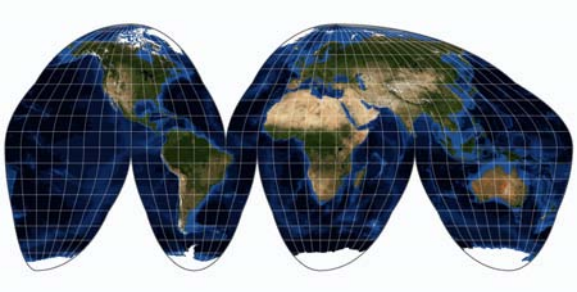

Figure 2.2: Goode Homolosine Projection (NASA [38])

ical equations for projecting the visualization, and the employment of the inverse projection to enable analysis, take extended amounts of time, which add up quickly. An interrupted mapping better suited for computer graphics visualization instead relies on a polyhedral projection.

2.2 Polyhedral Projections

The polyhedron has been considered as an alternative to a flattened planar projection dating back to the 16th century [49]. Durer, in 1538, [19] explored the flattening of polyhedron such as the tetrahedron, dodecahedron, icosahedron and other shapes however he did not connect their potential for employment with map projections.

It was not until the 1600s when Ignace-Gaston Pardies would apply the sky onto the cube, based on the gnomonic projection. With the equatorial poles, their midpoints and the equinoxes centered on their respective faces, the resultant representation aims to approximate the globe.

The attractive map produced by Pardies inspired a number of others, which are further discussed in Snyder (1997). These would sporadically be explored over the centuries. More recently, in 1976, L.P. Lee developed a fairly involved conformal polyhedral projection [34].

Bradley, in 1946, constructed the first equal area projection onto a polyhedron [7]. Applied approximately to the icosahedron, this mapping employed a cylindrical equal-area projection. Unfortunately, through this cylindrical use, the planar equilateral triangles are not identical in area to their respective equilateral spherical triangles.

Irving Fisher addressed this issue as an editorial footnote to Bradley's work. Within it, Irving employs the Lambert Azimuthal Equal Area projection[1], centering onto each of the faces. Unfortunately, the edges of symmetrical regions from this projection are not well aligned, resulting in discontinuities across the faces. Figure 2.3 illustrates the radiating discontinuities.

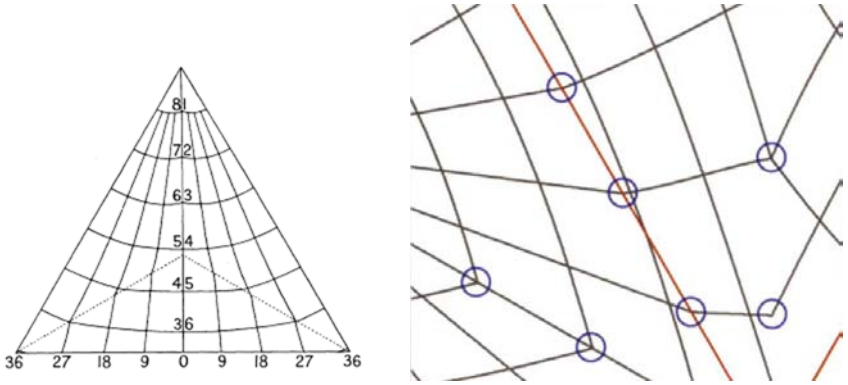

(a) Icosahedral Face [7] (b) Cusp Issues [57]

Figure 2.3: Irving's Projection - On the Icosahedral Face, with Cusps and Discontinuities Highlighted

Irving's derivation was re-derived by White et al in 1992 during their work on EMAP sampling for the U.S. Environmental Protection Agency [62]. This mapping was the one of the first times these polyhedral projections had been employed for computational use. More importantly, this projection laid the groundwork for developing a global grid system for effective and modern use within a computational environment [45].

Snyder extended this re-derivation to ensure a simplification of formulae that could be applied with nominal adjustments for each of the five platonic solids, in addition to the truncated icosahedron. It is consequently this projection that provides the foundation for an equal area polyhedral projection. By centering a known equal area projection, and providing slight adjustments to ensure edge matching - as encountered by Irving - and areal preservation, it offers a natural progression from the traditional planar to the employment of a polyhedron for mapping

[1]The Lambert Azimuthal Equal Area projection differs from the aforementioned Cylindrical projection. The former directly projects the sphere to a disk, while the latter converts it to a cylinder before projecting to a rectangular map.

purposes. While there are visual discontinuities from the resulting unfolded polyhedron, in folded form, it represents a very close approximation to the otherwise spherical Earth - a benefit in modern day computing which relies on the use of small faces or triangles to approximate the surface of a sphere.

Snyder's work is discussed in detail in Section 3.1.3 as its understanding is essential for the development of the work herein.

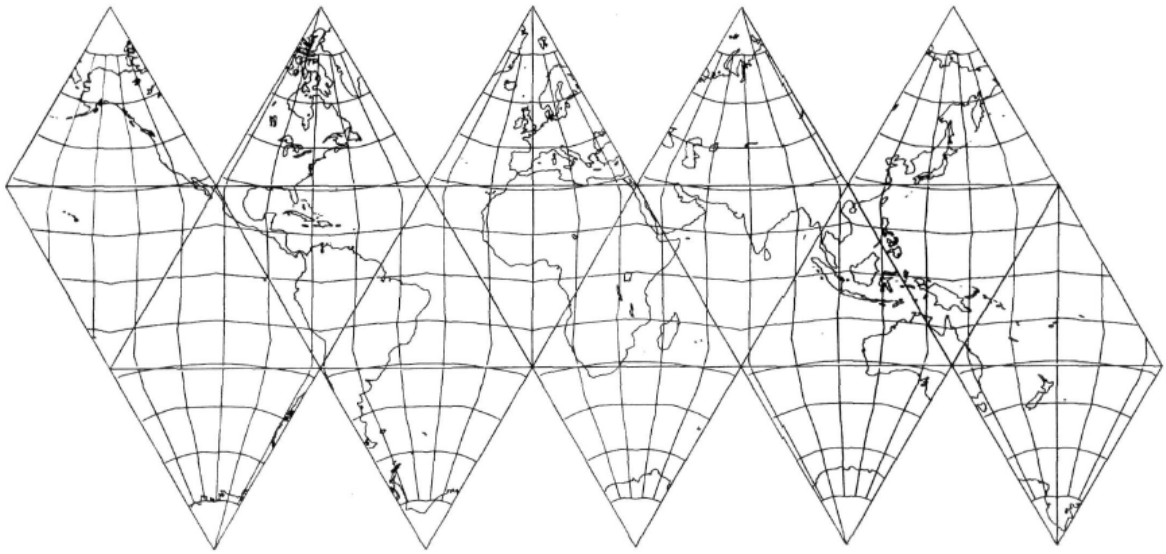

Figure 2.4: Icosahedral Mapping Using Snyder Equal Area Projection [48]

Through Snyder's use of a polyhedral surface, and its improved approximation of the spherical Earth, properties such as angular distortion, are reduced. For example, with an icosahedron - illustrated in flattened form in Figure 2.4 - Snyder achieved an angular deformation of less than 17.3^o and a scale variation of less than 16.3%. White et. al. and later Kimerling et. al.'s evaluation of the method against other approaches resulted in categorizing the Snyder projection as well suited for preserving area, and less well suited for ensuring compactness - a method for evaluating similarity in resulting shapes [31, 63]. This matches with the objectives of the projection, which aims to preserve area.

Leeuwen et al. [57] provide a more recent alternative to Snyder's polyhedral projection.

They demonstrated an areal preservation which distributes the angular deformation more uniformly across the surface of the projection. Their Slice and Dice approach explores areal calculations from a *derivative* approach. They start comparably to Snyder's projection, employing an initial polyhedral tessellation of the sphere, and reducing the problem to the smallest distinct region. From here, they divide the spherical triangle into four slices, maintaining areal ratios en route (See Figure 2.5). An intersection of these slices defines the position of the projected point. As a consequence, the distortions become less noticeable, eliminating discontinuities and reducing cusps. While their forward projection is defined - albeit fairly involved, requiring a dozen trigonometric calls - the inverse is not.

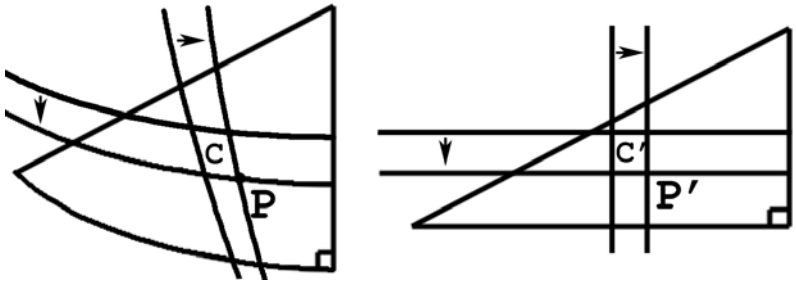

(a) Spherical Triangle (b) Projected Triangle

Figure 2.5: Slice and Dice Projection [57]. The projection has the area of $C \rightarrow C'$ as the segmenting lines approach each other.

Leeuwen et al explored the work of Snyder and sought to improve it. They re-oriented the projection, as well as deriving a fundamentally unique approach through their Slice and Dice technique. Because of the initial groundwork laid out by Snyder, Leeuwen et al were able to approach the practical task of expanding and improving this idea of a polyhedral projection. Consequently, their angular distortion through their fundamentally different approach (as sought by the disassociation with the original underlying Lambert Azimuthal Equal Area Projection employed by Snyder, and his predecessors) is reduced, and more uniformly distributed across the face of a given polyhedron. This ensures a visually improved projection. Unfortunately, unlike Snyder, an inversion - namely the projection from the planar map to sphere - is

not directly presented from their works. Such an inversion is required when employed with modern computers to ensure accurate conversions between the planar and 3D representation of the data.

It should be noted that Song et al. [51] presented an equal area small circle subdivision. This subdivision divides triangles directly on the sphere into subtriangles using small circle arcs. For a single triangle, dozens of trigonometric calls ensure that all subtriangles are identical in area. While these faces could be indexed to their respective planar triangle, finding the projected location of a point would require repeated subdivision until the vertex of a subtriangle is sufficiently close to the point of interest. This approach is quite costly, especially when requiring numerous subdivisions to obtain high precision. As the objective of this work is to compute an inverse projection quickly, Song's approach does not lend itself to a viable implementation.

Similar to Song et al. [51], the recent work from Rosca and Plonka [41], and that explored by NASA's Jet Propulsion Lab [26], defines a segmentation of the sphere through a polyhedron. With a polyhedron as the fundamental initial surface, the application, as with Song et al., occurs directly on the surface of the sphere. This spherical discretization differs from the approach taken herein. Since the visualization of the resultant sphere will require its own discretization, affecting the discretization directly offers a potentially robust approach. Ideally, as demonstrated by our work, the discretization of the sphere will not require re-calculation, since the underlying surface is already composed of a mesh-like structure. As such, the surface is immediately ready for visualization, and does not require secondary data transference, as would be the case in spherical discretization.

2.3 Subdivision Surfaces

Before getting too far into the detailed approaches available through the use of computer graphics, it is useful to define subdivision surfaces. This area of research is fundamental to a number of other topics, which in turn may be applicable for our work on new representations of the Earth's surface.

A mesh, in computer graphics, is a collection of faces, and each face a collection of vertices and edges. Vertices are points in space. An edge can be viewed as the line connecting two vertices. A face is a region bounded by a collection of edges. For a face, these edges form a non-intersecting cycle. Ideally, all edges (and consequently their respective vertices) reside on a singular plane for a given face, however this is not necessary.

Consider, for example, a cube. It is composed of eight points or vertices. These eight points define twelve edges, and these edges make up six square faces. In this case, a cube is a closed surface. A closed surface is one wherein every edge is adjacent to exactly two faces, with a distinct segmentation between an inside space, and an outside space (See Figure 2.6). In particularly, there are no boundary edges. For the purposes of our construction of a representation of the Earth's surface, we will be working only with closed surfaces.

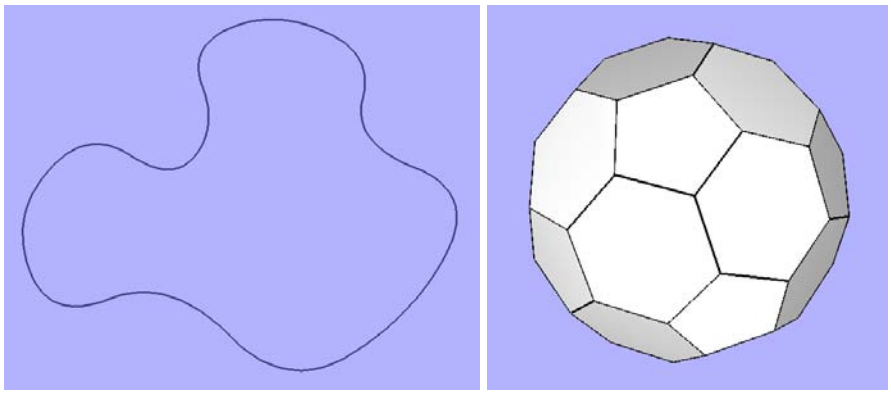

(a) Closed Curve (b) Closed Surface

Figure 2.6: Closed Examples. The curve and surface act as a boundary separating an inside and outside region.

Any three dimensional shape can be represented as a collection of faces. Within computer graphics, such a collection is generally referred to as a mesh[2]. Since a mesh is composed of a finite collection of faces, it is only capable of approximating the shape in question. Furthermore, if we have a very small number of faces employed to approximate the shape, we lose the details on the shape, such as scratches, or rounded edges. Such a mesh is often referred to as coarse.

Two commonly used subdivision techniques are Catmull-Clark Subdivision and Loop Subdivision. These approaches, employed later within the research, are discussed in further detail.

2.3.1 Catmull-Clark Subdivision

In 1978 Ed Catmull and Jim Clark [12] devised an approach for automatically constructing smooth surfaces given a coarse mesh. From the coarse mesh, the faces are split into smaller faces, and the vertices repositioned to gently round the edges. In this way, the mesh is subdivided into a higher quality mesh. After a few iterations of such a subdivision, the mesh, with its finite faces, becomes a better approximation of a smooth surface. Figure 2.7 illustrates the first two subdivisions on a triangular face, and the result of subdividing an icosahedron twice. Their work has been greatly used since then, particularly in the automotive and entertainment industries [16, 35, 37].

Within the Catmull-Clark method, faces of an arbitrary number of edges, are subdivided by their respective midpoints, converging to a surface that maintains G^1 continuity at extraordinary vertices, and C^2 everywhere else. Within the first level of subdivision, the surface becomes wholly constructed from quadrilaterals (which may not necessarily exhibit planarity).

As illustrated in Figure 2.7, each face is split according to the number of edges it contains. Vertices are inserted along each of the edges, and one within the interior of the face. The

[2]The definition of a surface also applies to parametric representations which need not necessarily have explicit faces. For our purposes, the terms ‘mesh’ and ‘surface’ are generally used interchangeably since one might consider a mesh as a coarse representation of a surface.

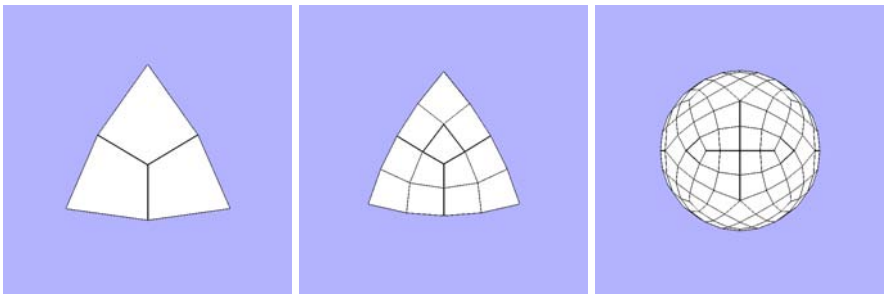

(a) Level 1 Face (b) Level 2 Face (c) Level 2 Subdivision

Figure 2.7: Catmull-Clark Subdivision

interior face vertex, v_f, is selected to be the midpoint of the face's original vertices. In other words:

$$v_f = \frac{1}{n}(\sum_{i=1}^{n} v_i),$$

where v_i are the face's original vertices. From this, vertices inserted along the edges, v_e, are averaged from their nearest neighbouring vertices - the two end points of the given edge, and the recently constructed face points from the adjacent faces. Mathematically we denote this as:

$$v_e = \frac{1}{4}(v_j + v_{j+1} + v_{f1} + v_{f2}).$$

For each original vertex from the original face, a new quadrilateral face is constructed by connecting the original vertex, the face vertex, and the two edge vertices, as visualized in Figure 2.7. Lastly, the original vertices are repositioned to ensure improved smoothness for the resulting surface. It is positioned based on its neighbouring vertices as such:

$$v'_f = \frac{1}{n}(\sum_{i=1}^{n} v_f) + 2(\sum_{i=1}^{n} v_e) + (n-3)v_f,$$

where $avg(v_f)$ represents the average face point from all neighbouring faces and $avg(v_e)$ represents the average edge point from all neighbouring edges, with respect to our vertex v_f.

2.3.2 Loop Subdivision

Another commonly used subdivision approach is the Loop Subdivision, developed by Charles Loop in 1987 [36]. Reliant on a triangular mesh, this subdivision creates interior faces based on the midpoints of edges. Designed for triangular meshes, the resulting subdivision surface exhibits G^1 continuity at irregular vertices and C^2 continuity elsewhere.

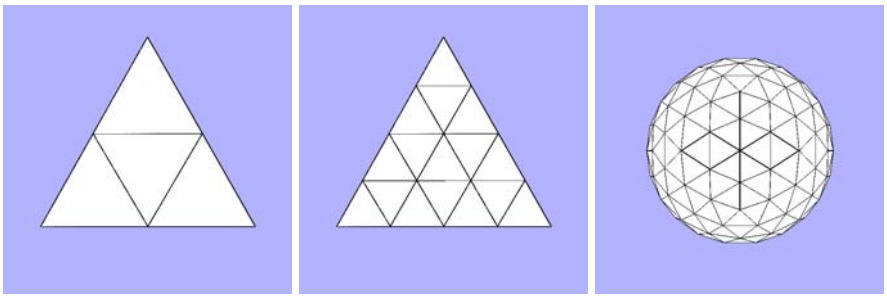

(a) Level 1 Face (b) Level 2 Face (c) Level 2 Split

Figure 2.8: Flattened Loop - Split (1:4)

The first three levels of vertex insertion upon an icosahedron (and not repositioning) is illustrated in Figure 2.8. For each face of the coarse mesh, interior faces are created by identifying the midpoint of each edge, and connecting them. For Loop's subdivision scheme, midpoints are repositioned according to the equation:

$$v_e = \frac{3}{8}(v_{e1} + v_{e2}) + \frac{1}{8}(v_{f1} + v_{f2}),$$

where v_{e1} and v_{e2} are the vertices bounding the given edge, and v_{f1} and v_{f2} are the third vertices from the two neighbouring faces. Original vertices are adjusted according to the equation:

$$v' = (1 - n)sv + s \sum_{i=1}^{n} v_i,$$

where v is our original vertex, v' our repositioned vertex, v_i the neighbouring vertices of v, and s is our scaling factor:

$$s = \frac{1}{n}(\frac{5}{8} - (\frac{3}{8} + \frac{1}{4}cos(\frac{2\Pi}{n}))^2).$$

In other words, the midpoints along edges are inserted as vertices, and the interior triangle constructed from these midpoints divides the original triangular face, resulting in four triangular faces. The original vertices are then repositioned to attain the smoothness qualities discussed above.

Several other subdivision approaches have been presented over the years. I refer the reader to Warren's book of subdivision methods for a more detailed discussion on the different variants [59].

2.4 Multiresolution Representations

One of the ways in which data can be visualized in modern day computing is through multiresolution representations. These have been gathering gradual support and employment in the visualization community [28, 32, 35, 42].

Under the motivation of Forsey and Bartels [21], an hierarchical representation of data structures has been expanded upon. A multiresolution representation refers to a surface which can be refined, or made coarser. While subdivisions, as discussed above, have been fundamental in refining a surface, the advent of reversing the subdivision process [4, 46] enables a full multiresolution representation. This way both a coarse or finer visualization may be available as desired. More importantly, the representation spans a fixed amount of space, corresponding to the highest degree of refinement. When a coarse representation is desired - either due to a desire for increased speed, or to visualize a low resolution version - the highly detailed version is decomposed into this low resolution version, and the supplementary details are stored in the residual memory space. To this end, the high resolution variant is able to be reconstructed by merging the coarse representation with the supplementary details. In this way, a fixed amount

of storage space, limited only by the highest quality representation, need be used at any time.

While wavelets have been well used for image processing [52], they may additionally be applied to multiresolution surfaces [27]. The important benefit is that wavelets retain the details of a more refined resolution, and can consequently maintain a fixed storage space. This is extremely useful since directly working with high level details can be costly and time consuming.

This multiresolution also lends itself to an hierarchical representation. Since one can transition between coarse and fine representations, determining parent-child associations between data becomes easier.

Conceptually, we can explain a three dimensional example through the visualization process of Google Earth. This piece of software allows the user to view the whole Earth, or to zoom in to a particular region. Numerous zoom levels are supported. From a multiresolution stand point, the coarse representation would be the visualization of the Earth as a whole, while a zoom into, say, a city, would visualize a high resolution of the given region. Storing the data between these different levels of resolution is what makes up a multiresolution representation.

The traditional visualization of Google Earth employs planar data which is projected onto the spherical representation of the Earth, through the use of the General Projection [23]. An advanced multiresolution representation would instead have data indexed to the faces of the mesh representation of the Earth. A low resolution would have a limited number of faces, while a high resolution would construct a fairly close approximation of the spherical representation of the Earth.

An additional desirable feature of such a multiresolution representation would be for the faces to retain areal equivalence throughout their construction. For the purposes of scientific research, where there may be overlaid data spanning different levels of resolution, having such an accurate representation is highly beneficial, as discussed earlier.

Furthermore, since a multiresolution representation of the Earth's surface is applied indis-

criminately, every point on the Earth is readily covered. This is of particular importance due to the polar issue often exhibited within the geographic community. Due to the construction of the different projection techniques, the north and south poles often have erroneous results. The popular indexing employing latitude and longitude, as well as the Universal Transverse Mercator (UTM) [47] both stretch, distort, and require alternative visualizations for plotting the poles. The UTM system, for example, explicitly recommends the use of the Polar Stereographic for representation. Through a subdivision, or multiresolution approach, all portions of the Earth's surface may be treated more uniformly.

2.5 Constraint-Based Subdivision

Through the various subdivision approaches, there has been the necessity to support a variety of desirable qualities. The flexibility of subdivision surfaces construction facilitates the incorporation of constraints. These constraints may vary from degree or style of smoothness, to interpolatory or analysis based. Catmull-Clark [12], for example, aimed to construct a highly smooth surface. The Dyn and Levin's Butterfly subdivision [20] ensures original vertices are interpolated (i.e. remain fixed in place) throughout the subdivision. Though deformations of surfaces have attempted to preserve volume or surface area, none of their work, to date, has applied to the subdivision process.

Notice that our definition of subdivision is that in which computer graphics employs through its mesh structures. Other subdivisions applied directly on the sphere have been proposed. Song et al.'s Small Circle Subdivision employs mathematically involved small circle arcs to divide the sphere [51]. Tregenza [54], in conjunction with the recommendation from the Commission International de l'Eclcairage (CIE) [14], described and recommended the use of a 145 segment equal-area subdivision. Since we are aiming to associate our data with our graphical visualization - which occurs through the use of faces - these approaches do not foster develop-

ment of our work.

Throughout all of the related literature, there has been no attempt made at preserving the area on a sphere-converging mesh. The discretization of the sphere - which is increasingly explored - negates the fact that the underlying data must eventually be visualized. Within a graphical context this will either require the mathematically accurate but slow process of ray tracing, or else result in the overlay of the data onto a mesh-based surface. Our approach aims to more neatly associate the data with the planar, yet small faces of the surface which will be resultantly employed for visualization. This promises faster visualization, as well as ad-hoc analytical evaluations.

Chapter 3

Background

This chapter explores background information relevant for understanding the work presented within the body of this work. A more detailed discussion on subdivisions, projections and their inversions is presented.

3.1 Equal Area Projections

As previously discussed, a projection, similar to a function, transforms a point in one domain, into another. For our purposes, we apply the cartographic variant wherein a projection specifically transforms a point in spherical coordinates to a point in planar coordinates. Such a mapping may be performed in a variety of methods, as illustrated in the previous chapter. Again, our purposes dictate the requirement of areal preservation throughout our representation.

The preservation of area spans two definitions. The first is that an equal area projection will ensure any bounded region will have the same area within each of its representations. A lake on the spherical Earth will have the same area (usually with some constant scale factor) as its projection onto the planar map. If the mapping of any infinitely small region on the sphere maps retains its area on the planar map, the projection is said to be equi-areal, or equal area. Within the body of this work, a preservation of infinitely small regions shall be referred to as maintaining local area. An interesting result of this local area property is that the summation of all infinitely small regions on the sphere - resulting in the surface area of the sphere - when projected will created a bounded region of the same size as the surface area. This way, the total or global area is preserved.

Mathematically, a projection is said to be equal-area if an areal scale is constant. From

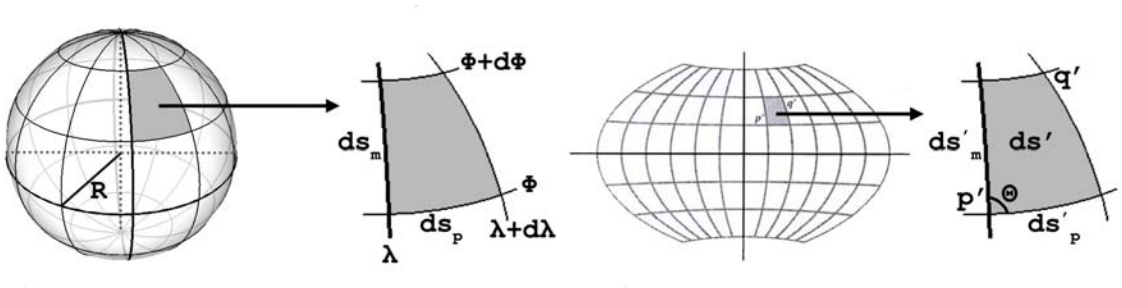

(a) Spherical Region (b) Projected Region

Figure 3.1: Elementary Quadrangle [10]

Canters' definition [10], an area scale σ is the ratio of the area of the infinitely small mapped quadrangle $ds'_m ds'_p \sin\Theta'$ over the area of the quadrangle $ds_m ds_p$ on the generating globe (See Figure 3.2). Values ds_m, ds_p represent the original width and height of the quadrangle, while ds'_m and ds'_p represent their values on the projected quadrangle. Angle Θ' is the angle between arcs ds'_m and ds'_p. To this end, σ can be defined as:

$$\sigma = hk \sin\Theta',$$

with h, the scale along the meridian and k, the scale along the parallel. These can be re-written as:

$$\sigma = \frac{\sqrt{EG - F^2}}{R^2 \cos\phi},$$

where R is our radius, and E, F and G our Gaussian fundamental quantities:

$$\begin{aligned} E &= (\frac{dx}{d\phi})^2 + (\frac{dy}{d\phi})^2 \\ F &= \frac{dx}{d\phi}\frac{dx}{d\lambda} + \frac{dy}{d\phi}\frac{dy}{d\lambda} \\ G &= (\frac{dx}{d\lambda})^2 + (\frac{dy}{d\lambda})^2. \end{aligned}$$

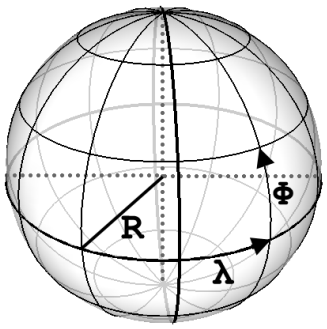

Figure 3.2: Generating Globe [10]

For an equal area projection, wherein $\sigma = 1$, we have the condition that:

$$\sqrt{EG - F^2} = R^2 \cos\phi,$$

or in other words, that the small quadrilateral region on the sphere should equal the area that would be computed for the bounding points. These angles are visualized in Figures 3.1 and 3.2.

Numerous projections preserve area as discussed in the previous chapter. The Lambert Equal-Area, Sinusoidal, Molleweide and Werner are just a few examples which preserve area. As they are projected to a single, uninterrupted planar map, they often exhibit high degrees of angular distortion, which will be mitigated by working directly on an approximation of a spherical surface.

3.1.1 Inverse Projection

While a forward projection refers to the mapping from a spherical surface to a planar one, its inverse refers to the mapping from a planar surface to a spherical one. Figure 1.3 from Chapter 1 visualizes this concept. The resulting coordinate conversion - in both directions - is commonly employed when working with a projective visualization. The forward projection

is used when transforming user input, which has occurred directly on the sphere, into its data-stored coordinates, namely the planar mapping.

The inversion of the forward projection is used ever more frequently as it does the actual transference of data to be visualized onto the surface in question. As the data requested for visualization changes - either from user input, or automatically from the system - the inversion accurately positions data onto the sphere for visualization.

For systems that employ projective visualizations, ensuring effective and efficient means of computing the forward and inverse mappings is necessary. Any reduction in speed directly impacts the amount of data that may be visualized within a reasonable time frame.

Traditional cartography has not been especially driven by the inverse projections. Since they are important for computer visualization, they are described in the following sections where possible.

3.1.2 Lambert Azimuthal Equal Area Projection

In 1772, Johann Heinrich Lambert presented numerous projections, including his Azimuthal Equal Area Projection [33]. This particular projection maps the Earth's surface to a disk. Well suited for limited regions of the Earth, such as the poles, it may also be simply modified to support some directional preservation. Consequently, it is well used by the National Atlas of the United States in its online Map Maker project [55]. The European Environment Agency recommends its use for statistical analysis and display [2]. Of particular interest, this projection is later employed by others when determining and equal area projection across a polyhedron (as is performed by Snyder in Section 3.1.3).

In Figure 3.3(a), we see the construction. Here, a point P on sphere S is projected to point P' on the plane. P' is the closest point intersecting the plane and circle C, whose origin is the point of tangency between S and the plane, and also passes through point P. As the point P gets closer to this point of tangency, the radius of circle C decreases. The antipode, or point

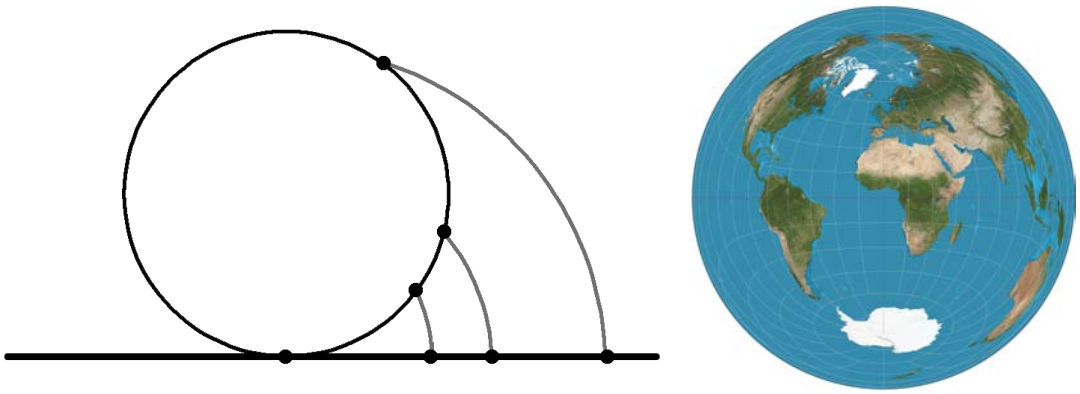

(a) Construction

(b) 15^o Graticule (Blue Marble Series, NASA [38])

Figure 3.3: Lambert Azimuthal Equal Area Projection Projection

opposite the point of tangency may not be uniquely mapped, whereas the point of tangency is mapped back to itself.

Explicit formulae for the projection $F(P) \rightarrow P'$, for $P = (x, y, z)$ on the sphere, and $P' = (X, Y)$ on the plane, are as follows:

$$F(x, y, z) = \left(\sqrt{1 - \frac{X^2 + Y^2}{4}}, \sqrt{1 - \frac{X^2 + Y^2}{4}}, \sqrt{-1 + \frac{X^2 + Y^2}{2}} \right)$$

The inverse, mapping $F^{-1}(P') \rightarrow P$ can be written as:

$$F^{-1}(X, Y) = \left(\sqrt{\frac{2}{1 - z}} x, \sqrt{\frac{2}{1 - z}} y \right)$$

From this construction, area is preserved since a small region on the sphere is mapped to the same area region on the plane. Since there is reduced angular distortion nearest to the point of tangency, the placement becomes important. Often, maps are constructed within a small region near this tangency point. Consequently, Snyder chose to center this projection on each of the polyhedral faces to retain this property, as discussed in the following section.

3.1.3 Snyder's Equal Area Map Projection for Polyhedral Globes

It is worthwhile to discuss the mathematics behind the Snyder projection, given its fundamental importance in this work. A further contribution is the clarification and correction of erroneously labelled elements within Snyder's paper.

Snyder's projection defines a function F which takes a point p on the sphere and determines its coordinates on a polyhedron. The main idea is to use a modified version of the Lambert Azimuthal Equal-Area projection (Section 3.1.2), centering it respectively for each face of the polyhedron. The modification corrects the otherwise imprecise edge matching between faces. Several steps define this projection. Firstly, the symmetric property of the polyhedral faces is recognized, and the problem is reduced to its smallest distinct region on the regular face - always a right-angle triangle (Figure 3.4). The second step ensures the triangle's area on the plane and on the sphere are equivalent through a scaling factor between the radius of the sphere, and that of the polyhedron's enscribing sphere. The third step generates a triangle on the polyhedron whose area exactly matches that of a spherical triangle bounded by point p. With this accomplished, the final step positions p' along this triangle's edge while maintaining areal scale (See Figure 3.5).

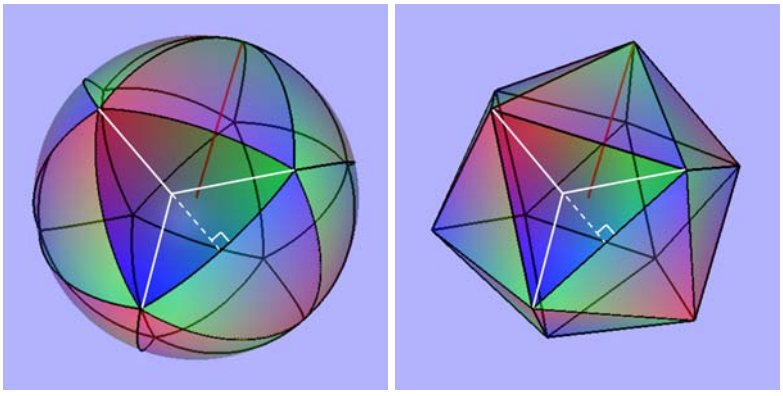

Figure 3.4: Spherical and Planar Icosahedron with Symmetric Decomposition (Red line indicates the radius)

Here we provide an initial definition of the triangles and their relevant variables through ac-

companying figures. The initial extraction of the smallest distinct region is visualized in Figure 3.4. This illustrates how the triangular face may be divided equally into three subtriangles and then further halved into a right angle triangle. This face splitting may be applied to any regular polygon, which in turn makes up the faces of the platonic solids and the truncated icosahedron - common sphere-circumscribing polyhedra.

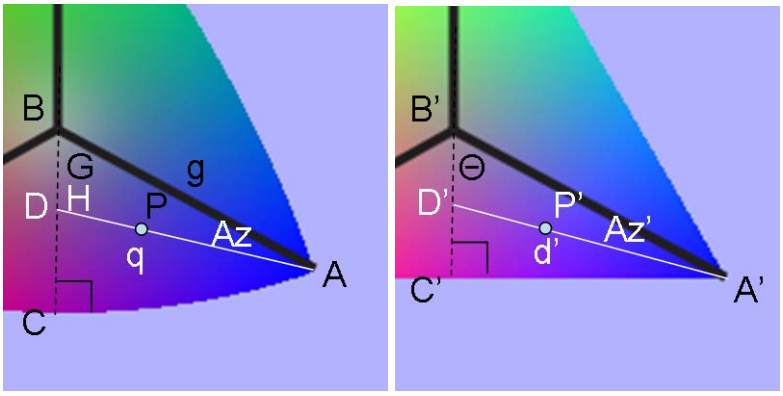

Figure 3.5: Snyder's Projection

Figure 3.5 illustrates several of the key angles and vertices used during the projection. The triangles $\triangle ABC$ and $\triangle A'B'C'$ represent the spherical and planar triangles of interest. These are associated with the underlying polyhedral face with A and A' the vertices, B and B' the centroids, and C and C' the midpoints along the edge. Our point of interest P is mapped to P' using the projection. In the net-area preserving step, the radius R of the spherical polyhedron is associated with the radius R' of the sphere circumscribing the polyhedron. The point D is intersection of a great circle arc from A through P, with arc BC. The resulting triangle, $\triangle ABD$, with angles $\angle G$, $\angle H$ and $\angle Az$, is used to determine $\triangle A'B'D'$ with the same area. $\triangle A'B'D'$ needs only be determined using the angle or azimuth $\angle Az'$, which may in turn determine the position of point D'. In the final step, the ratios between arc length $q = AD$ and edge length $d' = A'D'$ are used to position P' appropriately. It should be noted that due to the definition of the triangles, angles $\angle\Theta$ and $\angle G$ are fixed and known for the given polyhedron. These values are computed and presented within Snyder's paper [48].

With these terms defined, a discussion of Snyder's projection is readily facilitated. Recall the first step determines the approximate scaling factor between the radius R, of the sphere and the radius, R' of the sphere of which the polyhedron is circumscribed. This is uniquely determined for each platonic solid (and truncated icosahedron), and is based on areal calculations for the respective faces. For example, for the spherical or global triangle (GT) face, we observe that its area is:

$$A_{GT} = \frac{(G - \Theta)\pi R^2}{180^o}, \tag{3.1}$$

as defined through the spherical excess [22]. For example, on an icosahedron, where $G = 46^o$, and $\Theta = 30^o$, we compute an area of $A_{GT} = \frac{1}{30^o}\Pi R^2$. This is expected, given the twenty icosahedron faces, split six ways producing 120 global triangles. Dividing the surface area of a sphere, $4\pi R^2$ by 120, we get $\frac{1}{30^o}\Pi R^2$, or the area of our computed A_{GT}.

The area of the planar triangle is computed using the traditional form, adjusting the width and height with respect to its circumscribing sphere. In this case, with the face made tangent to a sphere of radius R', and therefore having side $A'B' = R' \tan g$, we observe that our planar or mapped triangle (MT) has an area:

$$A_{MT} = \frac{1}{2}(R' \tan g)^2 \sin\Theta \cos\Theta. \tag{3.2}$$

Equating 3.1 and 3.2, and applying the numbers for the icosahedron, we achieve the scaling ratio of:

$$R' = 0.9104R,$$

which will ensure overall areal preservation.

The next step is to preserve localized areal equivalence. Recall that point D is formed by the intersection of great circle arcs AP and BC. Rather than computing D geometrically, one can use $\angle H$, calculated through the spherical Law of Sines and Cosines [3, 5]:

$$\angle H = \arccos(\sin Az \sin G \cos g - \cos Az \cos G). \tag{3.3}$$

This in turn defines the area of $\triangle ABD$ as:

$$A_{ABD} = \frac{(Az + G + H - 180^o)\pi R^2}{180^o}, \tag{3.4}$$

again through its spherical excess. To associate the area of $\triangle A'B'D'$ with its circumscribing radius, and angles of interest, Snyder defines the area as:

$$A_{A'B'D'} = \frac{(R' \tan g)^2 \tan Az'}{2(\tan Az' \cot \Theta + 1)}. \tag{3.5}$$

Since we need $A_{ABD} = A_{A'B'D'}$, we can transform equations 3.4 and 3.5 to define our planar azimuth, Az':

$$Az' = \arctan(2A_{ABD}(R'^2 \tan^2 g - \cot \Theta)).$$

The final step is to position point P' along this calculated azimuth, Az' so that it preserves overall areal scale. Snyder modifies the proportionality factor from the Lambert Azimuthal Equal-Area projection so that it complies with the polyhedral face. This proportionality factor becomes:

$$f = \frac{d'}{2R' \sin(q/2)}, \tag{3.6}$$

for arc length q and edge length d'. (See Snyder [48] for more details). For any point at position z along the same azimuth Az', one obtains the ratio:

$$\rho = 2R' f \sin(z/2), \tag{3.7}$$

which in turn is used to compute x, y on the planar face:

$$\begin{aligned} x &= \rho \sin Az', \\ y &= \rho \cos Az'. \end{aligned}$$

The resulting process for calculation is as follows:

1. Using Snyder's tables and polyhedral layout, compute Az, and z from the given point's latitude and longitude.
2. Adjust Az to fall within the angular range supported by the given face's smallest symmetric region
3. Calculate q, and check that it resides on the given polygon
4. Apply the equations for calculating Az' (and unwrap accordingly) and the proportionality ratio ρ, to finally compute the mapped positions:

$$\begin{aligned} x &= \rho \sin Az' \\ y &= \rho \cos Az'. \end{aligned}$$

It should be noted that calculations presented by Snyder are reliant on the specific polyhedral layout and the x, y position and, latitude and longitude offsets for each polygon. These

constant values are visualized and listed within his paper [48]. Employing these values becomes a task in look up tables rather than generalizing for geometric shapes and points.

Though many calls to sine and cosine are required, this process results in an effective and closed form calculation of the forward projection.

3.1.4 Inversion of Projection

The inversion process finds the spherical coordinates of P given the polyhedral coordinates of P'. From the forward projection, we see that symmetric extraction and net areal scaling require nominal modification. In matching the areas of $\triangle ABD$ and $\triangle A'B'D'$, we have Az' and must compute Az. Thus, we can define the area of $\triangle A'B'D'$ as [48]:

$$A_{A'B'D'} = \frac{R'^2 \tan^2 g}{2(\cot Az' + \cot \Theta)}.$$

Setting this equal to the area of $\triangle ABD$, from equation 3.4, we observe that Az is involved both linearly, and trigonometrically, through $\angle H$'s reliance on the arccos of $\sin Az$ and $\cos Az$ by equation 3.3. Solving for Az results in a non-linear equation. Since a closed form is neither proposed nor easily determined, Snyder suggests the use of the Newton-Raphson iterative approach [39] in order to deduce an adequate value.

This approach takes the derivative of the equation and uses it to iteratively find an improved approximate solution. Conceptually, it takes uses the tangent to derive an improved approximation for finding the root, or zero, for the equation. Figure 3.6 illustrates the original curve, and an initial approximation x_0 (Figure 3.6(b)). Taking the tangent at the initial estimate, we use it to find our next approximation, x_1. We repeat until we are 'close enough' to our solution. Generally, this is determined when our change in x between iterations is trivially small.

Applying Newton's method within the inverse Snyder projection, we use the equations:

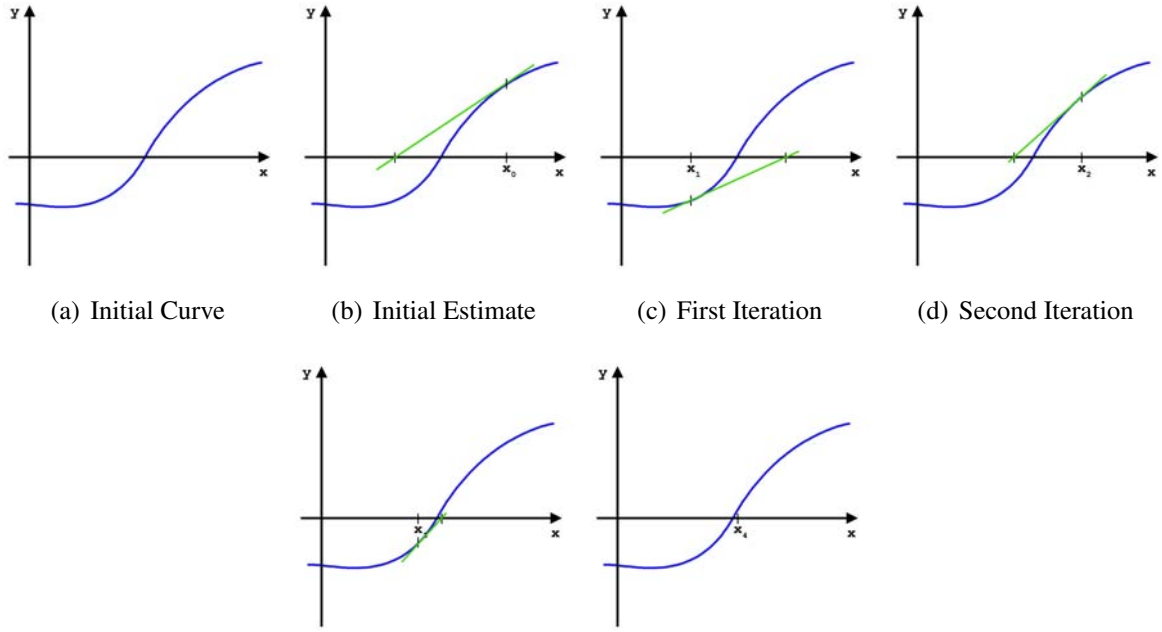

(a) Initial Curve (b) Initial Estimate (c) First Iteration (d) Second Iteration

(e) Third Iteration (f) Fourth Iteration

Figure 3.6: Newton Raphson Method. Green indicates tangent line at point of intersection.

$$g(Az) = \frac{180^o A_{A'B'D'}}{\pi R^2} - G - H - Az + 180^o \tag{3.8}$$

$$g'(Az) = \frac{\cos Az \sin G \cos g + \sin Az \cos G}{\sin H} - 1 \tag{3.9}$$

$$\Delta Az = -\frac{g(Az)}{g'(Az)}.$$

On each iteration, ΔAz is added to Az until ΔAz goes below some pre-determined threshold.

The final positioning of P along great circle arc AD is straightforward using the proportionality from equation 3.7.

These calculations are illustrated in Algorithm 1.

Due to the non-linear equation, this inverse projection requires a number of iterations to converge on a value within a required accuracy. It should be noted that some of the computations within the iteration process are often repeated and therefore redundant within a formal implementation. These repetitions must be identified and removed.

Algorithm 1 Inverse Snyder Calculation

Require: $face, Az'$
  $lat \leftarrow 0$
  $lon \leftarrow 0$
  // Convert Az' to symmetric subregion (not shown)
  // Determine initial estimate for Az
  $Az \leftarrow Az'$
  $A_G \leftarrow \frac{R'^2 \tan^2 g}{\cot Az' + \cot \Theta}$
  $\delta \leftarrow 1$
  // Iterate using Newton-Raphson
  while $\delta \neq 0$ **do**
    $F(Az) \leftarrow \frac{180^o A_G}{\pi R^2} - G - H - Az + 180^o$
    $F'(Az) \leftarrow \frac{\cos Az \sin G \cos g + \sin Az \cos G}{\sin H} - 1$
    $\delta \leftarrow \frac{F(Az)}{F'(Az)}$
    $Az \leftarrow Az + \delta$
  end while
  // Unwrap Az, so it falls in the correct symmetric region of the face (not shown)

3.1.5 Leeuwen et al.'s Slice and Dice

Leeuwen et al.'s Slice and Dice approach [57] is interesting in that it doesn't just provide a definition for a projection. Instead, it refers more broadly to a general approach for construction a variant of equal area polyhedral projections. They provide two concrete projections as examples from their underlying principle. These projections are the Parallel Small Circle and the Vertex-Oriented Great Circle. Interestingly enough, the construction of the Vertex-Oriented Great Circle approach mimics Snyder's Polyhedral Projection from Section 3.1.3, exchanging the use of the center of the face for the vertex.

In developing the Slice and Dice approach, Leeuwen et al. begin similarly to Snyder's construction. From the sphere, a spherical polyhedron is applied. The planar version of the polyhedron is then associated with the respective region on the sphere. Again, similar to Snyder's approach, from the given polyhedron - one of the platonic solids, or the truncated icosahedron - containing regular faces, the smallest symmetrical element may be extracted for projection development. This results in spherical right triangle, T and planar right triangle T'. (Recall

Figure 3.4 from Section 3.1.3).

The novel concept behind the Slice and Dice method is how the spherical and planar regions are constructed for equivalent area. To ensure areal equivalence, it is sufficient to demonstrate that spherical triangle T is mapped onto plane triangle T' with constant areal scale. Leeuwen et al. formulate this through a constructive proof, employing infinitely small regions or cells to illustrate their constant areal scale. These cells are constructed through the equal area partitioning (slice) and equal area positioning (dice) steps.

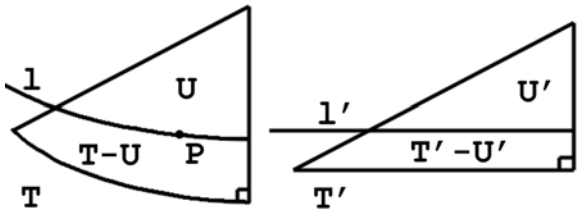

Figure 3.7: Partitioning Triangles T and T'

Initial partitioning occurs by constructing lines l and l' on the T and T' respectively. From Figure 3.7, these lines are selected so that the resulting regions maintain the ratio:

$$\frac{A_{U'}}{A_{T'}} = \frac{A_U}{A_T},$$

for $A_{U'}$, A_U the upper areas on triangles T' and T' respectively, and $A_{T'}$, A_T the areas of the given triangles.

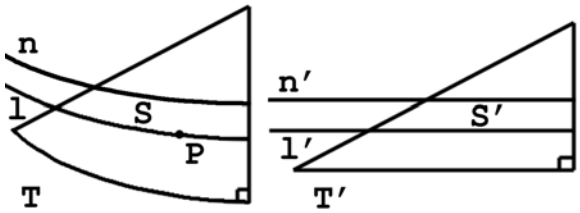

Figure 3.8: Partitioning Triangle Slices S and S'

To keep the scale constant, we ensure an infinitely thin slice will map from the spherical T

to the planar T'. Given secondary lines, n and n' on the sphere and plane respectively (Figure 3.8), as $n \to l$, we wish to have:

$$\lim_{n' \to l'} \frac{A_{U'} - A_{S'}}{A_{T'}} = \lim_{n \to l} \frac{A_U - A_S}{A_T},$$

for A_S and $A_{S'}$ the area of our respective slices. This means we wish to demonstrate that as the lines become close, the area of the slices equals. Or in other words:

$$\lim_{n' \to l'} \frac{A_{S'}}{A_{T'}} = \lim_{n \to l} \frac{A_S}{A_T}.$$

To determine the positioning of point P' on our planar triangle T, while retaining a constant areal scale, we dice our resulting triangles. As such, we construct an arc m and line m' which intersects with our points P and P' respectively. (See Figure 3.9).

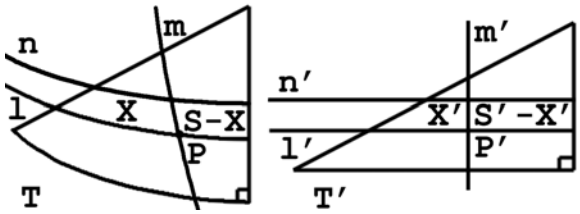

Figure 3.9: Positioning Triangle

To retain areal equivalence, we will construct these lines so that:

$$\lim_{n' \to l'} \frac{A_{X'}}{A_{S'}} = \lim_{n \to l} \frac{A_X}{A_S}.$$

We would like our area ratio to be equivalent between the compartmentalized slice regions. Once this has been accomplished, we must lastly ensure that our infinitely small region C and C' will have a constant area scale (Figure 3.10). More accurately, this means:

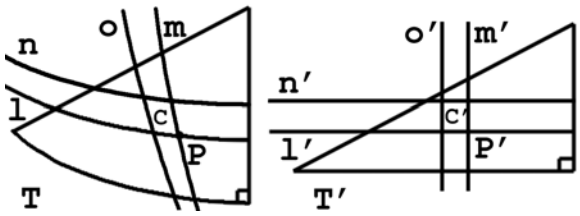

Figure 3.10: Positioning Triangle Dices X and X'

$$\begin{aligned}
\lim_{n'\to l',o'\to m'} \frac{A_{X'} - A_{C'}}{A_{S'}} &= \lim_{n\to l,o\to m} \frac{A_X - A_C}{A_S} \\
\lim_{n'\to l',o'\to m'} \frac{A_{C'}}{A_{S'}} &= \lim_{n\to l,o\to m} \frac{A_C}{A_S} \\
\lim_{n'\to l',o'\to m'} \frac{A_{C'}}{A_{T'}} &= \lim_{n\to l,o\to m} \frac{A_C}{A_T}.
\end{aligned}$$

Consequently, as the respective line pairs $n \to l$, $n' \to l'$, $o \to m$, and $o' \to m'$ approach one to the other, we result in C being mapped to C' with a constant areal scale. This illustrates how infinitely small regions on T are mapped to T' with constant areal scale, and therefore result in an equal area projection.

The next stage involves practical applications where these properties hold. Leeuwen et al. present two different approaches: the Parallel Small Circle Projection, and the Vertex-Oriented Great Circle Projection. A brief description is provided. For more in-depth details, I refer the reader to the paper [57].

3.1.6 Parallel Small Circle Projection

The parallel small circle projection employs small circle arcs for the initial slicing mechanism, and great circle arcs descending from the pole given by the small circle. To this end, right angle triangles on the spherical triangle are constructed, and more readily solved for the equal area projection. The planar triangle resultantly constructs numerous right angle triangles from the consequent straight lines. (Figure 3.11).

Point P on spherical triangle ABC is to be mapped to point P' on planar triangle $A'B'C'$.

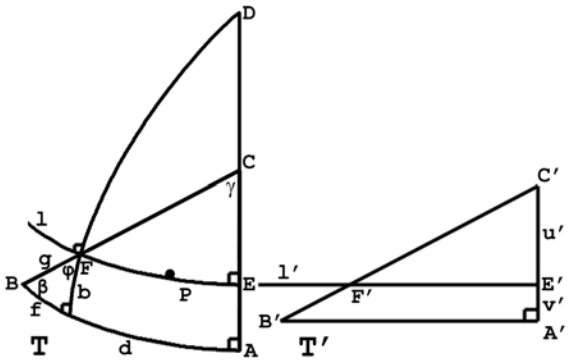

Figure 3.11: Parallel Small Circle Slice

Arc l represents the small circle parallel to arc AB, running through point P, and consequently F. Angle β comes from the original triangle, whereas angle γ is generated by the great circle arc passing through point D, the pole for our parallel circles. Through spherical formulas, and our equivalence ratios, we end up calculating our lengths u', v' via:

$$\left(\frac{u'}{u'+v'}\right)^2 = \frac{\gamma+\varphi-d\sin b}{\beta+\gamma-\frac{1}{2}\pi}.$$

Given we can construct the location of line l' from u', v', to position our point P' along l', we must dice our triangles accordingly (Figure 3.12).

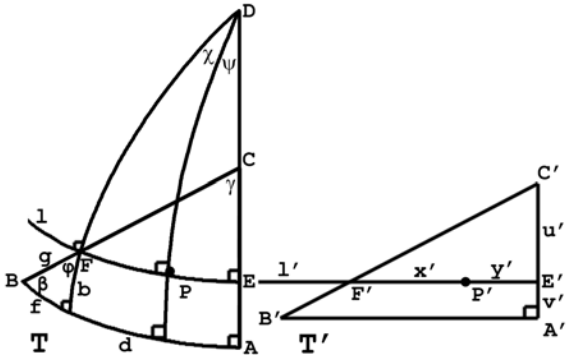

Figure 3.12: Parallel Small Circle Dice

An additional great circle arc running from D through P, generates angles χ and ψ. These, in turn, give us our values x' and y' for the positioning of P' from the equation:

$$\frac{x'}{x' + y'} = \frac{\chi}{\chi + \psi}.$$

The resulting projection eliminates angular distortion within the interior of the face - a feat unachieved by the Snyder projection. This comes at the expense of angular distortion radiating towards the boundaries of the planar faces, though still less than that attained by Snyder.

3.1.7 Vertex-Oriented Great Circle Projection

The vertex-oriented great circle projection is comparable to that of Snyder's polyhedral projection, with the importance of the vertex of a triangulated face being exchanged with the center of the face. Instead of exploiting right angles, this variant employs the fixed vertex, and the ratios en route.

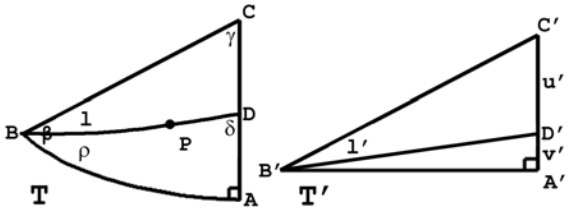

Figure 3.13: Vertex-Oriented Slice

For example, the slice component takes the great circle arc l passing through points B and P, and consequently through D, and constructs point D' from the line l' passing through B'. (Figure 3.13) The values u' and v', here, are computed by maintaining:

$$\frac{u'}{u' + v'} = \frac{\beta + \gamma - \rho - \delta}{\beta + \gamma - \frac{1}{2}\pi}.$$

In order to position point P' along line l', we simply relate the ratios of arc lengths x and y to planar lengths x' and y'. (Figure 3.14).

We accomplish this positioning, or dicing, by maintaining the ratio:

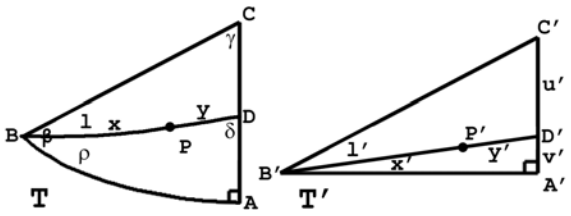

Figure 3.14: Vertex-Oriented Dice

$$\left(\frac{x'}{x'+y'}\right)^2 = \frac{1-\cos x}{1-\cos(x+y)}.$$

This projection results in a more even distribution of angular distortion across the face of the resultant polyhedron. While not entirely eliminated, the distortion is reduced greatly as compared with the Snyder projection, and on average, better than the Parallel Small Circle approach.

As noted in Leeuwen's results, and their comparison against Snyder's approach, the angular deformations are either reduced, or more uniformly distributed across the face of the triangle, thereby minimizing drastic angular changes, which induce undesirable cusps and distortions.

3.1.8 Inversion of Projection

As mentioned at the start of our discussion on projections, modern visualization relies not only on the traditional forward projection, but also with the inversion - namely the mapping from the plane to the sphere. In our case here, from the polyhedral planes to the sphere - such an inversion would be greatly beneficial, given the reduced angular distortion results from the forward projection. In this way, the Slice and Dice approach could be more accessible for use.

Unfortunately, while the paper is highly detailed on the forward projection, it is lacking in any definition for the inversion. Evaluation of the equations sees little potential for easily computable closed form approaches. As such, an intricate system of inverse trigonometric equations results. Employment of a numerical solution finding approach would be required,

and risk being highly time consuming for computation.

3.2 Subdivisions

As our method at approaching equal area projections is through the employment of subdivisions, it is worthwhile to explore one of the existing equal area approaches. Song et al.'s Small Circle Subdivision [51] is highly satisfactory at segmenting the surface of the Earth into consistently equivalent regions. Unfortunately, its approach is applied directly onto the surface of the Earth. From a visualization standpoint, planar faces tend to be employed in practise. Having a direct tie-in to this planarity would be of benefit for associating the data with the true visualization. Song et al.'s approach is briefly discussed for supplementary context.

3.2.1 Song et al.'s Small Circle Subdivision

In 2002, Song et al. [51] sought to construct an areally-uniform global grid. They employed small circles to approach their problem. As with the projections referred to in this chapter, the sphere is initially segmented through a spherical polyhedron. For their purposes, the icosahedron - limited to regular equilateral triangles - is chosen for their process.

Given the resulting spherical triangle, the objective applies a variant of the Loop subdivision [36], generating a 4-fold approach. Midpoints along the edges, or arcs, of the triangle are selected. However, instead of connecting the newly inserted vertices with great circle arcs - which would result in non-equivalent areal regions - small circle arcs are employed. (Figure 3.15).

In order to determine the small circle arcs, the three triangles, each containing an original vertex, are systematically computed. (This could actually be done in parallel). The mathematical calculation ensures that the small circle arc selected as its third bounding edge creates a region that is one-quarter the size of the original triangle. To compute this, an areal estimate

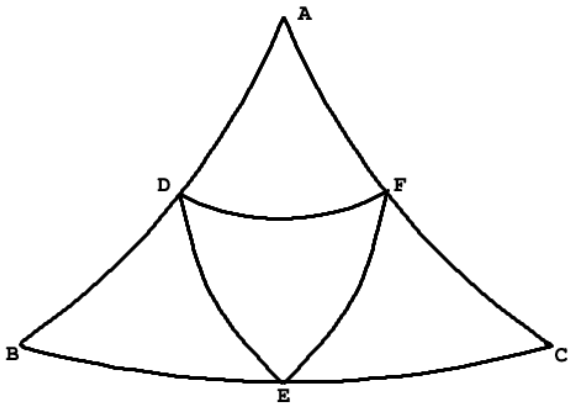

Figure 3.15: Small Circle Subdivision Face [51]

using a great circle arc is initially employed, and the difference from the lune is subtracted out. (Figure 3.16).

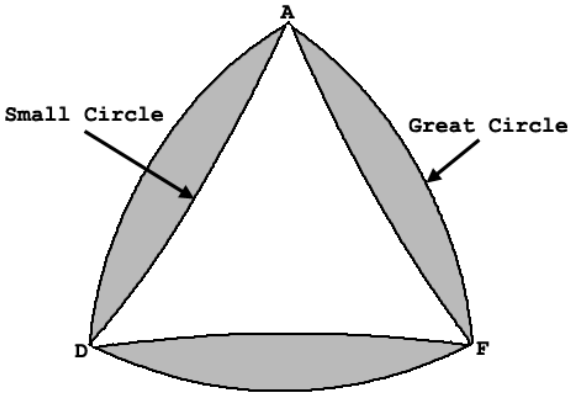

Figure 3.16: Small Circle Subdivision Lunes in Grey [51]

Once the three small circle arcs are generated, the resultant interior triangle, by deduction, must also have an area one-quarter the original size.

As this is applied iteratively, more original edges become constructions of small circle arcs. As a result, the one-quarter mathematics becomes increasingly sophisticated. A detailed discussion of the formulae are left for the reader to explore at their leisure [51].

An additional 9-fold subdivision is also presented. Consequently, it inspires us to consider such a subdivision during our own traditional approach. An evaluation of the subdivision by

Song et al. explores its compactness - or similarity of triangular shape - through six levels of subdivision. They conclude that while faces are not identical in size, the difference is fairly uniformly spread across the face. Snyder's projection, in contrast, would result in highly different triangles along the rays of the original triangular face.

As previously mentioned, while this approach is seemingly effective - reducing compactness, and establishing an equal area subdivision - its application directly on to the sphere is not of use for our visualization needs. While its 4-fold approach is interesting, it merely helps to inspire consideration when preserving area in our own experimentation. Supplementary to its lack of direct application to our own problem, the highly complex small circle arc formulae would make this implementation extremely slow, and in the long run not viable without optimization.

Chapter 4

Improving Equal Area Projections [1]

As the traditional form of constructing an equal area mapping involves the development of an projection, it is important within a computer driven visualization system to ensure a fast and robust technique is applied.

Given the popularity of the Snyder projection [40], it is important to ensure an effective forward and inverse calculation. The forward projection, as defined in Section 3.1.3, demonstrates a straightforward calculation.

The inverse Snyder projection, however, is lacking this mathematical ease of computation through the required iterative solution finding as outlined in Section 3.1.4. This is the underlying foundation for much of the computation of a geoscience system. Consequently, for the transference of planar data - upon which such computational processing occurs - to its spherical representation, it is imperative that we explore how to speed up such a process.

4.1 Snyder Optimizations

As the inversion of the Snyder Equal-Area Polyhedral projection has no known closed mathematical form, and therefore requires iterative root-finding techniques, it is necessary to optimize the calculations that occur within these iterations. For a visualization software, these inverse projection calls may be induced billions of times for each frame of information. Consequently, primitive implementations may impede the real-time requirements. Though an order of magnitude in reduction would be preferential, with such time costs, even a reduction of a constant form is helpful.

[1] Much of the content herein is adapted from Harrison et al. [29]

Notice that the aim of a real-time interactive system is the achievement of 24 to 30 frames per second (fps), or thereabouts. For an interval of 24 frames per second, this gives 0.04167 seconds per frame. For the generation of a single frame of data, there will be a constant initial set up time and a constant final wrap up time. Everything in between reflects the amount of data that can be visualized within the remaining portion of a second.

As the inversion time for data visualization increases, the number of points of data that can be projected and visualized decreases in order to maintain this real-time quality. For example, a single call of the original Snyder inverse projection can occur within 1/10,000ths of a second. Therefore, in achieving 24fps, assuming no constant overhead time, a maximum of 4,200 points of data may be inversely projected. If one wanted to simply visualize all of the 6,500 locations of Canada Post offices, this could not be accomplished on-the-fly, and without the facilitation of additional data storage, in real-time. By speeding up the inverse calculations facilitates an increased number of items to be visualized.

An icosahedron has been used throughout the optimization of this inversion process. Its uniform triangular faces offer a regularity easily suited for computer graphics applications, while its high face count more readily approximates a sphere over alternatives, such as the tetrahedron. As Snyder's projection is explicitly defined for the five platonic solids as well as the truncated icosahedron, these optimizations could be comparably adjusted for these respective base surfaces.

Three main optimizations are explored. Operation reduction, curve fitting and iteration removal are employed to speed up this inverse projection. While the most helpful optimization involves the use of curve fitting for improved initial approximating values, an initial discussion on operation reduction is presented due to its simplicity.

4.1.1 Operation Reductions

A very preliminary approach involves the recognition of repeated calculations and the respective storage of pre-computed information. Being able to identify such computations within the iterative process is particularly important due to its repetitive nature. For example, $\cos Az$ and $\sin Az$ are used in the calculation of H, and again in calculating $g'(Az)$. In turn, H is used in both the calculation of $g(Az)$ and $g'(Az)$ (See equations 3.8, 3.9). As such, these values should only be computed once within the iteration, and used accordingly. This is particularly meaningful given the extensive time it takes to evaluate sine and cosine operations.

Similarly, storing these repeated calculations into a single fixed value may assist in reducing the duration of iterations and consequently the overall computation time. For example:

$$\frac{180^o A_{A'B'D'}}{\pi R^2} - G + 180^o, \tag{4.1}$$

is a fixed value, yet it is calculated for each of the repeated calls to $g(Az)$ (Equation 3.8). As such, the net value computed once and then employed when necessary. Similarly fixed values $\sin G \cos g$ and $\cos G$, from $g'(Az)$, can be pre-determined and then used within the evaluation of $g'(Az)$. Other such repeated operations, including the conversion between radians and degrees, can additionally be simplified.

An additional speed up involves the recognition of calls to sine and cosine for the same angle value, as occurs in equations 3.8, 3.9. For example, both $\cos G$ and $\sin G$, as well as $\sin Az$ and $\cos Az$ are calculated. Under a variety of compilers, there exists a *sincos*, *sincosf* or *FSINCOS* directive enabling the simultaneous calculation of sine and cosine in the same amount of time it takes to compute sine. Alternatively, an implementation such as Wanhammar et al.'s [58], which calculates the two in the same time it takes to compute one, could be employed.

Overall, these basic operational modifications can improve the overall processing time.

Storing identified repetitions increases memory size - local to the method - for each uniquely repeated calculation (roughly three to five instances). The resultant decrease in calculations time corresponds to the number of operations. If we roughly approximate five to six operations per tracked instance, which is repeated one to two times, we reduce the floating point operations by thirty. Estimating an overall operation count of less than 1,000, this equates to a 3% reduction time. Incorporating the improvement in sine and cosine calculations, for the simultaneous calculations for G and Az, we would observe a roughly 4-5% improvement time.

4.1.2 Curve Fitting Numerical System

Ideally, and where possible, an iterative solution finding approach should be avoided as it leads to an unreliable and indeterminate increase in execution time. While a closed form is preferred, one is neither proposed by Snyder nor readily deduced by standard solution-finding methods. As such, we analyzed the data involved in the iterative process to better understand its mathematical evaluation. Notice, from the equations above, how the resulting azimuth relies entirely on the original input azimuth. Due to the decomposition of the polyhedron, there is a fixed angular range (60^o for the equilateral triangle that makes up an icosahedron face). Evaluating for each possible angle, we can find angle that is constructed as a result of the iterative process. Figure 4.1 plots this initial azimuth against the resulting azimuth that has been calculated by the iterative root-finding process.

While the result is not exactly linear, the one-dimensional smooth curve evokes the potential for a polynomial curve fitting approach. A cubic polynomial curve is fitted and visualized against the data to illustrate such a tactic[2]. To further illustrate its non-linear form, the exponential of its difference is also plotted: $y = e^{5(f(x)-x)}$. The calculated residuals of the data from this curve - scaled for visibility - are also shown. To determine the effectiveness of higher degree polynomials, curve fitting and tracking of residuals has been performed. These results

[2] At the expense of additional operations, Chebyshev polynomials [6] can be used. These will maintain stability in higher order polynomials.

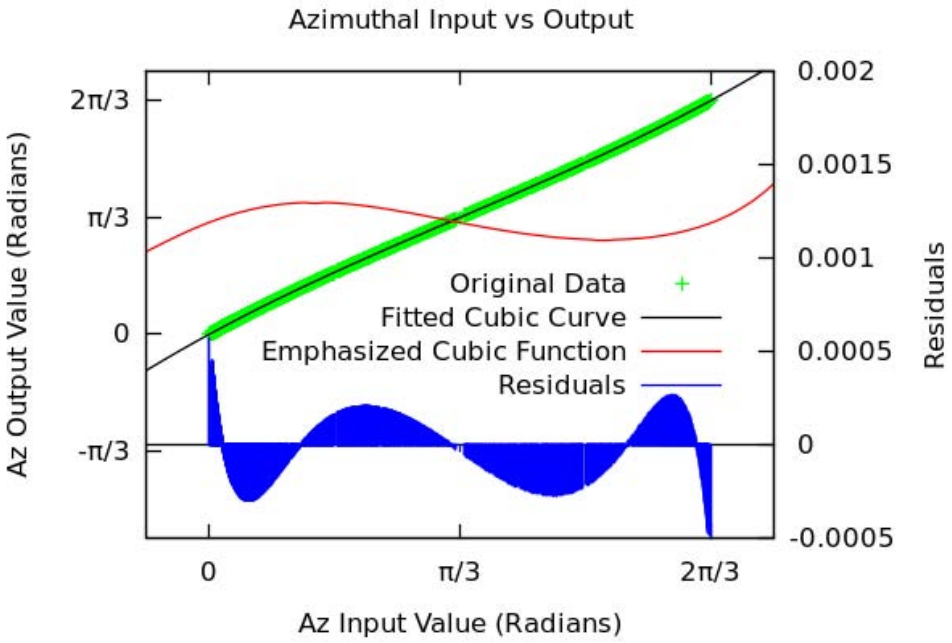

Figure 4.1: Azimuth Before Against Azimuth After Iterative Newton Raphson - Cubic Polynomial Approximation (Gnuplot 4.2)

Table 4.1: Polynomial Approximating Azimuthal Shift

Polynomial	Sum of Squares of Residuals	Variance of Residuals
Degree 1	1.19e+00	2.02e-04
Degree 2	9.27e-01	1.66e-04
Degree 3	2.30e-04	3.92e-08
Degree 4	2.06e-04	3.51e-08
Degree 5	2.51e-05	4.28e-09
Degree 6	2.20e-05	3.75e-09
Degree 7	9.06e-07	1.55e-10

are presented in Table 4.1.

These polynomials can be used to provide an improved initial estimate for the iterative system. It is expected that convergence will occur faster as a result. Horner's Rule [6] will further ensure a minimization of operations. His rule ensures that previously calculated values are effectively incorporated into the calculation. For example, the polynomial $f(x) = x^3 + 2x^2 - 3x + 1$ would be computed as:

$$\begin{aligned} f(x) &= x^3 + 2x^2 - 3x + 1 \\ &= ((x + 2)x - 3)x + 1. \end{aligned}$$

Notice how in the first equation, eight operations are required to compute $f(x)$. The second, from Horner's Rule, only takes five.

4.1.3 Iteration Removal

Rather than using the polynomial approximation computed in the prior section as the initial estimate for the azimuth, we could instead use the constructed polynomial as the definitive azimuthal value that would have resulted from the iterative process. In this way, we would use the result of the polynomial function and skip the iterative calculations entirely. As such, the computations will change from:

$Az \leftarrow poly(Az)$ // *initial approximation*

$\delta \leftarrow 1$

while $\delta > 0$ **do**

$\delta \leftarrow F(Az)/F'(Az)$

end while

to simply:

$Az \leftarrow poly(Az)$

It should be noted that a point projected with the iterative approach, and a point projected with this eliminated iteration approach will not coincide. If the offset is within a desirable precision, this error may be classified as negligible. Using a higher degree polynomial in this case will reduce the resulting error and attain increased precision. When validating this approach, analysis of the resulting positional and areal change will be necessary.

4.2 Results

Implementation of the respective inversions was completed in a Qt/C++ environment. Testing was performed on an Intel i7 quad core processor under Ubuntu 10.05. A cubic polynomial approximation was employed due to its low computation time, and relatively low residuals. The degree of polynomial selection is a balance between the error from the residuals, and the speed for evaluation. A cubic polynomial seems to best balance these attributes. An additional discussion on alternative degree polynomials is included at the conclusion of this chapter.

The original implementation was contrasted against approaches described using the fitted cubic polynomial. While basic operation reduction is useful, comparison against rudimentary implementations is not performed due to the inherent understanding that this will indeed speed up the calculations. As with any implementation, trivial optimization should be performed where possible, and therefore subsequent statistical comparison is unnecessary. By virtue of the reduction of operations, one can deduce that a speed up will have taken place.

Of more importance is the speed up of either a) reducing the iteration count through the use of an improved initial estimate with the polynomial approximation, and b) the elimination of the iterative process by wholly using the polynomial approximation. For the latter approach, an error analysis, exploring areal change and displacement is also provided.

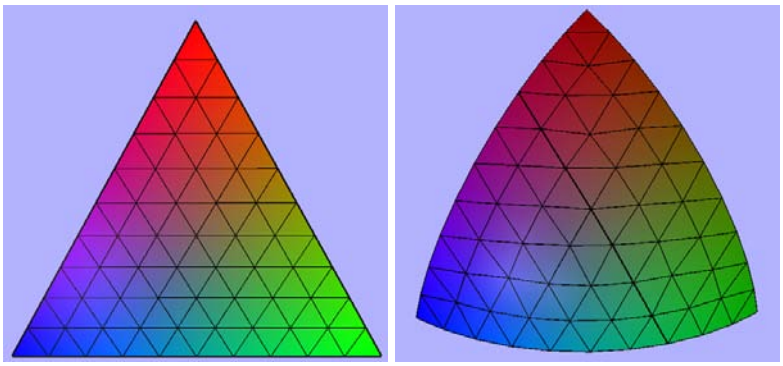

(a) Flat Face (b) Projected Face

Figure 4.2: Quality 10 of Triangulated Face

Table 4.2: Profiling Results

Method. Quality	Avg Iter.	Avg Time (s)	Std Dev.	Time Improv.	Iter Improv.
O.25	3.61	0.0082	0.0083		
I.25	2.76	0.0059	0.0081	28.05%	23.66%
E.25	0	0.0035	0.0058	57.32%	
O.50	3.70	0.0229	0.0149		
I.50	2.82	0.0205	0.0131	10.48%	23.83%
E.50	0	0.0138	0.0113	39.74%	
O.75	3.74	0.0546	0.0222		
I.75	2.81	0.0509	0.0217	6.78%	24.72%
E.75	0	0.0310	0.0167	43.22%	
O.100	3.75	0.0996	0.0329		
I.100	2.83	0.0832	0.0292	16.47%	24.50%
E.100	0	0.0524	0.0209	47.39%	

For each of the three approaches - the original, improved polynomial approximation, and iteration elimination - profiling, using *gprof* (v2.17), was performed 100 times for four resolution, or quality levels. Each of the respective resolutions refers to the number of times an icosahedron face is initially divided prior to vertex projection. For example, a 10x10 resolution generates 100 triangles (Figure 4.2) whereas a 100x100 resolution will split the face into 10,000. Resolution levels 25, 50, 75 and 100 were profiled, with each region representing approximately $4,080km^2$, $1020km^2$, $453km^2$ and $255km^2$ respectively on the Earth's surface.

Numerical results are presented in Table 4.2. The column "Method.Quality" indicates the method - original (O), improved initial estimate (I), eliminated iterative computations (E) - and quality level. The average iteration reflects the change in iteration distribution. As illustrated in Figure 4.3, as the quality improves, there are proportionally more expensive interior calls. The standard deviation of the average method time is also presented. Notice, especially with lower qualities, that the standard deviation has the same order of magnitude as the average method time. For the higher qualities, the results become more statistically significant. The percent time improvement compares the improved and eliminated approaches against the original. Similarly, the iteration improvement over the original is presented for the improved but

Table 4.3: Error Analysis of Elimination Approach

Quality	Avg Dist. Error	Max Dist. Error	Avg Area Error (%)	Max Area Error (%)
25	9.379e-05	6.193e-04	1.982e-05	1.596e-02
50	9.439e-05	6.193e-04	9.223e-06	2.195e-02
75	9.459e-05	6.193e-04	3.330e-06	3.504e-02
100	9.490e-05	6.193e-04	1.758e-06	5.073e-02

not the eliminated approach as no iterations occur.

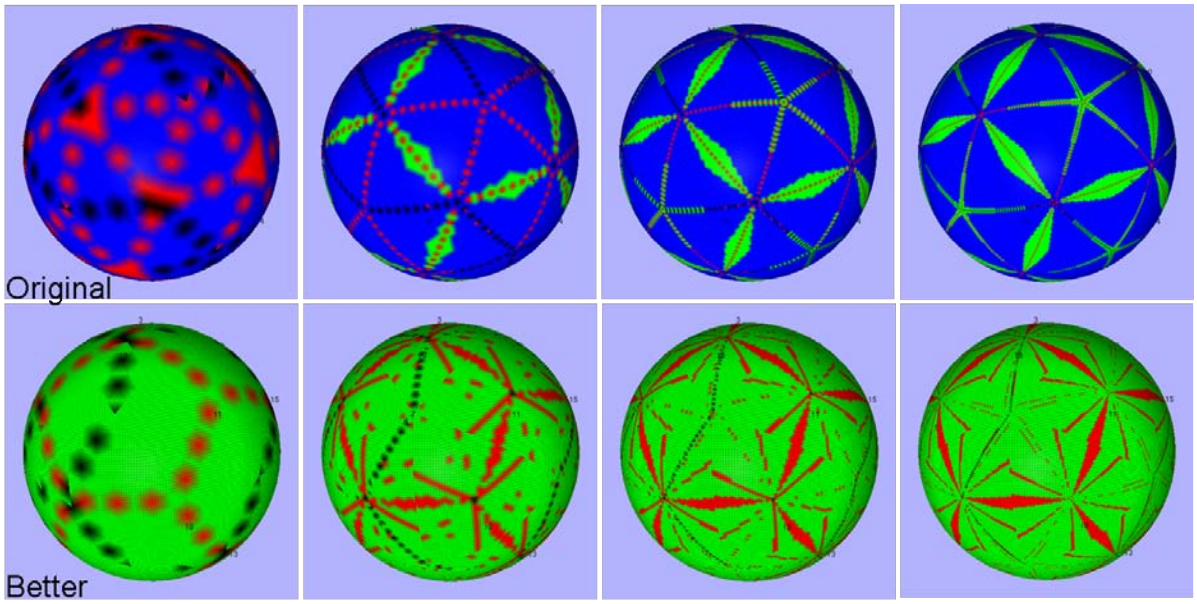

Figure 4.3: Iteration Distribution. (l-r) Quality 10, 30, 60, 100. Blue = 4, Green = 3, Red = 2, Black < 2

Error analysis for the elimination approach can be found in Table 4.3. Here, distance is an absolute value (Figure 4.4(b)). Given a point projected using the original inverse Snyder and the same point projected using the eliminated approach the distance error reflects the absolute displacement between the two points. For example, the average distance error of 9.379E-05% is the equivalent of 5.903 *m*, based on an average Earth radius of 6,369 *km*. The areal error is calculated using a quality-based triangle (Figure 4.4(a)). The error computes the triangle's change in area from its original to its eliminated projection. As the quality increases - and the

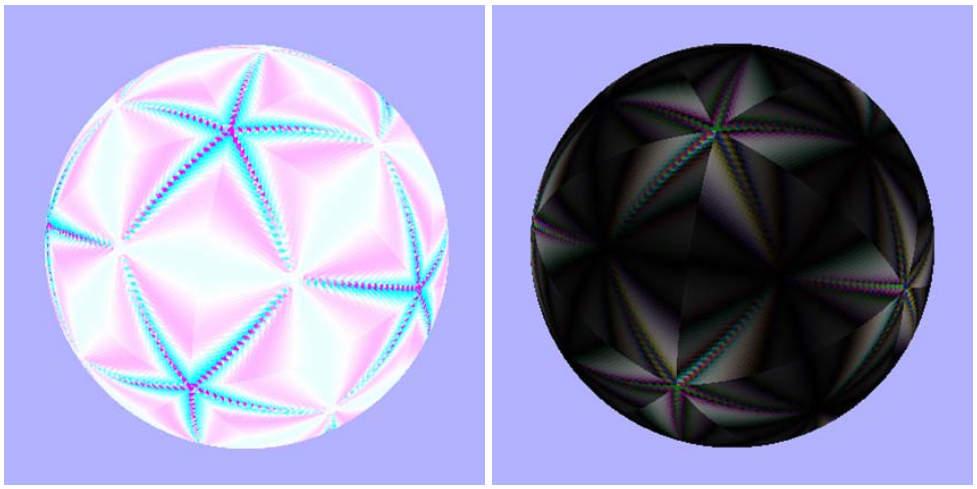

(a) Areal Change. Cyan - growth, Magenta - reduction

(b) Distance Change. RGB associated with XYZ displacement

Figure 4.4: Visualization of Skip Distortions

size of the quality-based triangles decrease - the error also decreases. For example, at quality 100, the average error of 1.758E-06% equates to 713.5 m^2 on the Earth.

4.3 Discussion

Based on the findings of the error comparison between the original projection and the eliminated approach, as illustrated in Table 4.3, the distance and areal errors are marginal. For example, with a resolution level of 100, the resulting displacement equates to a mere 5.9 m. In the case of areal changes, the shift is upwards of 0.7 km^2. Consequently, if these distances are indistinguishable from the given scale of graphical visualization, it may be computationally preferable to use this faster projection. If precision is of the utmost importance, with an error less than the eliminated approach, then the use of the iteration reduction technique offers a viable alternative, reducing the computation time by approximately 15%.

As a consequence of these increases in speed, we observe that amount of data that may be visualized, while maintaining real-time speeds, is increased. Recall Section 4.1 that at a resolution of 100, assuming no other overhead, a maximum of 4,200 data points may be

processed through the inversion call. As a result of the speed up, 5,000 points of data can be evaluated when the improved iteration approach is used, and 7,900 when the eliminated approach is used.

While a cubic polynomial was used, a higher degree polynomial may be alternatively employed. When a higher degree polynomial is used, an improved approximation occurs. Consequently, a better initial estimate is obtained, which in turn reduces the number of iterations required to converge. A separate trial using a seventh degree polynomial resulted in 45% fewer iterations, as compared with the cubic polynomial. When the higher degree approximating polynomial is used directly, vis-a-vis the eliminated approach, the errors are reduced. A simple trial concluded that for the seventh degree polynomial, the areal and distance errors were reduced by a factor of 10. This equates to a distance error of 0.59 m on the Earth's surface, as compared with the cubic polynomial's 5.9 m. Note that the seventh degree polynomial requires 14 operations to compute, whereas the cubic polynomial requires only 6. Both of these are significantly less than the calculations required for a single iteration, thereby offering an improvement in speed. If speed is of essence, employing the eliminated approach with a cubic polynomial will prove fastest and fairly efficient. If both speed and accuracy are of importance, a higher degree polynomial to better approximate the azimuth may prove more efficient. The level of precision, amount of acceptable error and need for faster calculations must be considered when selecting the appropriate approach and polynomial.

4.4 Conclusion

For large visualization systems that require real-time accurate and equal area information, an effective projection mechanism is critical. With modern data acquired in planar forms, such as field surveys or satellite images, preserving area during spherical conversion is of the utmost importance. This area preserving quality assists researchers and businesses with accurate

analysis of the respective data (Figure 4.5 illustrates our improved implementation of the inverse Snyder projection, where, for example, Greenland has the same area on the icosahedron as on the sphere). Yet this analysis, along with the visualization, relies on underlying feature outlines which may span thousands or even millions of points for a single region. Converting these feature points into their spherical coordinates becomes an exercise in accuracy and speed. Frequent use of the inverse Snyder projection, as occurs in industry, can be very time consuming to the point of limiting the system. Any form of speed up assists in achieving the real-time qualities expected by modern clients.

The largest limiting factor in the inverse Snyder projection is the non-linear calculations of the spherical azimuthal value, requiring an iterative root-finding approach. A polynomial function may improve this by providing a more accurate initial estimate. For a cubic polynomial, this reduces the number of iterations by 25%. Use of a seventh degree polynomial improves this by a further 45%. For the cubic polynomial, this results in a 15% reduction in time spent performing the inversion. Though the iterations are reduced, this improved estimate still converges to an accurate azimuth, as per Snyder's initial approach.

If some degree of error is permitted, an additional decrease in time may be achieved by using the polynomial to directly represent the azimuthal value. For the cubic polynomial, this results in a 47% reduction in time. In exchange for this large decrease, inherent errors occur. A higher degree will reduce the resulting error at the expense of its own increased calculations. In the case of the cubic polynomial, a displacement error of 5.9 m results, whereas a seventh degree polynomial will decrease this by a factor of 10. While this error may be undesirable, for primitive visualizations or calculations, these faster approaches may be beneficial.

Recall that we tend to consider the Earth as a perfect sphere, when in fact gravitational forces make it thicker around the equator, and geological features leave it uneven. Incorporating these into our model may demonstrate the fallibility of the area preserving approach taken. However without precise measurements, this spherical model of the Earth generally suffices.

Based on these non-spherical qualities, it would be of value to consider a multiresolution areal preservation that takes into consideration the precise, and imperfect shape of our home planet.

Similarly, these optimizations can be applied to other non-linear iterative projections. If the non-linear equations satisfy the conditions for Newton-Raphson convergence, then they may comparably be employed to find a representative equation. Of course it is worth confirming that for a highly detailed sampling of points, the resultant values approximate a well-defined curve, as was computed and illustrated in Figure 4.1.

Unfortunately, this is not readily applicable to the Slice & Dice projection. While a clear forward projection is explored, no inverse projection is presented. Furthermore, a clear inversion process, despite best efforts, has yet to be determined by this author. Until an inverse projection is defined, such curve-fitting optimizations are unavailable.

While we have demonstrated an approach for improving inverse projections, it does not demonstrate a full exploitation of many of the advantageous strategies possible through computer graphics, and computer science in general. As previously mentioned, the development of an hierarchical representation of the Earth's surface offers an improved process for data analysis. Using current computer graphics data structures we can enable such a feature. Attempting to support hierarchical elements, through a projective approach would be extremely involved, and computationally expensive. To this end, we progress in our discussions to explore a subdivision approach for equal area visualization.

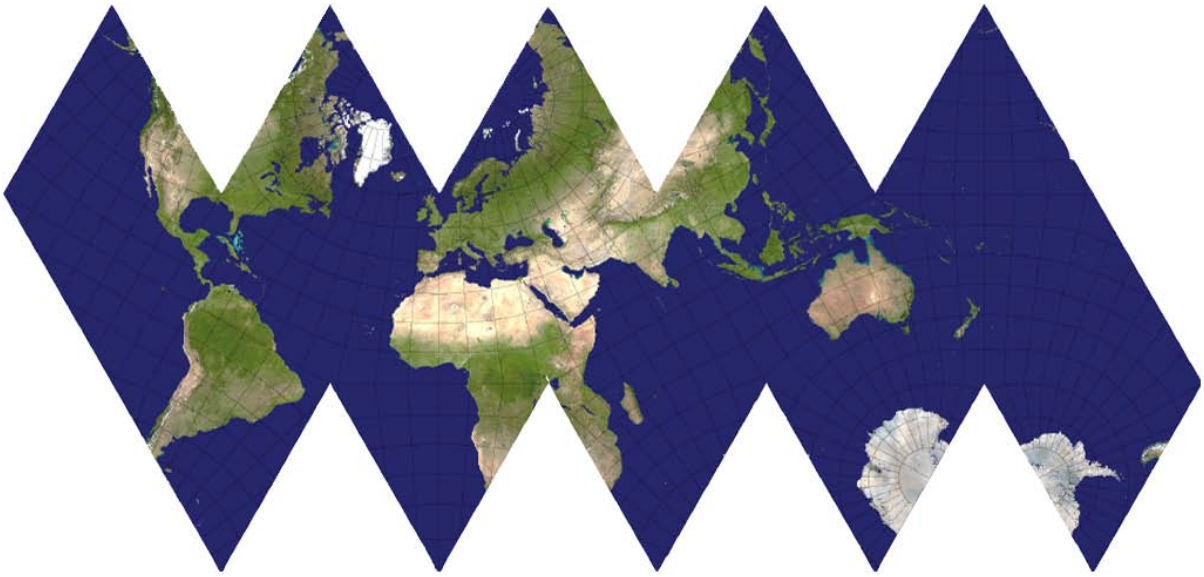

(a) Initial Texture

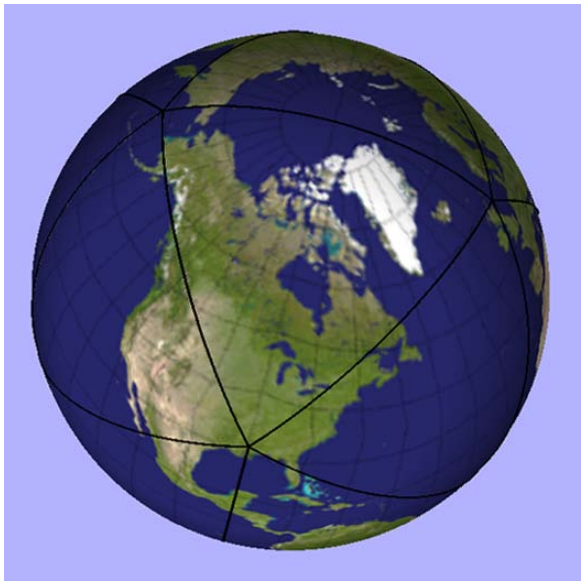

(b) Inverse Snyder

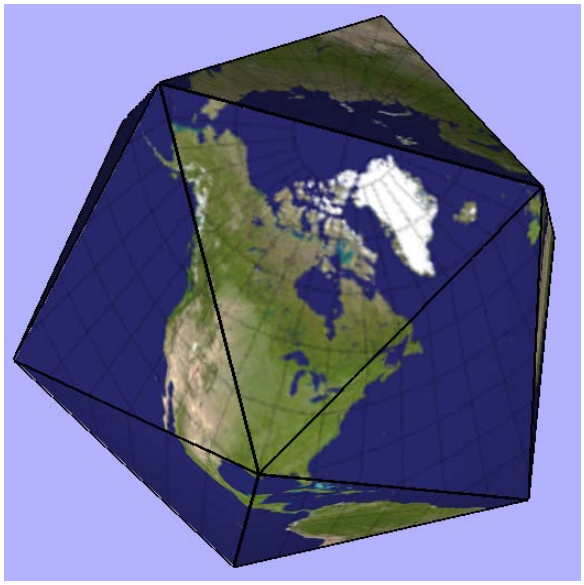

(c) 3D Icosahedron

Figure 4.5: Inverse Snyder Projection (Blue Marble, NASA [38])

Chapter 5

Arc-Based Equal Area Subdivision

The previous chapter explored the employment of projections for preserving area during the visualization process. While forward projections from the sphere to the plane are often easily computed, the inverse projections are exceedingly slower, and require optimizations. Furthermore, while these optimizations are helpful, they still do not address the underlying issue of associating visualization with the related data. Until this is addressed, the use of projections will continue to be a limiting factor for the visualization of data on the Earth's surface.

As mentioned in the Background, there are a variety of existing approaches for associating data with the visualization. Unfortunately, these either lack a hierarchical structure, or alternatively are directly applied to a spherical surface. This use of a spherical surface results in a need to transform it to the polygonal mesh required for standard computer graphics techniques for visualization. The requirement of this intermediary transformation results in essentially a mathematical replacement for the otherwise projective approach.

To make the data more meaningful, direct association with the visualization structure will offer a variety of benefits. Recognize that traditionally data, even stored in relational databases, still have a linear form. Having a three dimensional representation, or at least indexed associated representation, will facilitate an improved neighbourhood finding (e.g. when establishing where a road goes, and over what additional data sources), and through a multiresolution approach, where its parental and child information is located (e.g. a finer and coarser representation of the data at a given level of storage.)

In such a case, a cell of data may be represented by the face of the polyhedron. This will either store directly, or be indexed for reference, the data associated with that particular region on the surface. Consequently a coarse cell may span a larger region, and maintain a high level

of associated data. Higher resolution levels, wherein faces decrease in size, may instead have more specific information. Due to the hierarchical nature of this approach, each cell will be able to ascertain increased levels of information by accessing the respective data from its child cells.

To facilitate the above objectives, there are a series of elements we must construct when developing our sphere-based subdivision surface:

1. Develop a subdivision which maintains overall area of the cells throughout the subdivision process

2. Ensure overlapping and hierarchical association of the data

3. Develop an indexing for data association

The first item has undergone vigorous exploration within the works presented herein. The initial exploration evaluated how to preserve the net area of first a circle-inscribing curve, and then a spherical surface. It should be noted that an extension to the preservation of arbitrary curves has been evaluated. The natural translation to arbitrary surfaces was not explored as it significantly detracted from the objectives of the research. It is instead left as a possible future extension.

Upon acquisition of preservation of the net length, and net surface area during subdivision, ensuring the preservation of area across curve segments or cells was explored, and meets the objectives of item two. The development along the curve proved to be trivial, and is briefly discussed. The construction of a surface-based equal area subdivision underwent an experimental approach. Hypothesis, approaches taken, and the results are provided.

The third item, while important for the data association, relates more closely to the development of an indexing system. Traditional indexing often associates latitude and longitude coordinates to the given data. However, this negates the ability to easily and accurately determine neighbours and hierarchical data within the representation. An alternative approach to

retain these properties is being explored as another project within our research group. As the objective of the work herein is to explore the visualization approach, this indexing research is not discussed further, with the exception that it is still under development and will be presented within the coming year.

Consequently, the first two items are broken down in to their respective subsections. Furthermore, these are broken down into their respective curve, and surface approaches. The discussion of area preservation of arbitrary curves is explored in the conclusion of the first subsection.

5.1 Net Area Preservation

Equal area preservation ensures that a bounded region on one surface has the same area as the comparable bounded region on the second surface for all possible bounded regions, as discussed in Section 3.1. For the purposes of cartographic mapping, this is extended to incorporate a potential constant scaling factor in the area when mapped. In this way an area on the Earth may be projected to a small planar map and retain area equivalence with some scaling factor. A 1:100,000 ratio will construct a map that is physically easier to handle than a 1:1 ratio. With the advent of computers, we may be better able to accomplish a 1:1 mapping for the Earth's surface; however for the purposes of our equal area definition, we choose to incorporate the scaling factor to enable greater flexibility in visualization.

This definition means that for an infinitely small bounded region on one surface, the mapping of that bounded region on the secondary surface must have the same area. Provided this property holds across the surface of our domain, for a given region of interest, it may be represented by the summation of such infinitely small regions. As such, the mapped area of the given region will be equivalent to the area of the original region.

Again, by this definition, at its coarsest, one can take the area of the entire domain and

recognize that when mapped it will have an equivalent area. We start our investigation of area preservation by retaining this overall or net area property, harnessing advantages from working within a spherical domain. Upon completion, we extend our exploration to ensuring an infinitely small region will map to an identical area, which we refer to as local area equivalence - explored in the following section.

In Snyder's Equal Area Polyhedral Projection, and in Leeuwen et al.'s Slice and Dice approach, the representing spheres are scaled a) to ensure edges match accurately without overlap or gaps, and b) to ensure the area, with scale, is maintained. Consequently, we adopt this idea for our own purposes. Assume we have a polygon circumscribed with a circle, or a polyhedron circumscribed within a sphere. When we subdivide our shape, it will have a greater number of vertices and faces. If we construct this subdivision in such a way that the new vertices also lie circumscribed within the original circle or sphere, we have better approximated our circle or sphere. For two polyhedron circumscribed within a common sphere, the polyhedron with fewer faces will have a smaller surface area than the more refined polyhedron, which better approximates the sphere. If we take the radius of the circumscribing sphere and reduce it - and accordingly the vertices of the more refined polyhedron - in such a way as to preserve the surface area as compared with the original polyhedron, then we will have retained net area between the two polyhedra. As we continue to subdivide the surface, ensuring it resides on a common sphere, and then reducing the radius of said sphere so it matches the surface area of the original polyhedron, then we have constructed a net equal area subdivision.

This concept is discussed in details in the following subsections. We start with the simple construction of the 2D curve, and expand it to the 3D surface. In the final subsection, we explore its potential extension to non-spherical subdivisions.

5.1.1 Curves

The objective of this work is to develop an equal area subdivision for a polyhedron that will approximate a sphere. Beginning development with the simpler 2D curve case clarifies the approach, while exploiting its simplicity for overall development.

For the 2D curve case, we start with a polygon that is circumscribed within a circle (Figure 5.1). The circumscribing circle, C, we denote as having origin at point O, and a radius of length R. As such, this curve is closed, does not self-intersect, and is not a degenerate (e.g. neither a point, nor a single line). Consequently we have a minimum of three distinct points defining our coarse curve.

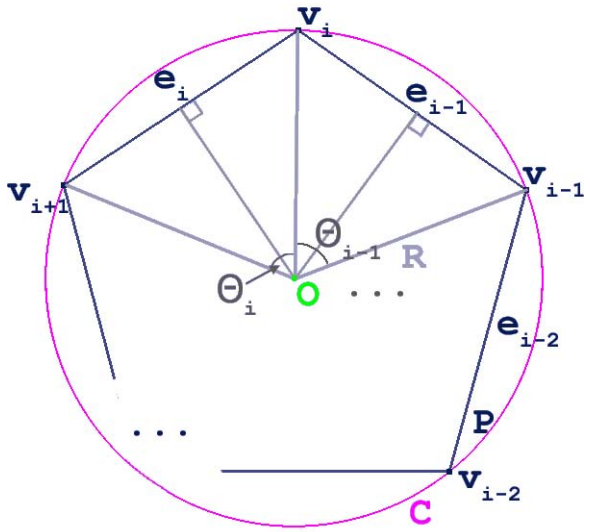

Figure 5.1: Global Curve Definition

The polygon, has n vertices, namely $v_0, ..., v_{n-1}$ where $n > 2$. Edges of the polygon are the collection $e_0, ..., e_{n-1}$, with $e_i = (v_i, v_j)$, $j = (i + 1)\%n$, the two vertices composing the edge. Furthermore, let Θ_i be the angle $\angle v_i O v_j$, where $0 < \Theta_i < 180^o$.

Since the polygon is circumscribed within C, then we have:

$$\|v_i - O\| = R, 0 <= i < n.$$

As such, we can define all vectors v_iO as having the same magnitude (namely R), and that Θ_i is simply the rotation that takes vertex v_i to v_j about the origin O.

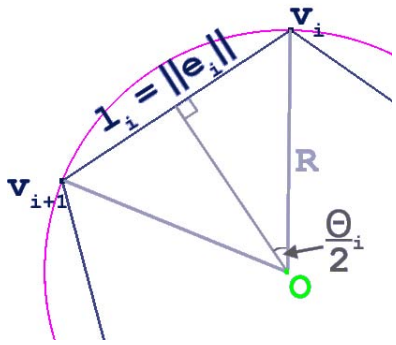

Figure 5.2: Global Curve Edge Length

Thus, we can write the edge length of edge e_i as (Figure 5.2):

$$l_i = 2R\sin(\frac{\Theta_i}{2}).$$

Then, the sum of all the edge lengths, or the perimeter of the polygon, can be written as:

$$\begin{aligned} L &= \sum_{i=0}^{n-1} \|e_i\| \\ &= \sum_{i=0}^{n-1} l_i \\ &= \sum_{i=0}^{n-1} 2R\sin(\frac{\Theta_i}{2}) \\ &= 2R\sum_{i=0}^{n-1} \sin(\frac{\Theta_i}{2}). \end{aligned}$$

Recall that our objective is to define a subdivision which ensures all vertices remain circumscribed on a circle. Geometrically, we can define such a subdivision by:

1. Insert a midpoint along each of the edges e_i thereby constructing two new edges for each original edge

2. Projecting the inserted vertex to the original circumscribing circle.

For our purposes, this second repositioning step can be simply chosen. If we reposition the point, p so that its vector $(p-O)$ has magnitude R, then we have constructed such a continuously circumscribing subdivision.

Geometrically, we can merge these two steps by taking the intersection of the bisector of vectors $(v_i - O)$ and $(v_j - O)$ with the circle arc between vertices v_i and v_j as our point p, and then replacing edge e_0 with the two edges, v_i to p, and p to v_j.

Given the above subdivision definition, let us define polygon P_1 as the original polygon, and P'_1 as the subdivided polygon. These two polygons have the same radius, R, but P'_1 has twice as many edges, while still being circumscribed on circle C. Since it better approximates the circle, and therefore has a larger perimeter, it must be modified so that its perimeter is equivalent to that of polygon P_1.

Recall from our prior discussion that the length L for a circumscribing polygon is defined as:

$$L = 2R \sum_{i=0}^{n-1} \sin(\frac{\Theta_i}{2}).$$

Let L_1, and L'_1 represent the perimeters for our polygons P_1 and P'_1 respectively. Then $L_1 \neq L'_1$ as previously discussed. To have them equivalent, we will scale the radius for P'_1, constructing polygon P_2 with perimeter $L_2 = L_1$. Therefore, we have:

$$\begin{aligned} L_1 &= L'_1 \\ \Rightarrow 2R_1 \sum_{i=0}^{n-1} \sin(\frac{\Theta_i}{2}) &= 2R'_1 \sum_{i=0}^{2n-1} \sin(\frac{\Theta'_i}{2}) \\ \Rightarrow R'_1 &= \frac{L_1}{2\sum_{i=0}^{2n-1} \sin(\frac{\Theta'_i}{2})}. \end{aligned}$$

Thus, in order to have $L_1 = L'_1$, we would have to adjust polygon P'_1 so that it lay on a circle of radius R'_1. This is the circumscribing circle to which P_2, constructed from P'_1, would reside within.

To simplify calculations, if we were to track the edge lengths of P'_1 during the subdivision process, then for length L'_1, we have:

$$L'_1 = 2R \sum_{i=0}^{2n-1} \sin(\frac{\Theta'_i}{2}).$$

To define our scaling radius R'_1, which will become our radius R_2 for polygon P_2, we simply have:

$$\begin{aligned} R'_1 &= \frac{L_1}{2 \sum_{i=0}^{2n-1} \sin(\frac{\Theta'_i}{2})} \\ &= \frac{L_1}{\frac{L'_1}{R}} \\ &= \frac{L_1 R}{L'_1}. \end{aligned}$$

Thus, upon completion of constructing P'_1, we can construct P_2 by adjusting all vertices so that they lie at a radius of R'_1 away from the origin O, rather than a radius of R_1. This can be accomplished by repositioning each vertex v_i along its respective normal vector with respect to the origin O. The resulting polygon, P_2 will then have the same perimeter as the original polygon P_2, though its circumscribing circle has been reduced to reflect the improved circle approximation through the increased number of edges.

To this end we have constructed a global equal area, circle-converging subdivision. The above can be applied to any subsequent subdivision. Upon midpoint construction and repositioning, the vertices then become scaled along their respective normals to achieve a radius of:

$$R_{i+1} = \frac{LR_i}{L'_i},$$

where L is the original perimeter, L'_i is the perimeter for the subdivided polygon, and R_i is the radius of the prior polygon, P_i. The first three such subdivision are illustrated in Figure 5.3. The initial triangle clearly has a different size circumscribing circle than the third iteration of

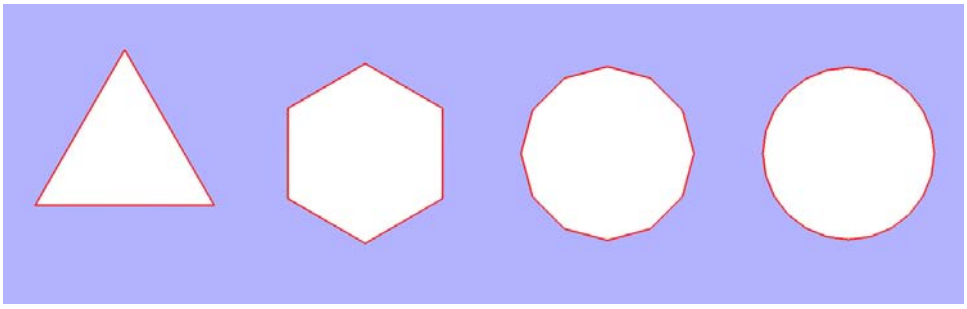

Figure 5.3: Iterative steps in the length-preserving circular subdivision

the subdivision. However the net perimeter remains the same. At the limit, the polygon will approach a circular shape with a total length equivalent to that of the original polygon.

At this point we have defined a constructive perimeter-preserving subdivision for curves. The steps involved in performing this subdivision are illustrated in Figure 5.4, and are as follows:

1. Compute the original polygonal edge length
2. Divide the polygon, in such a way as to generate more vertices along the edges
3. Reposition all vertices such that they lie on a common circle (this can most easily be performed by treating the points like vectors and normalizing their positions, or by geometrically finding the intersection between the circular arc and the bisector of the original vertices)
4. Compute the new polygonal edge length - can be calculated during the previous step
5. Compute the radius (as defined above), and scale these new vertices along their normals, by the calculated radius

At the end, we achieve equal area net edge length. Preserving the overall, or global perimeter of the curve ensures that mapping regions between subdivision resolutions will be retained.

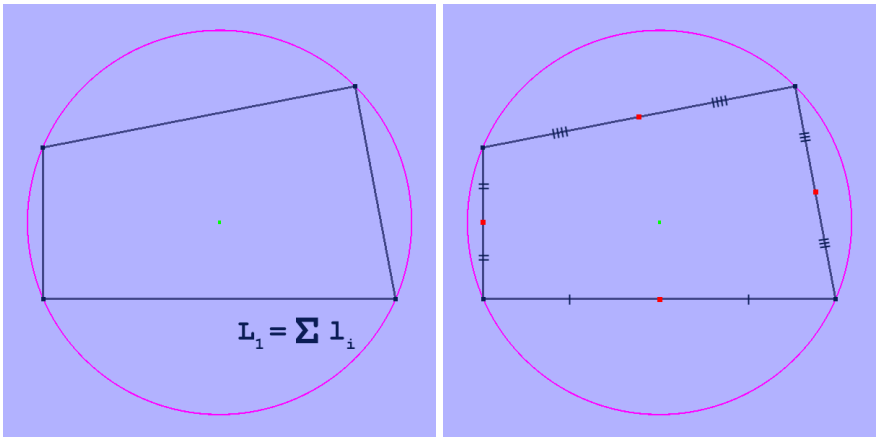

(a) 1. Compute Perimeter (b) 2. Insert Vertices

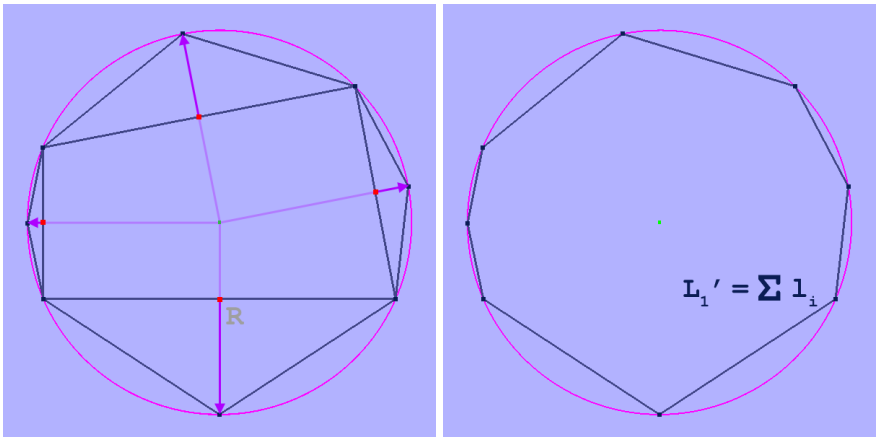

(c) 3. Reposition on Circle (d) 4. Compute Perimeter

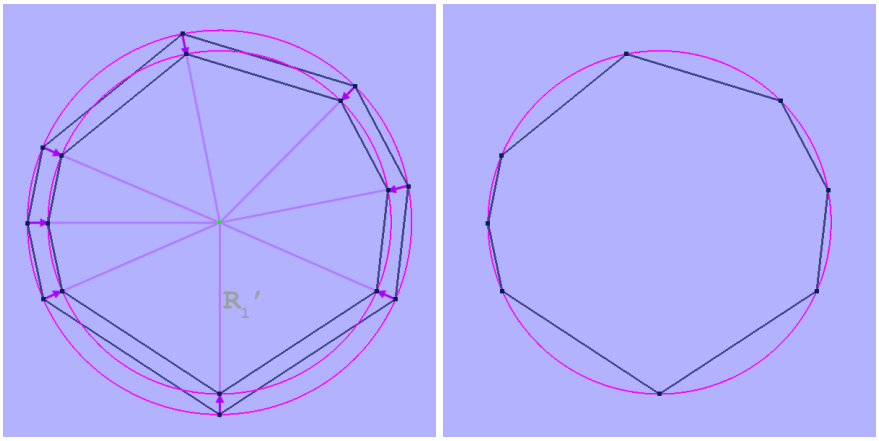

(e) 5. Reposition on Updated Circle (f) 6. Final

Figure 5.4: Steps for Global Curve Preservation

At this stage we have yet to demonstrate areal equivalence. Instead, we have simply retained net edge length equivalence upon a polygon. As we continue to subdivide the polygon, we continue to bring the vertices to a circle. Since the curve converges to a circle, it is inherently smooth.

While this approach begins to meet the needs of this research, its applicability to arbitrary curves is marginal. The main drawback with this global perimeter preservation is that the vertices of this initial convex curve must lie on a common circle. While this is to the benefit of our Earthly application, it restricts the process from being applicable to generic initial shapes. Applying this to general curves requires a modification to the approach. Such a modification is presented in Section 5.1.3, discussing arbitrary polygon equal area subdivision.

5.1.2 Surfaces

From a conceptual level, transitioning from the two dimensional curve to the three dimensional surface is trivial. A similar process of subdividing the mesh, and repositioning based on a common sphere may also be applied. However the mathematics of such a transformation must be adjusted accordingly.

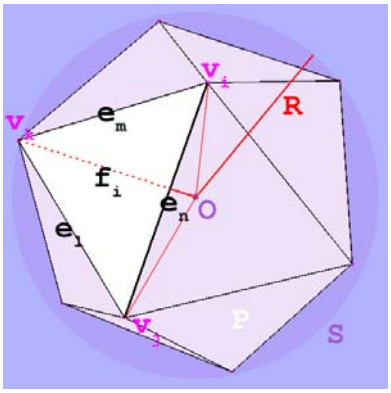

Figure 5.5: Global Surface Definition

First, we must define the geometric objects upon which we are working with. For the three dimensional case, we have a sphere, S of radius R centered around origin O (Figure 5.22). Our

surface, which is circumscribed within S, is defined as a three dimensional closed mesh. In other words, vertices v_i are positioned along the surface of the sphere, retaining the property from the two dimensional case that:

$$\|v_i - O\| = R,$$

for all v_i. Pairings of vertices are joined together by edges e_i. Since the vertices lie on the sphere, these edges act as chords through S. These edges, in turn, bound faces, f_i. It is the summation of the area of the faces upon which we are primarily concerned. For the purposes of our work, the polyhedron, P generated by such a mesh is convex, has genus 0, does not self-intersect, and has neither degenerate edges nor faces (i.e. neither the length of an edge, nor the area of a face equals zero). The justification for this simplification is that the approximation of the Earth's surface, for the time being, will not self intersect, nor have a genus greater than 0[1]. Furthermore, let us assume all faces are triangles. Within a computer graphics context, this is not unreasonable, and indeed can be generated from any existing mesh through triangulation, such as Delaunay's triangulation [15].

Therefore, let us take a sample face, or triangle, with bounding vertices A, B, and C as illustrated in Figure 5.6 lying on S. Notice M is the point on chord $\overline{BC}$ nearest to A. Therefore $\overline{AM}$ is perpendicular to $\overline{BC}$, and their values can be used to compute the area for triangle ABC.

To find such an area, the traditional approach involves the determination of formula for the area of a triangle, namely:

$$A = \frac{hw}{2},$$

with h the height, and w the width. For our triangle ABC, we can denote the height:

[1] A realistic representation of the Earth's surface would not be limited by these properties. However, such a representation would not be modeled after a sphere, as it is not perfectly spherical. Indeed, at the very least, the Earth is thicker around the equator than other regions due to magnetic and gravitational forces.

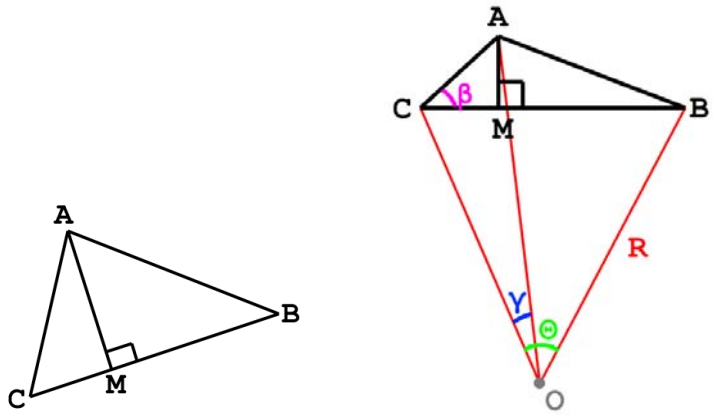

(a) Top View (b) Side View

Figure 5.6: Triangle on the Sphere with Variables Marked

$$\begin{aligned} h &= ||\overline{AM}|| \\ &= ||\overline{AC}|| \sin\beta \\ &= 2R \tan\frac{\gamma}{2} \sin\beta, \end{aligned}$$

where angles $\beta = \angle ACB$, and $\gamma = \angle AOC$. Then the width may be computed as:

$$\begin{aligned} w &= ||\overline{BC}|| \\ &= 2R \tan\frac{\Theta}{2}, \end{aligned}$$

where angle $\Theta = \angle BOC$. Thus, we have the area of a triangle on the sphere as:

$$\begin{aligned} A &= \frac{hw}{2} \\ &= \frac{1}{2}(\overline{BC} * \overline{AM}) \\ &= \frac{1}{2}(2R \tan(\frac{\Theta}{2}))(2R \tan(\frac{\gamma}{2}) \sin(\beta)) \\ &= 2R^2 \tan(\frac{\Theta}{2}) \tan(\frac{\gamma}{2} \sin(\beta)). \end{aligned}$$

If we summate the area of all triangles across the mesh, we get the surface area of the mesh as follows:

$$\begin{aligned} SA &= \sum_{i=0}^{n-1} A_i \\ &= \sum_{i=0}^{n-1} 2R^2 \tan(\frac{\Theta_i}{2}) \tan(\frac{\gamma_i}{2}) \sin(\beta_i) \\ &= 2R^2 \sum_{i=0}^{n-1} \tan(\frac{\Theta_i}{2}) \tan(\frac{\gamma_i}{2}) \sin(\beta_i). \end{aligned}$$

In order to generate a more refined mesh which more closely approximates the sphere, we must select a subdivision process which will increase the number of triangular faces across the mesh, and whose vertices yet again reside on a common circumscribing sphere. Let us take, for example, Loop subdivision, wherein vertices are added to the midpoints of each of the edges, and they are in turn are connected in such a way as to produce four triangular faces - three exterior faces composed of an original vertex, and the mid points of the adjacent edges constructing each of the three exterior faces, and one interior face composed of the three midpoints (Figure 5.7). The Loop subdivision can be adjusted for our purposes by only repositioning the newly introduced vertices so that they lie on the same circumscribing sphere. Similar to the 2D case, these added vertices can be positioned at a radius of R away from the origin O along the vector $v_i - O$, for each new vertex v_i.

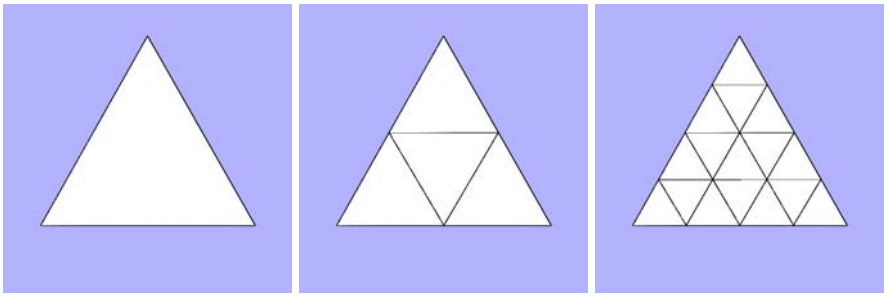

(a) Starting Face (b) Level 1 Face (c) Level 2 Face

Figure 5.7: Flattened Loop Subdivision

As with the two dimensional case, at this point we have a polyhedron, P'_1 which better ap-

proximates the sphere, as compared with the original polyhedron P_1, while still circumscribed within a sphere of identical radius. Consequently, the surface area of P'_1 will be greater than that of P_1. Similar to the two dimensional case, if we can scale our polyhedron P'_1 so that it still resides within a sphere, but of a smaller radius R_2, then we will have constructed a more refined polyhedron, P_2 which better approximates a sphere, while preserving the surface area as from the original polyhedron.

Since we want the surface area SA_1 and SA'_1 to be equivalent, we determine how best to adjust the radius. Mathematically, we observe that:

$$\begin{aligned} SA_1 &= SA'_1 \\ \Rightarrow 2R_1^2 \sum_{i=0}^{n-1} \tan(\frac{\Theta_i}{2}) \tan(\frac{\gamma_i}{2}) \sin(\beta_i) &= 2R_1'^2 \sum_{i=0}^{n'-1} \tan(\frac{\Theta'_i}{2}) \tan(\frac{\gamma'_i}{2}) \sin(\beta'_i) \\ \Rightarrow R'_1 &= \sqrt{\frac{SA_1}{2} \sum_{i=0}^{n'-1} \tan(\frac{\Theta'_i}{2}) \tan(\frac{\gamma'_i}{2}) \sin(\beta'_i)}. \end{aligned}$$

Thus, comparable to the two dimensional case, in order to have $SA_1 = SA_2$, we would adjust polyhedron P'_1 so that it lay on a sphere of radius R_2. This is the circumscribing circle to which P_2, constructed from P'_1, would reside within.

To simplify calculations, if we were to track the areas of the triangles within P'_1 during the subdivision process, then for net surface area SA'_1, we have:

$$SA'_1 = 2R_1^2 \sum_{i=0}^{n'-1} \tan(\frac{\Theta'_i}{2}) * \tan(\frac{\gamma'_i}{2}) \sin(\beta'_i).$$

To define our scaling radius R'_1, which will become our radius R_2 for polyhedron P_2, we have:

$$\begin{aligned} R'_1 &= \sqrt{\frac{SA_1}{2} \sum_{i=0}^{n'-1} \tan(\frac{\Theta'_i}{2}) * \tan(\frac{\gamma'_i}{2}) \sin(\beta'_i)} \\ &= \sqrt{\frac{SA_1}{\left(\frac{SA'_1}{R_1^2}\right)}} \\ &= \sqrt{\frac{SA_1 R_1^2}{SA'_1}}. \end{aligned}$$

Therefore, we can construct polyhedron P_2 from subdivided P'_1 by adjusting all vertices so that they lie at a radius of R'_1 away from the origin O, rather than a radius of R_1. As with the curve case, this can be accomplished by repositioning each vertex v_i along its respective normal vector with respect to the origin O. The constructed polyhedron P_2 will have the same surface area as the original polyhedron P_1. However, it will better approximate the surface of a sphere through its increased number of faces, despite its smaller circumscribing spherical radius.

Resultantly, we have constructed a definition for preserving net equal area sphere subdivision. The above can be applied to additional subdivision levels by scaling based on the calculated radius of:

$$R_{i+1} = \frac{\overline{SAR_i^2}}{SA'_i},$$

where SA is the original surface area, SA'_i is the perimeter for the subdivided polyhedron, and R_i is the radius of the prior polyhedron, P_i. As a polyhedron is increasingly subdivided through this approach, at the limit, the surface will approach a sphere with a total surface area equivalent to that of the original polyhedron.

The steps, therefore, for performing this subdivision are as follows:

1. Compute the original polyhedral surface area
2. Subdivide the polygons in such a way as to generate more vertices along the faces
3. Move all vertices such that they lie on a common sphere (again, this can most easily be performed by treating the points like vectors and normalizing their positions)
4. Compute the new total surface area
5. Compute the radius (as defined above), and scale these new vertices along their normals, by the computed radius

These are further illustrated in Figures 5.8 to 5.10.

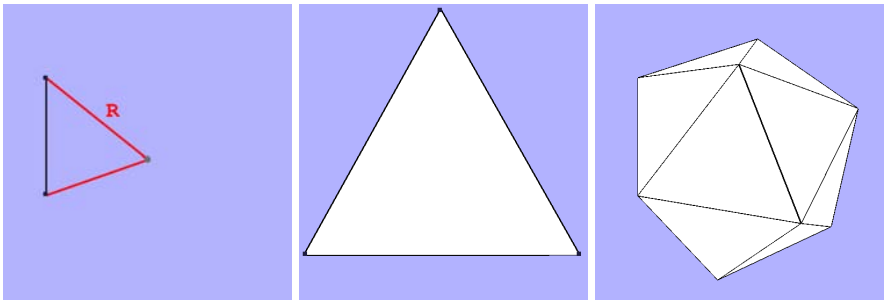

(a) Side (b) Front (c) Whole

Figure 5.8: Steps for Global Surface Preservation - Initial Setup

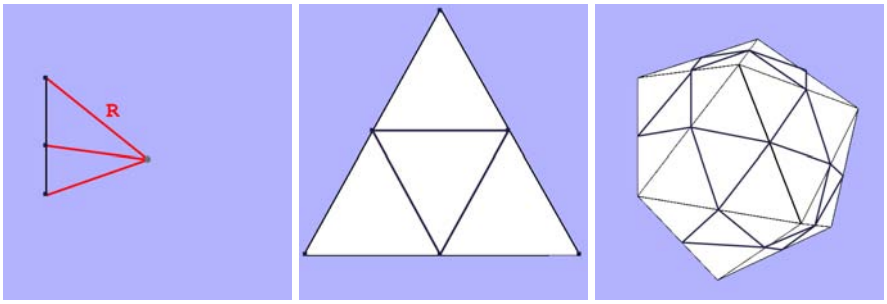

(a) Side (b) Front (c) Whole

Figure 5.9: Steps for Global Surface Preservation - Refine Surface

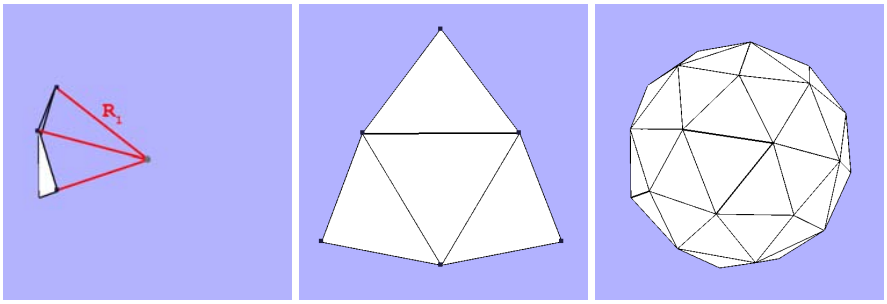

(a) Side (b) Front (c) Whole

Figure 5.10: Steps for Global Surface Preservation - Reposition on Sphere

Consequently, we have achieved net surface areal equivalence. As with the two dimensional case, the remaining problem is to find an adequate subdivision in order to retain equal areas across all faces of the polyhedron at a given subdivision level. Before we explore this, a brief exploration of an extension of net area preservation on arbitrary curves and surfaces is presented.

5.1.3 Global Preservation Applied to Arbitrary Curves

Digressing slightly from the main objective of this work is the immediate potential for an extension of the above net length preserving subdivision. Given the potential to preserve the total length of a curve, this might be applicable to curves whose intensions are not to converge to a circular shape. Instead, might there be a way of simultaneously satisfying the smoothness of a subdivided curve while preserving its total length. Potential applications for this fall within the realm of CAD/CAM and architecture for those who have fixed supplies or regions.

Within traditional subdivisions which result in a smooth curve, it can be noted that all possible tangents can be found on a circle. Indeed, as one traverses along the curve, the respective tangents can be represented by the traversal of the tangents of a point along a circle, as illustrated in Figure 5.11. As a point on the curve (Figure 5.11(b) moves from position 1 to position 4, the tangency can be found along the circle as illustrated in Figure 5.11(a). The greater the curvature, the faster the point on the circle travels. Upon a change in curvature (e.g. transition from convex to concave), the point on the circle switches places with its radial opposite, and continues on its way. This effect is illustrated in Figure 5.11 as the point transitions through position 3. As the curvature changes, the tangency on the circle must switch to its radial opposite. From this ability to represent curves as tangent transitions along a circle, one approach for designing a subdivision is the attempt at constructing a smooth transition from one circular arc to another.

Exploiting this effect, the exploration of circular arcs along a B-spline curve was researched.

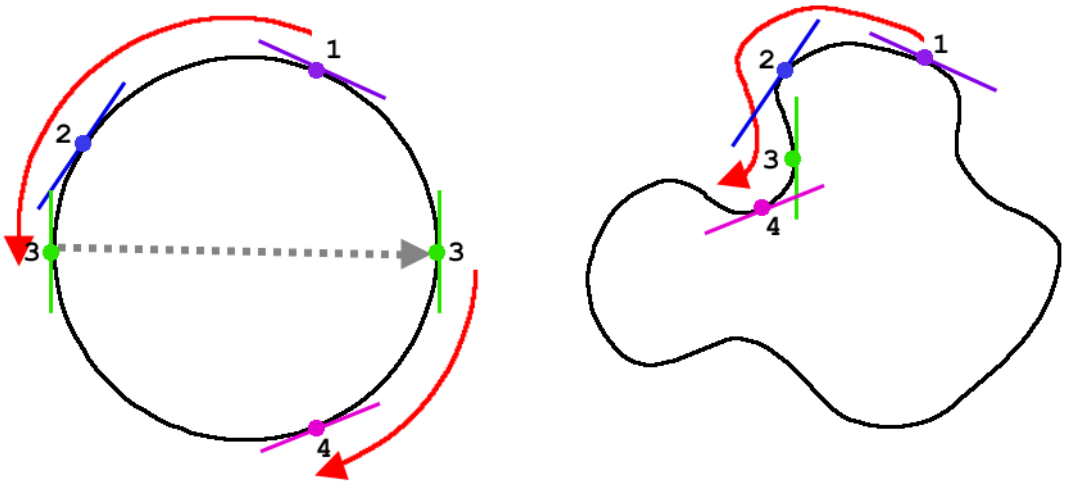

(a) Circle of Tangents (b) Curve of Interest

Figure 5.11: Circle Arc Transition

Figure 5.12 illustrates three stages of subdivision of an initial curve. As with the prior work, a closed curve is employed to simplify the equations and avoid determination of end conditions. In each of the stages, circular arcs in pink are overlaid. These arcs are constructed by taking each set of three consecutive vertices and constructing the arc between them. Notice that initially the circle arcs between each of the three consecutive vertices are very large, and greatly displaced from the line segments of the coarse curve. However, as the subdivision generates a finer approximation of a smooth curve, the circle arcs reduce in size until they overlay the smooth curve. This illustrates the above proposition regarding the transition tangents along a curve resembling the transition of the tangents of a point traveling on a circle.

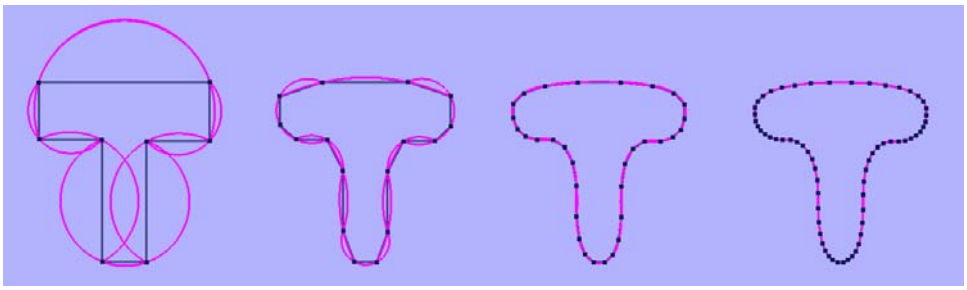

Figure 5.12: Cubic B-Spline Subdivision. In pink, the circle arcs of each set of three consecutive vertices are drawn.

The immediate potential of such a use of circular arcs is a potential to improve curve drawings. Instead of continuing to refine a subdivided curve until it easily resembles a smooth curve, we can instead employ a coarser representation connected by circular arcs. A particular threshold of arc displacement from the given points may be useful, but is left to the reader to evaluate further.

An alternative potential of this observation is the use of a subdivision generation. If we were to calculate the circular arcs generated by a coarse representation of a curve, by averaging out the circular arcs over the subdivision process, a smooth curve might be achieved.

5.1.4 Subdivision Definition

In order to construct a subdivision, we must insert and reposition a collection of new vertices, which in turn generates the new edges. The repositioning must ensure that a better approximation to the desired properties - in our case a smooth curve - is being respected. For our purposes, let us explore a single edge of a coarse polygon (Figure 5.13(a)). A similar construction will take place for all edges around the base curve.

Given the edge, e, we recognize, for our closed curve, neighbouring edges e_{-1} and e_1. Between these three edges are four consecutive vertices v_0, v_1, v_2, v_3 specified in counter-clockwise order and we define $e_{-1} = (v_0, v_1)$, $e = (v_1, v_2)$ and $e_1 = (v_2, v_3)$. Since a circle can be uniquely defined from three points, the two sets of triplets of consecutive vertices each are employed to generate a circular arc. In other words, v_0, v_1, v_2 defines our first circle, C_0 with origin O_0 and radius R_0, while v_1, v_2, v_3 define our second circle, C_1 with origin O_1 and radius R_1. Arcs are then defined as the counter-clockwise angle starting at v_0 and v_1 and ending at v_2 and v_3 respectively for the two circles.

Notice how these two circular arcs pass over our edge e of interest. To generate and position our new vertex, which will split edge e into two subsequent edges, we determine the average of the midpoint of these two arcs. In this way, we aim to average the two neighbouring circle

arcs, and thereby have a smooth transition between them.

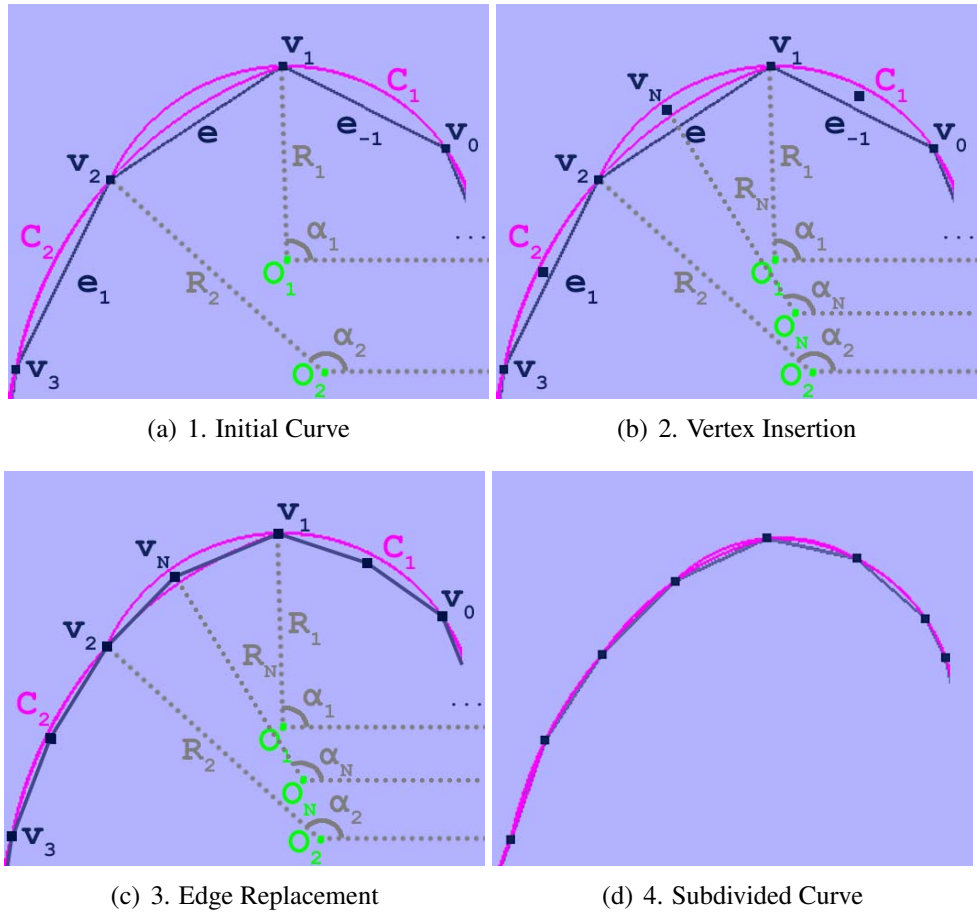

(a) 1. Initial Curve (b) 2. Vertex Insertion

(c) 3. Edge Replacement (d) 4. Subdivided Curve

Figure 5.13: Arbitrary Curve Construction

To begin, our new vertex is inserted as the midpoint m_N along edge e. For this subdivision scheme, this inserted vertex must be repositioned so as to represent an average of the two arcs influencing this middle edge, e. In this way, we construct a smooth transition between adjacent arcs, which in turn will generate a smooth curve due after several subdivisions. A geometric discussion is employed to define the location of vertex.

In order to construct an averaged position, we will first construct an averaged circle. Given the two origins O_0 and O_1 we will compute the average origin location as: $O_N = \frac{1}{2}(O_0 + O_1)$. Similarly, given the two radii, R_0 and R_1, we compute an average radius through: $R_N = \frac{1}{2}(R_0 + R_1)$. We use these two values, origin O_N and radius R_N to define our new averaged circle.

Notice, however, that since this circle is constructed using the average of the two neighbouring circles, it does not pass through v_1 and v_2.

We must then position our midpoint m_N so that it lies on this new circle. Observe that vertex v_1 is angle Θ_1 from some base axis - say the x-axis - around its defining circle C_0. Similarly vertex v_2 is angle Θ_2 from this same base axis, for its own defining circle C_1. Averaging these values, we obtain $\Theta_N = \frac{1}{2}(\Theta_1 + \Theta_2)$ as the angle for our repositioned midpoint from its base axis for its averaged circle. From this, we then position our midpoint to be a distance of R_N away from the origin O_N along the given direction (Figure 5.13(b)). Note that this direction differs from the perpendicular bisector of edge e due to the angle selection employed within this averaged sphere. This new point can then be connected to its neighbouring vertices v_1 and v_2, thereby replacing edge e with the resulting two edges (Figure 5.13(c)).

At this stage, the circle arcs are recalculated for each of the vertices. The result is illustrated in Figure 5.13(d).

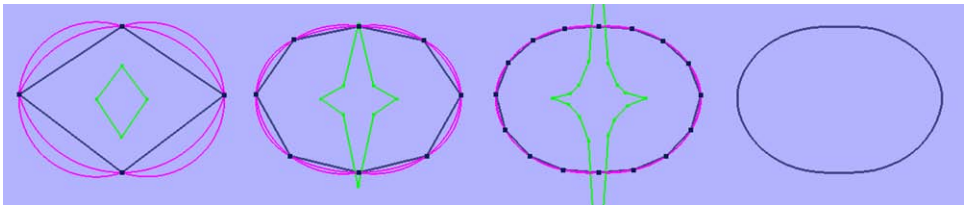

Figure 5.14: Radial-Based Subdivision. In pink, the circle arcs of each set of 3 consecutive vertices are drawn; in green the origins are drawn. Notice how the circle arcs are recalculated after each iteration. Further notice how vertices that are co-linear result in circle arcs with larger radii.

To present this method as a full subdivision scheme, we apply such a midpoint creation and repositioning for all edges along the curve, building off of the base mesh for each midpoint creation. (See Figure 5.14). Determination of circle arcs can be minimized by associating each coarse vertex with an origin and radius computed using its neighbouring vertices. Thus, for each edge split, we simply use the circular arc data from the two bounding vertices. A double-linked data structure - while adding to the memory requirement - would speed up this

process.

5.1.5 Contrast with Radial Basis Functions

One of the key questions for this process is exploring its relation to the radial basis function. A radial basis function is defined as a real-valued function which relies only on the distance from an origin [8]. For example, $\phi : \mathfrak{R}^n \rightarrow \mathfrak{R}^n$, where $\phi(x) = \phi(\|x\|)$, or $\phi(x, c) = \phi(\|x - c\|)$, for c, a central point. Usually, $\|\cdot\|$ is the Euclidean distance.

Radial basis functions are commonly used to approximate functions. Carr et al [11], for example, use the radial basis function to approximate surface interpolation for 3D volumetric data.

For radial-based subdivision discussed here, there is a deviation from the radial basis function. While the position of the midpoints is a linear combination of the different radially-influenced curves, the coefficients used are related to the angular midpoints between the two neighbouring vertices. Similarly, not only is the radius affected, but the origin is also a combination of the neighbouring vertices' influencing arcs. Consequently the radial-based subdivision presented differs from the standard use of radial basis functions for curve approximation.

5.1.6 Length Preservation

Notice that the definition of this subdivision does not currently ensure length preservation. However, we can use the association of each vertex with a circle to reposition all vertices while maintaining the overall length of the curve. As with the main section, Section 5.1.1, exploring length preservation for a circle-converging curve, the repositioning is performed through the identification of the constructed radius, and a scaling down for each respective vertex.

Recall that for each vertex v_i, there is an associated origin O_i and radius R_i. Then, the length of a segment, or edge e can be written as:

$$l_j = 2R_j \sin(\frac{\Theta_j}{2}), 0 < \Theta_j < 180^o,$$

for some Θ_j. These circular properties, however, are associated with vertices and not the edge e. The explicit determination of Θ_j for our edge e could chose to either use the circles associated with the bounding vertices, or alternatively their averaged circle. However, it is unnecessary to explicitly compute Θ_j, as it is merely used to illustrate the rational behind the process. The end result is the determination of a length which may be computed from some radial distance from an origin, in the direction of some angle.

Thus, we can compute the net length as:

$$\begin{aligned} L &= \sum_{i=1}^{n} l_i \\ &= 2.0 * R_i * \sum_{i=0}^{n-1} sin(\frac{\Theta_i}{2}). \end{aligned}$$

Using logic similar to the circular case from Section 5.1.1 (Figures 5.4(e),5.4(f)), we determine a radius from which to scale the given polygon. Let A represent the coarse polygon, and B represent the preliminary subdivision process with midpoint insertion and circle-based repositioning. To have the resulting lengths equal, we determine a radius upon which to scale or shrink the resulting polygon. Thus, we have:

$$
\begin{aligned}
L_A &= L_B \\
2.0 * R_{i_A} * \prod_{i=0}^{n-1} \sin(\frac{\Theta_{i_A}}{2}) &= 2.0 * R_{i_B} * \prod_{i=0}^{n-1} \sin(\frac{\Theta_{i_B}}{2}) \\
\Rightarrow R_{i_A} &= \frac{2.0 * R_{i_B} * \prod_{i=0}^{n-1} \sin(\frac{\Theta_{i_B}}{2})}{2.0 * \prod_{i=0}^{n-1} \sin(\frac{\Theta_{i_A}}{2})} \\
\Rightarrow R_{i_A} &= \frac{L_B}{\frac{L_A}{R_{i_A}}} \\
\Rightarrow R_{i_A} &= \frac{L_B}{L_A} * R_{i_A}.
\end{aligned}
$$

Therefore, we need only scale each influencing radius by the ratio of the original length to the new length (i.e. $\frac{L_B}{L_A}$). Figure 5.15 demonstrates an example curve scaled to preserve net length. Notice the slight repositioning of the original vertices towards the interior of the curve.

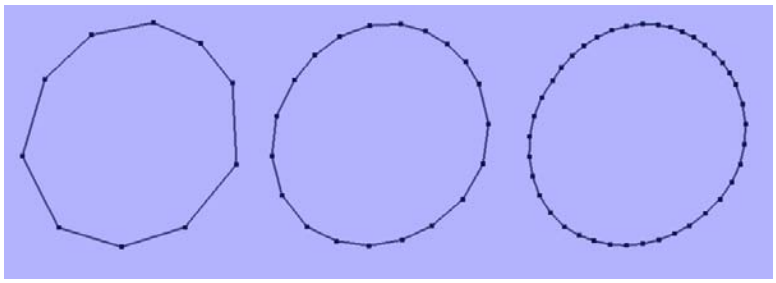

Figure 5.15: Radial-Based Subdivision - Preserving Length.

It should be noted that if the curve is concave in nature, then $\Theta_i > 180^o$. Consequently, when applying this scaling to a concave curve, the vertex will reposition towards the origin of the influencing circle. Since this origin is outside the curve, the vertex will pull away from the enclosing curve. Thus, the length of the curve will increase rather than remain constant as desired. Figure 5.16 illustrates the consequences of scaling such a concave curve.

Additionally, there is the potential to generate a concave curve from an otherwise convex curve. Notice that, but for the overall scaling, the subdivision scheme is interpolating. Consequently the vertices, once fixed, remain positioned proportionally with respect to their

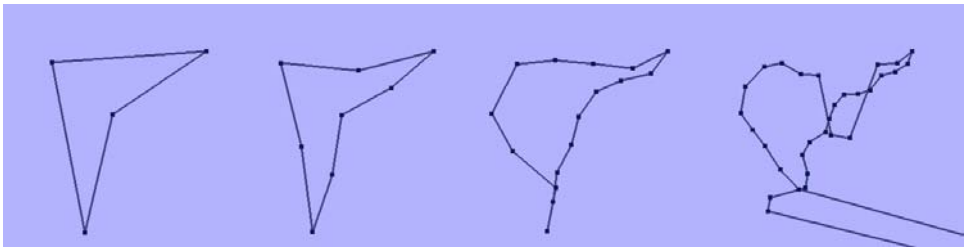

Figure 5.16: Attempting to preserve net length on a concave curve.

neighbouring vertices. Thus, when there is a large discrepancy between the circular arcs of two consecutive segments, (for example, one that is almost linear, and one that is almost perpendicular to it), the resulting curve will become concave. Figure 5.17 illustrates the potential for this development.

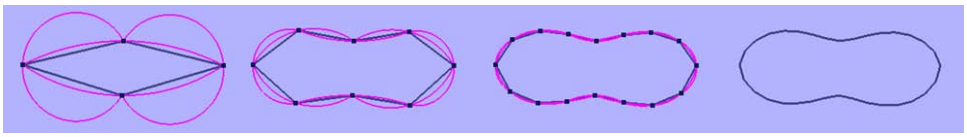

Figure 5.17: Radial-Based Subdivision: Initial convex may create concave curve when large discrepancies between adjacent circles are averaged.

5.1.7 Analysis

Of immediate notice with this subdivision scheme is that using the basic scheme, without length preservation, this subdivision technique interpolates the original and subsequent vertices. This is due to the fact that once placed, vertices remain unmodified in subsequent subdivision iterations. While of interest, there are other elements of this technique that must be discussed to provide a comprehensive evaluation.

One of the relevant aspects of a subdivision scheme is to determine its smoothness. In other words, we must demonstrate the level of differentiability the limit curve or surface maintains. For example, Catmull-Clark subdivision [12] maintains G^1 continuity at irregular vertices and C^2 continuity elsewhere. Evaluating our curve subdivision for smoothness is an essential component of presenting it as a subdivision technique.

Interestingly enough, under certain conditions, the basic subdivision generates cusps. Understanding how and why they form is beneficial to those who may use this subdivision for modeling.

Lastly, an analysis on length preservation is particularly helpful, to justify its success, and establish conditions under which it fails.

The analysis and evaluation of this subdivision process was developed on Windows XP, using Visual Studio 2008. OpenGL and Qt 4.6.2 were used for the rendering of the system, and user interaction, respectively.

5.1.8 Basic Subdivision - Smoothness

In the simplified circle-converging subdivision, the scheme was simplified so that all vertices, at all levels of the process, lay on a common circumscribed circle. Here, this is not the case. Visually, however, there appears to be a component of smoothness resulting from the technique. This smoothness comes from the linear interpolation of the two circle arcs employed for midpoint repositioning. As these arcs are continuously averaged, the resulting curve eventually converges to a smooth transition of these circle arcs.

To demonstrate smoothness, we either need to construct the associated subdivision matrix, or provide a more round-about proof of continuity.

In attempting to construct the subdivision matrix, we may observe a variety of factors. Using terms from Catmull-Clark subdivision, vertex-vertices are left unchanged, hence the interpolating property [12]. Thus the initial subdivision curve (for closed curves) will have the following banded structure:

$$S = \begin{pmatrix} & \vdots & & & \\ 1 & 0 & 0 & 0 & \cdots \\ a & b & c & d & \cdots \\ 0 & 1 & 0 & 0 & \cdots \\ e & f & g & h & \cdots \\ 0 & 0 & 1 & 0 & \cdots \\ i & j & k & l & \cdots \\ 0 & 0 & 0 & 1 & \cdots \\ & \vdots & & & \end{pmatrix}.$$

To determine the values for the inserted edge-vertices, we must review how their position is computed (Recall Figure 5.13). Note that length preservation is not included in this discussion.

- Point P_i, inserted into to edge e_i, as defined above, is located some radius R_{i_N} away from some origin O_{i_N}, about a direction angle Θ from a standard axis (e.g. the x-axis), and can be written as:

$$\begin{aligned} P_i &= O_{i_N} + R_{i_N} \begin{pmatrix} \cos\Theta & -\sin\Theta \\ \sin\Theta & \cos\Theta \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix} \\ &= O_{i_N} + R_{i_N} \begin{pmatrix} \cos\Theta \\ \sin\Theta \end{pmatrix} \\ &= \frac{1}{2}(O_{i-1} + O_i) + \frac{1}{2}(R_{i-1} + R_i) \begin{pmatrix} \cos\Theta \\ \sin\Theta \end{pmatrix}. \end{aligned}$$

- The origin can be found by taking the intersection of the perpendicular bisectors of two of the three edges. The resulting intersection point is the origin. Geometrically, for three

input points, P_1, P_2, P_3, one can compute the perpendicular bisector of $v = \overline{P_1P_2}$ and also $u = \overline{P_2P_3}$. The intersection of vectors u, v defines our origin, O.

- Finally, in terms of our construction, the radius, R can be computed as the magnitude of the vector $O - P_1$.

Due to the inherent geometric description, this results in a unclear breakdown of the weights involved in positioning the respective inserted points. Consequently, an alternative approach for evaluating smoothness is explored.

There are two possible arrangements for a given edge. The first is that the origins of both contributing circular arcs lie on the same side of the edge. In this case, the midpoint will simply be an average of the two neighbouring arcs. As the arcs are recalculated at each consecutive step, the influencing arcs will continue to be averaged. Iteratively applying this process, at the limit curve, we see that the resulting normals will have developed into a smooth quadratic b-spline curve. In other words, there will be a consistently smooth transition of normals along the curve.

Conceptually, one may visualize the subdivided curve as establishing a smooth transition from the coarse circular arc of one vertex to the coarse circular arc of the neighbouring vertex, for a given edge. The iterative process for subsequent subdivisions will continue using the updated and previously averaged curves, thereby smoothing the curve.

The other possible arrangement is for the origins of the circular arcs on a given edge to be on opposing sides. In this case, we see the development of a hard reversal of the normal directions. This may give rise to cusps, which are discussed in further depth in the following subsection.

5.1.9 Basic Subdivision - Cusps

One of the interesting side effects from this subdivision is under specific circumstances, rounded cusps may be formed. (See Figure 5.18).

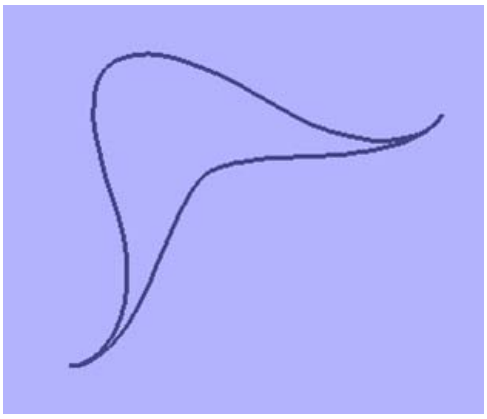

Figure 5.18: Rounded cusps formed by the subdivision.

This happens as a result of the positioning of the origins, upon which midpoints are located. Take, for example, a simple triangle. If the origin of the enclosing circle is located within the interior of the triangle, the mid points will be positioned onto the circumscribed circle, just as with the circle-converging variant of the subdivision.

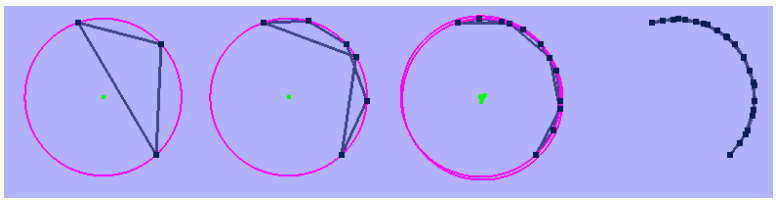

Figure 5.19: Triangle with origin outside the curve. Notice the resulting degenerating curve.

However, if the common origin (See Figure 5.19) is outside, the midpoints, on averaging, will position away from the origin. When working with arbitrary curves, these cusps form as a result of both contributing circular arcs having their origins reside outside their generating three points. On iterative calls to the subdivision, the cusp becomes more pronounced. (See Figure 5.20)

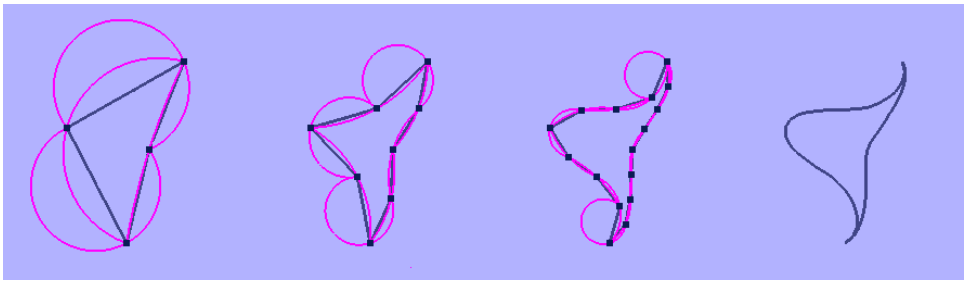

Figure 5.20: The development of a cusp. The origins for the external circles remain outside the closed curve.

Table 5.1: Length Analysis

Shape	Initial	2 Iters	3 Iters	Error (%)
Triangle	442.709	442.709	442.708	2.26e-4
Circum. Quad	456.760	442.797	456.765	3.05
Convex 5-gon	803.23	803.359	803.283	6.60e-3

Notice with this approach that it requires the absence of length preservation. The construction of a concave region will result in undesirable curvature calculations, as discussed below.

5.1.10 Length Preserving

Much of the validation of the length preservation has been presented when it was introduced. It is, however, important to explore the results of length preservation. On the whole, when the initial shape is convex, and concavity is not introduced (by extreme differences in consecutive circle arcs), we see net length is generally preserved.

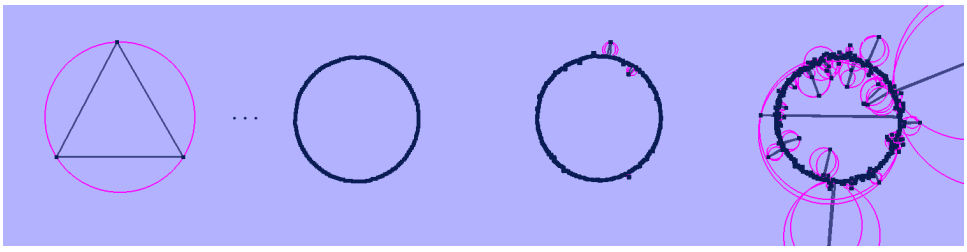

Figure 5.21: Floating point error as a result of co-linear vertices. From left to right: Initial shape, followed by subdivision levels 6 to 8.

Notice, however, that as the curve becomes sufficiently smooth, adjacent vertices become

increasingly co-linear. Consequently, the radius becomes much larger, and at some point, through floating point error, develops concavity, flipping the origin away from the shape's center. (See Figure 5.21). This can be minimized by identifying co-linearity, and preventing the origin from flipping away from the edges.

5.1.11 Surface Extension

There has been some preliminary investigation on extending this subdivision concept to convex 3D surfaces. The idea would be to average spheres, rather than circles, which are uniquely defined given 4 points that are not coplanar. For example, if the origin of a circumscribed sphere lies within a convex tetrahedron, then introduced vertices will be positioned to lie within this common sphere. After a certain number of levels of subdivision, the surface will converge to a sphere, since all averaged spheres will be the same. For an arbitrary polyhedron, this shape may change, but due to the averaged spheres, a smooth surface has the potential for generation.

There are a variety of options for performing the subdivision. One can define the average between two spheres over a shared edge. Alternatively, each vertex can find an averaged sphere by averaging all the spheres that would be defined by the neighbouring vertices (Figure 5.22). Further, vertex insertion, and face splitting may use any style of subdivision technique. The main contribution with this system would be the weighting and positioning of the respective vertices.

Figure 5.22 shows the results of some preliminary exploration in this direction. The initial polyhedron of Figure 5.22(a) is constructed, and the resulting circumscribed sphere is illustrated (Figure 5.22(b)). Through a rudimentary refinement of the mesh, and a repositioning, based on the averaging of adjacent sphere vertices, we construct the shape presented in Figure 5.22(c). Again, the averaged vertex spheres are visualized in Figure 5.22(d). The second level of subdivision (Figures 5.22(e),5.22(f)) demonstrates the unresolved limitations of this approach. The resulting averaged spheres for the vertices become increasingly large, and their

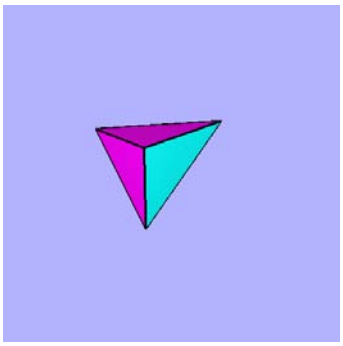

(a) Initial Polyhedron

(b) Initial Polyhedron with Averaged Vertex Spheres

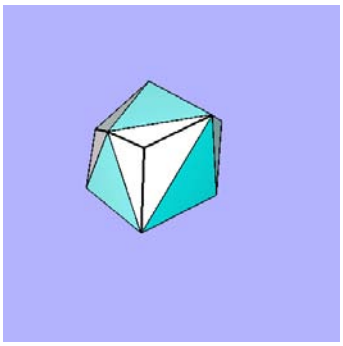

(c) Subdivision Level 1

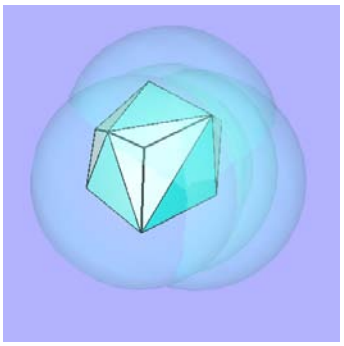

(d) Subdivision Level 1 with Averaged Vertex Spheres

(e) Subdivision Level 2

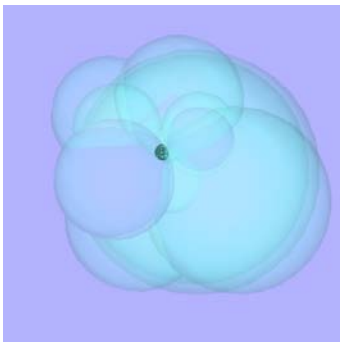

(f) Subdivision Level 2 with Averaged Vertex Spheres

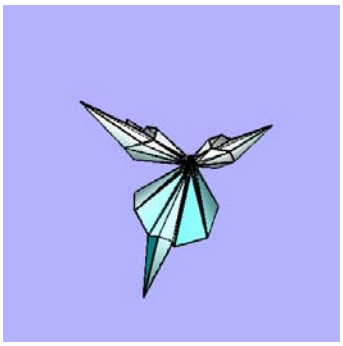

(g) Subdivision Level 3

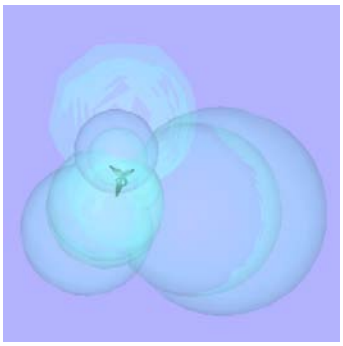

(h) Subdivision Level 3 with Averaged Vertex Spheres

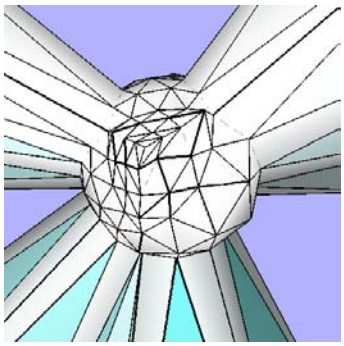

(i) Subdivision Level 3 Close Up

Figure 5.22: Arbitrary Surface Construction

resultant origins further from the mesh of interest. Consequently, the following subdivision, in Figures 5.22(g),5.22(h), have non-smooth, and erratic extrusions. A close up of the subdivision (Figure 5.22(i)) - to the zoom level of the original mesh - demonstrates the potential for smoothness, provided the ability to overcome the aforementioned issue.

However, there are a number of unresolved issues with averaging the spheres (See Figure 5.22(f) to 5.22(i)). For example, with the 2D curve case, an origin is either on one side or another of any given edge, thereby influencing induced concavity. With the 3D case, however, between four vertices, the origin may be inside, outside one, or outside two faces. This causes inconsistencies with the resulting possible positions. As such, the 3D variant would benefit from further investigation, but due to the goal of this work, is not contained herein.

5.1.12 Summary

The analysis of the properties of the base, and length-preserving curve offers an element of brief discussion. It is of particular interest that while the curve appears to exhibit great smoothing abilities, it is also able to generate cusps. These cusps offer an avenue of artistic expression not commonly offered in traditional curve processes, which rely on multiplicities or taking added advantage of the piecewise property of b-splines to define sharp points. These cusps that fall naturally out of this process eliminate the need for such explicit direction.

Much has been accomplished from the curve-aspect of this radial-based subdivision. It may benefit from a more descriptive mathematical breakdown, or an implementation that visualizes more of the length-averaging, or smoothing processes. Further, modifying the implementation so that the coarse level of subdivision represents the initial control points, and only the resulting finer curve is provided may be beneficial for artists, or modellers using this style of curve. Further research into length preservation support for concave curves may be of interest, as might the extension of this general concept to open curves.

Separate from the curve approach, a greater potential for future development is an improved

and more thorough extension to surfaces.

5.2 Local Area Preservation

At this stage we have demonstrated the preservation of the total area for a sphere-converging subdivision process. In order to meet the objective of an area preserving subdivision process, it is necessary to subdivide the shape in such a way that a region on a coarse representation maintains the same area as its area on a finer representation. We refer to this construction method as local area preservation, in contrast with the otherwise global preservation.

The simplest approach is to ensure that at each level of subdivision, every face has an equivalent area. This way, as the subdivision approaches an infinitely divided surface, the infinitely small regions will also be equivalent. Therefore the summation of the infinitely small faces that comprise a bounded region will have the same area as the bounded area on the coarsest representation.

To develop such a subdivision, one must ensure the surface is segmented into regions of equivalent area. As the surface is subdivided, it must illustrate an equal-area division. We start by exploring the trivial effect of such a division process on a curve. The discussion is then extended to surfaces, which takes on a more scientific and experimental approach. Potential subdivision techniques are discussed in this section, and the experimentation is explored in the subsequent chapter.

5.2.1 Curves

When exploring the two dimensional case for local preservation, rather than preserving area, we are instead preserving the length of the curve, or the polygon. As with the simpler global preservation case, we will work with closed, non-intersecting polygon to represent our coarse curve. The aim for the subdivision, as before, is for our coarse polygon to converge to a circle.

As mentioned in the opening section above, in order to have local length preservation for our subdivision process, one must start with sections of comparable size. This means that we will be working with regular n-gons, or regular polygons having n vertices (e.g. Figure 5.23). As with the simpler case, these vertices will be circumscribed within a common circle. By the definition of a regular polygon, all edges are the same length.

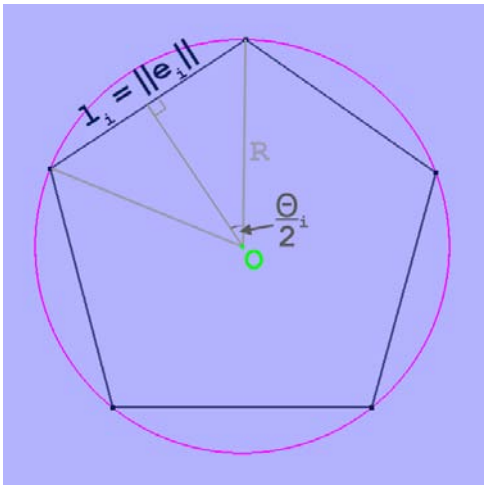

Figure 5.23: Local Equal Area Curve Construction

In constructing our subdivision, the simple midpoint insertion as defined above is employed. In other words, for each edge, $e_i = (v_i, v_{i+1})$ bounded by vertices v_i, v_{i+1}, a midpoint is inserted at the intersection of the perpendicular bisector of e and the circular arc connecting v_i and v_{i+1}. Consequently, these midpoints lie on the same circumscribing circle. As the last step for global length preservation, a radius R_2 is computed to scale the circumscribing circle (and the vertices lying upon it) to ensure a common length is maintained.

Notice that during the midpoint insertion, the two constructed edges are of equal length. This is due to the selection of the bisector for the angle $\Theta = \angle v_i O v_{i+1}$. For the new edge e'_i, we have $\Theta' = \frac{\Theta}{2}$. Therefore, the edge length of e'_i is:

$$l'_i = 2R \sin \frac{\Theta'}{2}.$$

Notice that $\Theta = \frac{360^o}{n}$, a constant for our regular n-gon. Thus, on subdivision, the next iteration will have:

$$\begin{aligned} \Theta' &= \frac{\Theta}{2} \\ &= \frac{360^o}{2n}, \end{aligned}$$

or more generally that:

$$\Theta^{[i]} = \frac{360^o}{2^i n}.$$

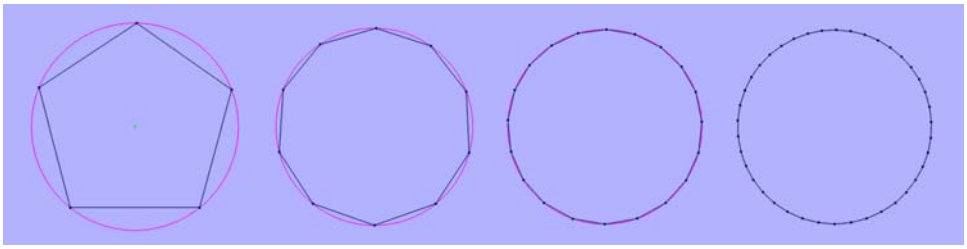

Figure 5.24: Local Equal Area Curve Subdivision

Therefore, for our subsequent iterations for the subdivision, we are simply generating an N-gon, where $N = 2^i n$, circumscribed within a radius of $R^{[i]}$. Therefore, at $i \to \infty$, we observe that all edge lengths will be equivalent, and through our global area preservation, that the edges will equally divide the total length. Consequently, for a segment on the original coarse polygon, it will be subdivided into a region of equivalent length on the more refined representation. This is visualized in Figure 5.24. Notice how for each level of subdivision, all edges have equivalent length, yet the curve still converges to a circle.

5.2.2 Surfaces

The next stage of this work is to explore the effects of an equal area subdivision on a sphere-converging surface. Borrowing from the approach taken for two dimensional curves, we ob-

serve that there are a handful of stages and requirements that must be met:

- Start with a surface that is circumscribed within a sphere
- Subdivide the surface so that there are:
  - Additional faces constructed
  - These faces are circumscribed within a sphere
  - The area of each of these faces is equivalent to all other faces at the same subdivision level
- Apply the global area preservation, from Section 5.1.2, to retain global areal maintenance

The crux of this system is the retention of areal equivalence between faces at a given subdivision level. Within the two dimensional case, we observed how this is trivially retained. Selecting a regular polygon with n edges, upon applying the subdivision, a regular polygon with $2n$ edges is constructed. Due to the symmetric nature of the subdivision, the resulting edge lengths are equivalent.

The development of a symmetric subdivision for surfaces is particularly challenging. The actual objective of such an approach is to generate faces which are all the same - area, angular, and distance-wise - at a given level of resolution. From the sphericity of this work, it becomes particularly difficult to meet these objectives.

As such, an experimental approach was employed for the evaluation of a subdivision's potential for supporting local area equivalence. A selection of commonly employed and alternately generated subdivisions was evaluated against their ability to support equal area preservation. Chapter 6 documents this experiment, detailing the approaches taken, the methodology, and the results.

Chapter 6

Experiments of Equal Area Spherical Subdivision

An experimental approach has been adopted to explore the possibilities of selecting an equal area subdivision which preserves area between faces. The subdivision, which as an additional requirement, must converge to a sphere - due to the motivating needs of the geoscience visualization community. Consequently, the attainment of smoothness is inherent within the subdivision process.

The experiment's objectives are two fold. Initially, it must validate the global area equivalence method as outlined in Section 5.1.2. To this end an initial surface area will be computed, and all subsequent layers of subdivision will be contrasted against this.

The second objective is the selection of an appropriate equal area subdivision. Since we are unable to readily discern a symmetric surface-based subdivision, as was available for curves in Section 5.2.2, some of the traditional refinements are instead adapted for our purposes. Their feasibility is measured and reported.

6.1 Methodology

A variety of approaches are evaluated for their potential as a preserving local area throughout the subdivision process. Where applicable, a geometric discussion of their construction is presented. This is the preferred choice for traditional cartographers who define projections for maps through their geometric properties. This dates back to the 12th century, when cartographers constructed maps through compass and ruler [49].

Unless otherwise discussed, the icosahedron (Figure 6.1) is employed as the base mesh for this experimentation. Its circumscription within a sphere and its high face count enable it to

better initially approximate a globe, as compared, for example, with the tetrahedron. Further, its uniformity of faces - namely all twenty triangles - are highly desirable for their efficient mesh construction within computer graphics.

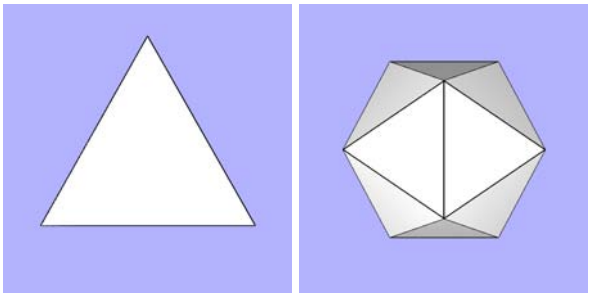

(a) Mesh Face (b) Polyhedron

Figure 6.1: Base Mesh - Icosahedron

Where applicable, from a qualifying subdivision approach, alternative polyhedron may be considered for employment, but are left for the reader to explore. In general, such alternatives provide a coarser approximation of the sphere, and will lend themselves to higher distortion during the subdivision process.

6.1.1 Split (1:4)

As discussed in Section 2.3.2, the Loop Subdivision (Figure 6.2) is well established within the computer graphics industry. Working within, and generating, a triangular mesh, this approach seems well suited for our equal area approach, which also depends upon triangular faces.

To adapt this refinement to our sphere-converging subdivision, we eliminate the repositioning of the vertices. The triangular face is still divided into four subtriangles. However, original vertices initially remain on the common sphere, S, and the midpoints are projected outwards along their normals (from the origin) to coincide with S. Once these are all circumscribed within a common sphere, they are repositioned according to the global area preservation quality outlined in Section 5.1.2.

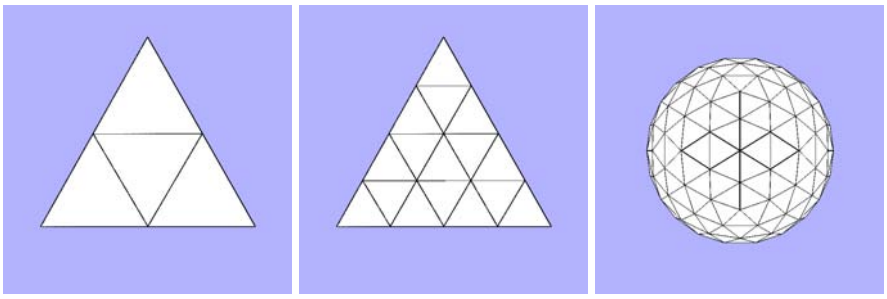

(a) Level 1 Face (b) Level 2 Face (c) Level 2 Split

Figure 6.2: Flattened Loop - Split (1:4)

If divided, without repositioning, on the base mesh, the new triangles will all have equivalent area, as they are symmetrically constructed (Figure 6.2(c)). However, it is expected that the repositioning for residing on a sphere will remove this similarity. Unfortunately, the lack of freedom in selecting the midpoints along the edges, due to the influence of the neighbours, eliminates the ability to reposition this construction for equal area. Figure 6.3 illustrates the result of vertices on the first level of subdivision repositioned onto the sphere. The blue of the interior face indicates a larger area than the slightly pink exterior faces. If the mid points are repositioned, the areas of the resulting triangles will be inconsistent and not equal.

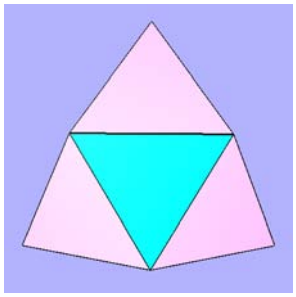

Figure 6.3: Split (1:4) Face on Sphere

This process is tested to confirm this hypothesis, and quickly eliminate it as a potential subdivision candidate. It does, however, enable the validation of the global area preservation approach.

6.1.2 Split (1:9)

Loop's construction of four triangular faces from the original is limited due to its fixed midpoint location. Inspired by the supplemental subdivision from Song et al. [51], adjusting the face count so that nine faces are constructed instead of four is explored. This approach has the potential to increase the flexibility or a degree of freedom during construction to ensure equal area triangles throughout the division process.

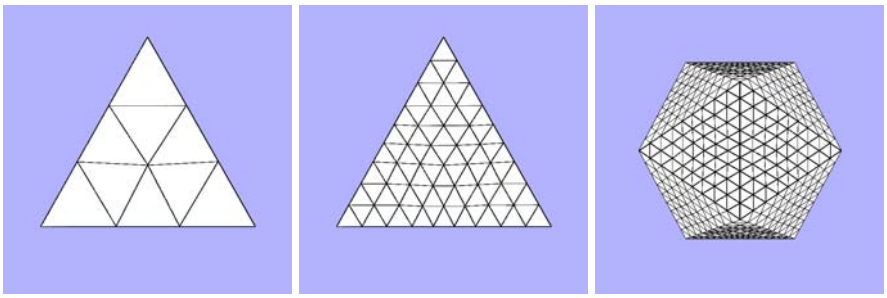

(a) Level 1 Face (b) Level 2 Face (c) Level 2 Split

Figure 6.4: Split (1:9)

From a planar perspective, each edge is split into thirds, and connected according to Figure 6.4. As with the modification of Loop's subdivision, the inserted vertices are projected to the base mesh's common circumscribing sphere, and then all vertices reposition according to the global preservation approach from Section 5.1.2.

It is expected that the resulting vertices will offer a degree of flexibility for adjustment, to ensure equivalent area throughout the subdivision process. This is due to the increases face count, and in particular, the interior vertices, which may be repositioned to account for inequalities.

6.1.3 Catmull-Clark

As described in Section 2.3.1, the Catmull-Clark Subdivision (Figure 6.5) is widely employed within the computer graphics industry. Therefore, it is important to evalute its refinement approach to establish its potential for equal area support.

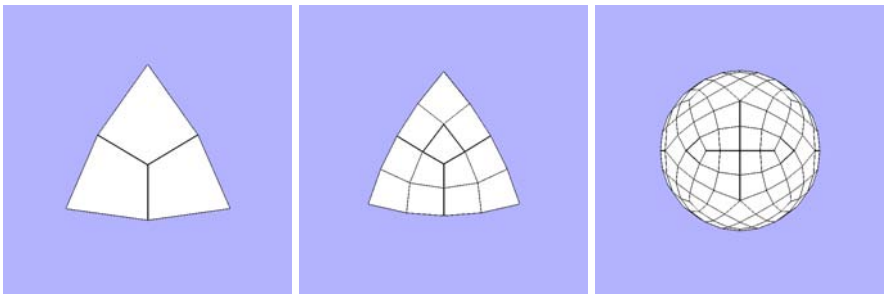

(a) Level 1 Face (b) Level 2 Face (c) Level 2 Subdivision

Figure 6.5: Catmull-Clark

For our purposes, similar to the approach with Loop's subdivision, this approach is adjusted so that each iteration of the subdivision resides within a circumscribing sphere. As with prior approaches, vertices are repositioned according to the global area preservation approach from Section 5.1.2 and evaluated for local area preservation amongst the remaining faces.

Since this subdivision quickly creates quadrilaterals, it is anticipated that the global area preservation approach may not hold. Recall that the global area preservation is described through triangular faces (Section 5.1.2). As such, areal equivalences can be computed independent of the angles used to construct the respective faces. With the Catmull-Clark subdivision constructing quadrilaterals, which are not planar due to the vertices residing on a common sphere - as per the construction process - it is expected that areal equivalence will not be able to hold.

The Catmull-Clark method is evaluated as one of the control groups for this global preservation, and secondarily evaluated for the potential for local area preservation. It is expected that the local area will not preserve well throughout the subdivision as the faces quickly become extremely distinct.

6.1.4 Split (1:3)

As an additional base case, a simple modified version of the Catmull-Clark subdivision is also presented. As seen in Figure 6.6(left), this refinement similarly creates a single interior vertex.

However, instead of connecting it to the midpoints of the edges, it is instead connected with the face's original vertices.

Here, the midpoint is computed as the average of all the face's vertices. Namely:

$$v_f = \frac{1}{n}(\sum_{i=1} nv_i),$$

where v_i are the face's original vertices.

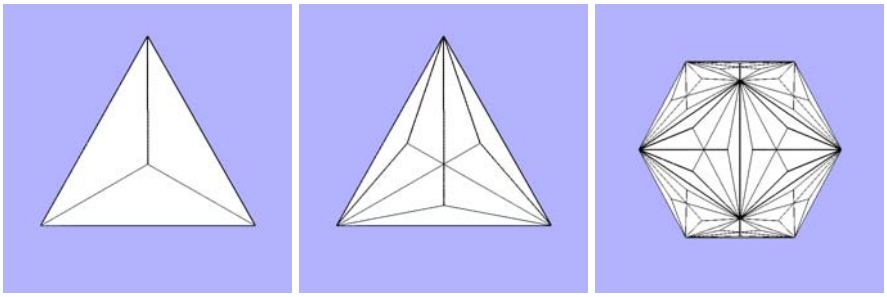

(a) Level 1 Face (b) Level 2 Face (c) Level 2 Split

Figure 6.6: Split (1:3)

As with prior approaches, the inserted vertices are projected along their normals, with respect to the origin, so as to reside on the same circumscribing sphere as the original vertices. All vertices are then repositioned according to the global area preservation of Section 5.1.2, and faces are evaluated for their potential at local preservation.

This simple approach is anticipated to validate the global area preservation, while offering an alternative approach at subdividing the surface. It is further anticipated to quickly lose a common face shape, and therefore quickly lose any local area preservance. Furthermore, the one degree of freedom throughout the process, as with prior approaches, offers little in the way of addressing inequalities of area. However, as it is uninfluenced by the neighbouring faces, this may be of less concern and offer unforeseen flexibility.

6.1.5 Split (1:6)

A merger of the above two approaches is presented in this alternative modified Catmull-Clark approach. Here, six faces are generated from the original, all reliant on the interior vertex, but also additionally reliant on one of the edge midpoints. In this way, the edges are broken up during the subdivision process, adding a flexibility of smoothness as illustrated in Figure 6.7.

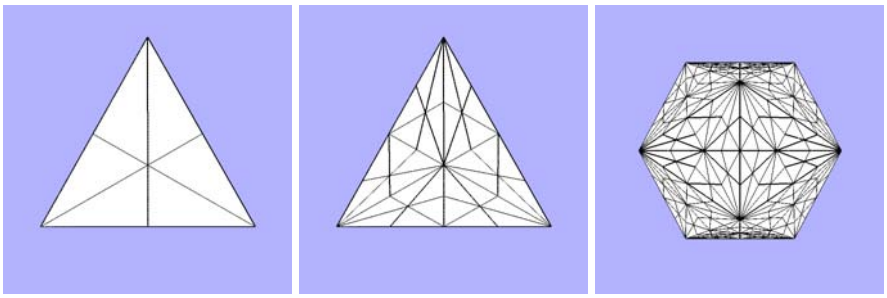

(a) Level 1 Face (b) Level 2 Face (c) Level 2 Split

Figure 6.7: Split (1:6)

The interior vertex is constructed, similar to the aforementioned 1:3 process as the average of the face's vertices:

$$v_f = \frac{1}{n}(\sum_{i=1} nv_i),$$

while the edge points are defined as the midpoints of the given edge. In other words,

$$v_e = \frac{1}{2}(v_i + v_{i+1}),$$

where v_i and v_{i+1} are the vertices that define our edge.

As with prior approaches, the inserted vertices are projected along their normals, with respect to the origin, so as to reside on the same circumscribing sphere as the original vertices. All vertices are then repositioned according to the global area preservation of Section 5.1.2, and faces are evaluated for their potential at local preservation.

It is anticipated that the resulting surface will preserve local area well. This is due to the increase in faces, and subsequent symmetric approach for subdivision.

6.1.6 Split (1:5)

In the Doo-Sabin subdivision [17, 18], the different geometric elements of a mesh are replaced with faces. Faces are replaced by miniature versions of themselves, edges are replaced with quadrilaterals, and vertices are replaced with faces having the same edge count as the vertex's valence. These are all joined together, and their constructed vertices positioned in such a way as to achieve G^1 continuity at irregular vertices, and C^1 everywhere else.

Due to the inconsistency of created faces, it is expected that such a subdivision would not support local area preservation. Ensuring every triangle, quadrilateral, and other n-gons across every subdivision level are equal in area is a challenging task, which may not be solvable. As such, the Doo-Sabin construction approach has not been evaluated for global preservation.

A variation of the Doo-Sabin subdivision, however, has been presented as a possible approach. Suggested by our industrial partner [40], who employ a similar subdivision on their planar representation, it has been considered for analysis. Starting with a dodecahedron (Figure 6.8), vertices are replaced by hexagons, and faces are replaced with smaller versions of themselves.

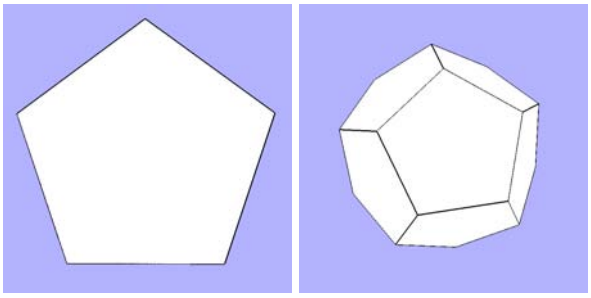

(a) Mesh Face (b) Polyhedron

Figure 6.8: Alternative Mesh - Dodecahedron

Figure 6.9 illustrates this construction between subdivision level 1 (black outline) and 2

(blue outline). The blue hexagons are added at the vertices of the underlying pentagon, while the pink faces are that of the original face, reduced in size. In this way, the surface becomes more composed of hexagons, while retaining the original twelve pentagons, as illustrated in Figure 6.10. For example, at the first level of resolution, a truncated icosahedron is constructed.

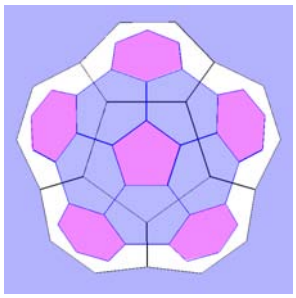

Figure 6.9: Split (1:5) Construction

This refinement approach is also akin the employment of a duality mechanism. Faces are constructed based on the positions of the underlying vertices, and old edges intersect with the newly constructed edges (though not the resized edges of the prior faces).

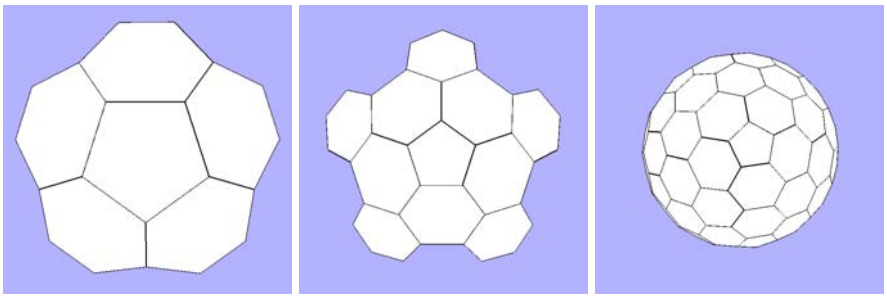

(a) Level 1 Face (b) Level 2 Face (c) Level 2 Split

Figure 6.10: Split (1:5)

Similar to the polyhedral subdivisions as above, once the initial subdivision takes place, the vertices are then repositioned to retain net surface area equivalence.

As with Catmull-Clark, and the construction of non-triangular faces, it is expected that this refinement may not be well suited for global area support. Since the hexagonal and pentagonal faces have vertices residing on a sphere, the non-planarity of the face results in inconsistent

area calculations, which may not be able to keep angles independent - a requirement from the global area construction (Section 5.1.2). However, this otherwise mostly hexagonal-face constructed mesh seems to provide some great potential for exploration, as discussed by our industrial partner [40].

6.1.7 Split (2:6)

From a suggestion from our industrial collaborator, a supplementary approach has been taken for refinement. In this approach, the icosahedron is initially divided symmetrically through the Split (1:6) approach. Subsequent iterations pair up neighbouring faces and divide them into six respective faces (Figure 6.11). It is observed that this can alternatively be viewed as dividing an individual face into three components. However, the division of one of the edges strongly influences the neighbouring face's subdivision.

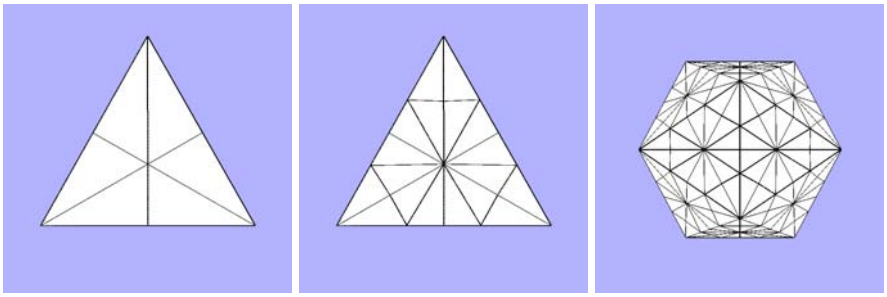

(a) Level 1 Face (b) Level 2 Face (c) Level 2 Split

Figure 6.11: Split (2:6)

The repetitive refinement employs an original vertex, and inserted vertices along two of the edges of the triangular face. Based on the icosahedron face, a unique symmetric triangle is generated after the initial Split (1:6) (Figure 6.12(a)). This triangle has angles 30^o, 60^o and 90^o. Splitting the 60^o angle symmetrically, and splitting the residual $30^o/30^o/120^o$ triangle symmetrically, we generate the refinement as visualized in Figure 6.12(b). Consequently, we create three equal subtriangles, all smaller, yet conformal with the original triangle. Since we have generated smaller version of the original triangle, we can comparably decompose through

subsequent iterations (Figure 6.11(c)).

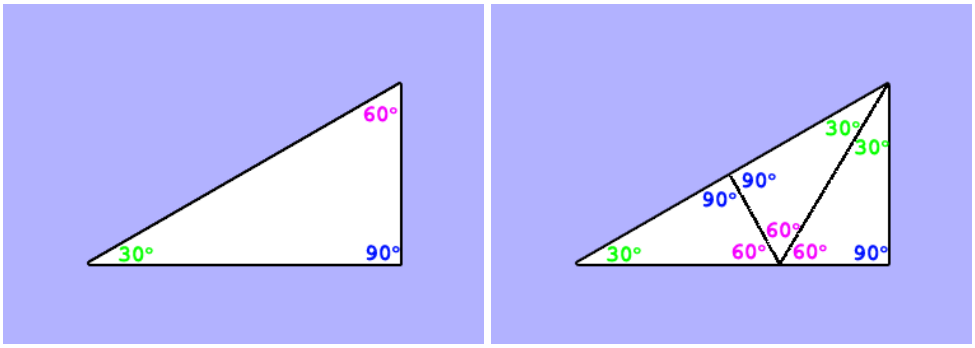

(a) Level 1 Face Close Up (b) Level 2 Face Close Up

Figure 6.12: Split (2:6) Construction

The initial employment of the 1:6 approach symmetrically divides the polyhedron, and constructs faces of equivalent shape. This lends itself to an improved spherical approximation while retaining shape, and consequently local area. The remaining steps are initially selected and positioned based on the preservation of local area. Consequently it is expected that at the coarser levels, equal area will be better maintained. As the faces exhibit increasingly different shapes, it is expected that the vertex placement will be difficult to balance the areal requirements of neighbouring faces. However, since this is occurring for smaller and smaller faces, it is expected that proportionally any resultant error will be minimal.

In summary, seven different refinement techniques will be explored for their ability to support local area, while verifying global area preservation. Splits (1:3), (1:4), (1:5), (1:6), (1:9), (2:6) and a modification of Catmull-Clark will be evaluated. It is expected that triangular refinements will have better support for global area preservation, as they are more aligned with the angle-independent requirements of the approach. A description of the methodology is presented, and then followed by the results.

6.2 Implementation

The application was developed on Windows XP, using Visual Studio 2005. OpenGL and Qt 4.1.2 were used for the rendering of the system, and user interaction, respectively.

A half-edge data structure [30] was employed to represent the underlying surface. The base polyhedra were borrowed from John Burkardt's collection [9], and converted to the half-edge data structure within the implementation.

Each of the aforementioned subdivisions was implemented. Upon performing the respective approach, unless otherwise stated, the vertices were re-projected to the surface of the inscribing sphere. This projection occurs along the normal defined by the vertex v, and the origin, O of said sphere. Once all vertices lay within a common sphere, the vertices were re-projected, or scaled, so as to maintain global area equivalence as outlined in Section 5.1.2.

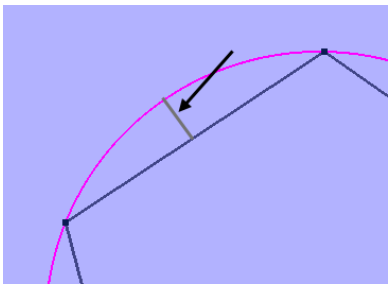

Figure 6.13: Face Displacement

For each approach, the base mesh was subdivided until the resultant surface would be represented within 1km of the Earth's surface (Figure 6.13), or until the surface demonstrated irreconcilable distortions. The summation of the area of all faces was calculated, and its percent deviation from the original surface area was computed. Recall that these resultant polyhedrons are aiming to approximate a spherical surface. As such, the maximal and average percent distance from the circumscribing sphere with respect to the midpoint of a face at a given resolution level was also recorded. Furthermore, the maximal percent error of a face's area at a given resolution was extracted. Finally, a distribution of the percent areal error was evaluated

and plotted.

6.3 Results

The results of the implementation and experiments performed are presented in the following sections. An initial exploration of the global area preservation approach from Section 5.1.2 is presented. This is followed by a local area preservation exploration, as from Section 5.2.2. For each approach, a visualization of the local area distortion, in graph and graphical form is incorporated.

6.3.1 Global Area Preservation

In order to validate the global area preservation approach as defined in Section 5.1.2, it is important to determine the amount of deviation the surface area of one subdivision level is from the original surface. To this end, Table 6.1 presents the deviation of surface area from the original.

Recall that surfaces are only subdivided until they closely approximate the sphere in which they are circumscribed. As such, Split (1:9) is only evaluated through three levels of subdivision, while Split (1:3), Split (2:6), and Split (1:5) are evaluated through seven subdivision levels. Further recall that Split (1:5) has a base mesh of a dodecahedron, differing from the icosahedron of all others. The original surface area is calculated for this original surface, and subsequent subdivisions are compared against it for this approach.

Notice from this Table 6.1 that subdivisions which begin and retain only triangular faces, namely Split (1:4), Split (1:9), Split (1:3), Split (1:6) and Split (2:6), maintain a nominal deviation from the original surface area. With an average deviation between these approaches of 9.36e -09%, we observe for the Earth's surface area of 5.10e08 km^2, an absolute error of 4.78e-02 km^2 occurs. This equates to the area of a small residential community (approximately

Table 6.1: Global Area Preservation Results - Percent Deviation (%)

Approach	Level 1	Level 2	Level 3	Level 4	Level 5	Level 6	Level 7
1:4	-2.49e-06	-2.33e-06	-9.00e-07	-8.26e-07	6.68e-07	1.90e-06	
1:9	-3.89e-08	9.05e-07	6.44e-08				
1:3	6.23e-07	-2.80e-06	1.95e-07	2.09e-06	-5.54e-07	4.54e-07	-3.98e-06
1:6	2.49e-06	-2.14e-06	1.06e-06	-2.08e-07			
2:6	2.49e-06	1.48e-06	4.28e-07	2.41e-06	-4.50e-07	9.80e-07	-2.66e-06
1:5	12.47	17.12	18.70	19.24	19.42	19.48	19.50
Catmull	55.78	-11.16	49.75	-12.05	49.36	-12.11	

160 houses, assuming 300 m^2 per house), spread across the surface of the subdivision surface.

6.3.2 Local Area Preservation

The results of the local area preservation experiments across the different subdivision approaches are presented. For each approach, for each subdivision level, the area of each face is calculated. The percent deviation refers to the difference the face is with respect to its expected face area. The computation for this percentage is:

$$P = 1 - \frac{(A_{face} - A_{expected})}{A_{expected}},$$

with $A_{expected}$ computed through:

$$A_{expected} = \frac{A_{net}}{n_{faces}},$$

where A_{net} is the surface area of the base polyhedron, and n_{faces} represents the number of faces for the given subdivision level.

6.3.3 Overview

Numerical analysis of their results is presented below. To assist in determining convergence of the sphere, a face count for each approach at the computed subdivision levels is presented

in Table 6.2. As expected, approaches that divide the original face into a larger number of subsequent faces quickly subdivide the polyhedron into small face. Split (1:3), on the other hand, takes a while to generate its 43,740 faces.

Table 6.2: Face Count

Approach	Level 1	Level 2	Level 3	Level 4	Level 5	Level 6	Level 7
1:4	80	320	1280	5120	20480	81920	
1:9	180	1620	14580				
1:3	60	180	540	1620	4860	14580	43740
1:6	120	720	4320	25920			
2:6	120	360	1080	3240	9720	29160	87480
1:5	32	92	272	812	2432	7292	21872
Catmull	60	240	960	3840	15360	61440	

Upon computation of each approach's face count, it was important to determine when a subdivision reached within the 1 km threshold of the surface of the Earth, or when it severely deviated. As such, the spherical displacement, as a percentage of the circumscribed sphere, was computed.

Geometrically, the computation is the percent difference between the midpoint of a face, with the radius of its circumscribing sphere. Since these polyhedra are convex, the midpoints of faces reside within the interior of the circumscribing sphere (Figure 6.13). For the icosahedron, the midpoints its faces are 20.5% away the sphere in which it is circumscribed.

As illustrated in Table 6.3, these approaches quickly reduce this 20.5% difference. For Split (1:9), the first level of subdivision, on average, as its faces only 2.66% away from its circumscribing sphere. All subdivisions converge to within 0.4% of their circumscribing spheres before additional subdivision is ceased.

Since the faces differ across the subdivision surface, particularly at the later levels of subdivision, it is helpful to be aware of the maximal displacement from its approximating sphere. As such, the maximal value is presented in Table 6.4.

Notice that although Split (1:3) had converged to an average displacement of 0.34% in Table 6.3, its constructs faces that remain 9.96% away from its circumscribing sphere. This

Table 6.3: Average Spherical Displacement (%)

Approach	Level 1	Level 2	Level 3	Level 4	Level 5	Level 6	Level 7
1:4	5.84%	1.51%	0.38%	0.10%	0.02%	0.01%	
1:9	2.66%	0.30%	0.03%				
1:3	11.35%	6.33%	3.53%	1.97%	1.10%	0.61%	0.34%
1:6	4.79%	1.07%	0.24%	0.05%			
2:6	4.79%	1.74%	0.60%	0.20%	0.07%	0.02%	0.01%
1:5	7.64%	2.64%	0.89%	0.30%	0.10%	0.00%	0.01%
Catmull	5.31%	1.32%	0.33%	0.08%	0.02%	0.01%	

is a fairly large discrepancy between resultant faces, only cutting down face distance by half from the original 20.5%.

Table 6.4: Maximumal Spherical Displacement (%)

Approach	Level 1	Level 2	Level 3	Level 4	Level 5	Level 6	Level 7
1:4	6.58%	1.78%	0.45%	0.11%	0.03%	0.01%	
1:9	2.86%	0.33%	0.04%				
1:3	11.35%	10.15%	9.98%	9.96%	9.96%	9.96%	9.96%
1:6	4.79%	1.49%	0.53%	0.19%			
2:6	4.79%	1.86%	0.67%	0.23%	0.08%	0.03%	0.01%
1:5	8.09%	2.77%	0.98%	0.33%	0.11%	0.04%	0.01%
Catmull	5.31%	1.76%	0.55%	0.17%	0.05%	0.02%	

With the Earth's average radius set at 6,371 *km* [10], we observe the following average absolute spherical displacement within the subdivision. This refers to the average distance from the surface of the Earth to the constructed spherical approximation. From Table 6.5, 0.71, for example, means that we are 0.71 km from the Earth's surface, or in other words, a mere 710 m away from the surface of the Earth.

Notice from Table 6.5 that in general we quickly approach less than 1 km distance from the Earth's surface. It should be pointed out that Split (1:3) ceases evaluation when it is unable converge to within 1 km of the Earth's surface. Consequently, it is no longer subdivided after level 7.

Having established the subdivision level for cessation of evaluation, it is important to establish the percent deviation for each of the faces from the expected face area. Notice, from

Table 6.5: Absolute Spherical Displacement (km)

Approach	Level 1	Level 2	Level 3	Level 4	Level 5	Level 6	Level 7
1:4	372.31	96.27	24.27	6.08	1.52	0.38	
1:9	169.68	19.19	2.14				
1:3	723.06	403.10	224.89	125.37	69.83	38.87	21.63
1:6	304.98	68.44	15.24	3.39			
2:6	304.98	110.65	37.96	12.78	4.28	1.43	0.48
1:5	486.97	168.23	56.85	19.06	6.37	2.13	0.71
Catmull	338.15	84.25	20.98	5.24	1.31	0.33	

our general area evaluation (Table 6.1) that on the whole, most subdivisions retain their overall surface area. As a result, an average face areal error will average out to an equivalent value. Ergo, an average face areal error is not presented.

However, a maximal face error is meaningful, and can be presented. Table 6.6 illustrates the maximal error an approach has for each of the subdivision levels. The absolute value of the deviation is presented to eliminate confusion between maximal reduction and expansion of a face.

Notice how Split (1:5) and the Catmull-Clark Split fluctuate wildly. This corresponds to their inability to retain consistent global area preservation. Interestingly Split (1:4) holds steady at an approximate 20% maximal deviation, whereas Split (2:6) retains a reduced percent change, and maintains a gradual increase over the subdivision. It is important that while this may be the maximal face areal error, it does not represent the number of faces that experience this maximal error. Distribution of the areal deviation is presented in the following subsections.

As with the calculation for the convergence of the subdivision surfaces, associating these maximal face errors with a region bounded on the Earth's surface is of importance. With the Earth's surface set at 510,072,000 km^2, Table 6.7 converts the percentages presented in Table 6.6 into real world values. The achievement of 970 km^2 from Split (2:6) is roughly the size of San Diego, California which is approximately 965 km^2 (2011, US Census [56]).

Table 6.6: Maximum Face Areal Error (%)

Approach	Level 1	Level 2	Level 3	Level 4	Level 5	Level 6	Level 7
1:4	13.42%	18.64%	20.14%	20.53%	20.62%	20.65%	
1:9	8.12%	9.51%	9.67%				
1:3	0.00%	33.73%	172.39%	552.13%	1469.73%	3619.13%	8525.78%
1:6	0.00%	2.81%	5.39%	8.84%			
2:6	0.00%	0.21%	2.28%	5.37%	8.92%	12.96%	16.63%
1:5	22.17%	23.77%	35.34%	48.83%	59.60%	68.12%	74.84%
Catmull	55.78%	40.04%	145.48%	75.65%	261.67%	157.67%	

Table 6.7: Maximum Face Areal Error - Absolute (km^2)

Approach	Level 1	Level 2	Level 3	Level 4	Level 5	Level 6	Level 7
1:4	8.56e05	2.97e05	8.02e04	2.04e04	5.14e03	1.29e03	
1:9	2.30e05	2.99e04	3.38e03				
1:3	85.0	9.56e05	1.63e06	1.74e06	1.54e06	1.27e06	9.94e05
1:6	42.5	1.99e04	6.36e03	1.74e03			
2:6	42.5	3.02e03	1.08e04	8.45e03	4.68e03	2.27e04	970
1:5	3.97e06	1.54e06	7.87e05	3.66e05	1.49e05	5.69e05	2.09e04
Catmull	7.38e06	7.56e05	1.16e06	8.84e04	1.30e05	1.15e04	

From the aforementioned definition of percent deviation, and colouration and visualization of the distribution across the surface of the subdivision is presented. Purple, blue and white are used to represent a given face's deviation. The closer to white, or the paler the colour of a face, the closer its areal error is to 0.0%. Blue represents an increase in size, whereas purple represents a decrease. It was established that by keeping a constant colouring scale across all approaches, many of resultant polyhedra exhibiting nominal area error would appear entirely white. To better visualize the distribution of the face error, colours are instead scaled from 0% to the maximal error for the given approach at the given subdivision level.

Edges are visualized until they occlude the visualization of the error distribution across the faces of the subdivision. This is particularly helpful in approaches which become rapidly divided.

6.3.4 Split (1:4)

Figure 6.14 illustrates the areal distribution across the first five levels of subdivision for our Split (1:4) which is undergoing global area preservation. Recall from Table 6.6, this scheme had a maximal error of approximately 20% through the subdivisions. The symmetrical nature of the error distribution across the faces is plainly visible. A dichotomy between an increase and a decrease in area is also visualized.

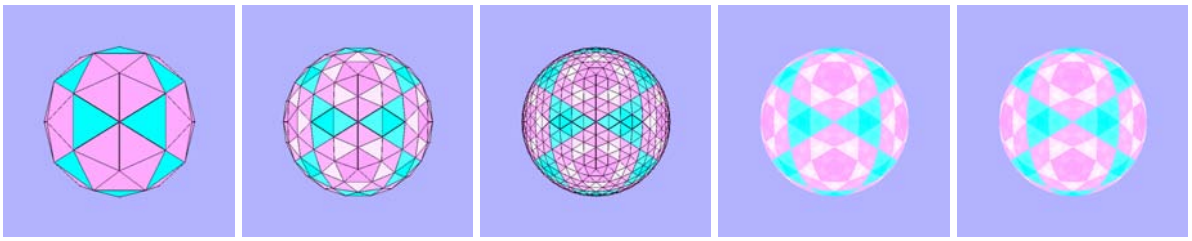

(a) Level 1 (b) Level 2 (c) Level 3 (d) Level 4 (e) Level 5

Figure 6.14: Area Highlighting - Split (1:4)

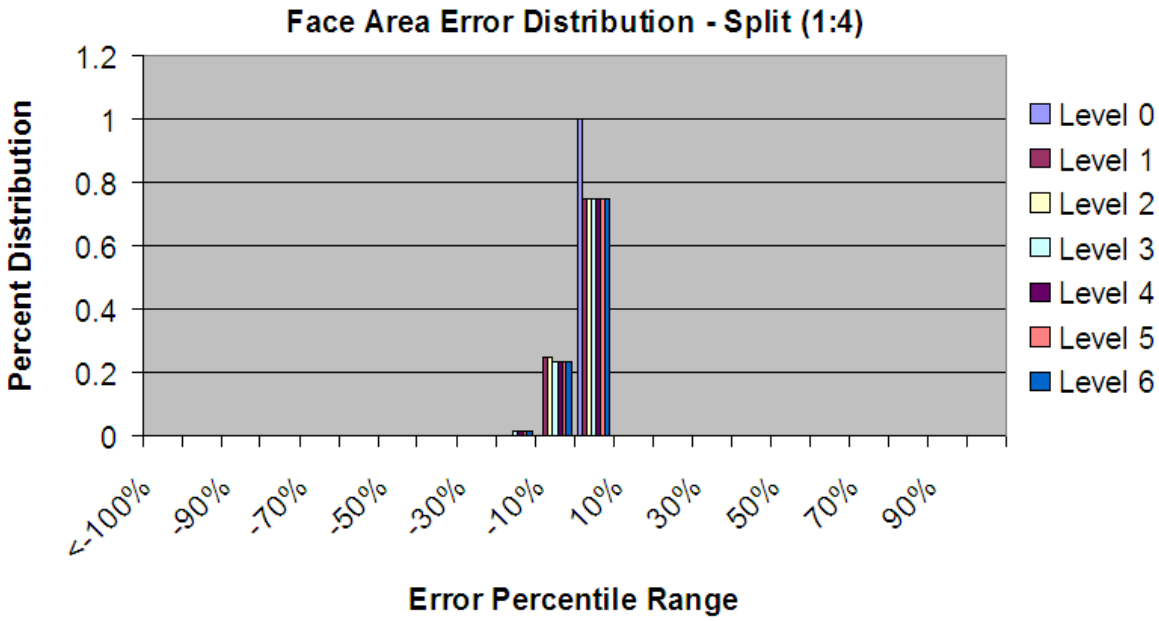

Figure 6.15: Face Area Error Distribution - Split (1:4)

The graph in Figure 6.15 visualizes the distribution - increases and decreases. Approxi-

mately 25% of faces reside outside of the bulk of the faces whose average error is closer to 0%.

6.3.5 Split (1:9)

Split (1:9) has a face error distribution visualized in Figure 6.16. As with Split (1:4), a clear symmetry of the error distribution is visible. This indicates a pattern of too large and too small faces tiling the surface.

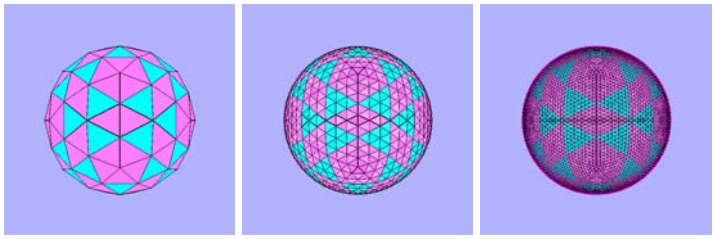

(a) Level 1 (b) Level 2 (c) Level 3

Figure 6.16: Area Highlighting - Split (1:9)

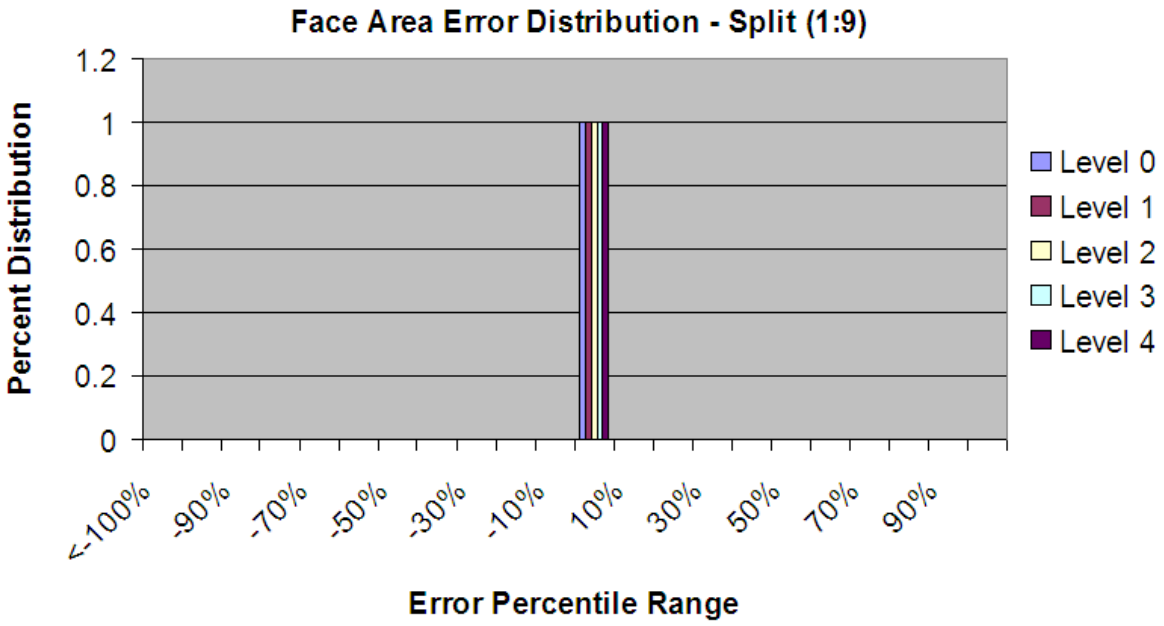

Figure 6.17: Face Area Error Distribution - Split (1:9)

However, as illustrated in Figure 6.17, this approach exhibits a very uniform distribution of cells. All faces, at all subdivision levels, reside within close proximity to the 0% percentile.

As was mentioned during its initial introduction, Split (1:9) appears to have a few degrees of freedom. A trivial test concluded that this is not actually the case. The central point is restricted to reside within the middle of the face, especially early on when the face is a regular triangle. Each outer vertex on an edge must be symmetric with respect to its partner along that same edge. If not, issues in edge matching will result, as a consequence of vertices of odd valence. This has been visualized in Figure 6.18(b)

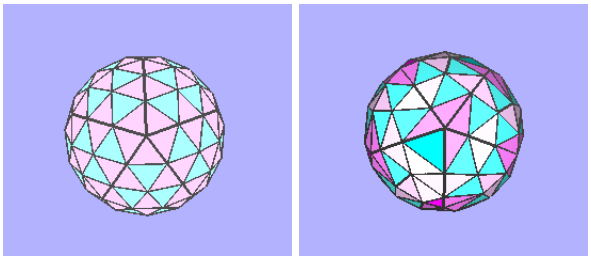

(a) Current Approach (b) Adjusted Midpoint

Figure 6.18: Split (1:9) Issues. Modification causes edge matching issues around odd-valence vertices

6.3.6 Catmull-Clark Split

The Catmull-Clark Split is presented in Figure 6.19. Notice how the size of the subdivision surface alternately increases and decreases, rather than a consistent decrease in radius as it converges to a sphere. This is validated with the global preservation noted in Table 6.1. Since Catmull-Clark general constructs quadrilateral faces, and does not guarantee planarity, our adjustment at attempted spherical preservation ensures that these such faces are indeed not planar. Consequently, our mathematics from Section 5.1.2 do not hold.

As a result, the expected area for a given face is altered due to its calculation. Since it relies on the preservation of the original global area, the expected face area is not preserved as the whole of the surface area is also unpreserved. The first subdivision (Figure 6.19(a)) illustrates

that as a result of the whole surface increasing in size, all faces are resultantly larger than their expected area. This holds in the third and fifth subdivision level, while the second and fourth have a closer areal error balance. This latter balance is validated by the results in Table 6.1, where the area is only scaled by 11%, rather than in excess of 50%.

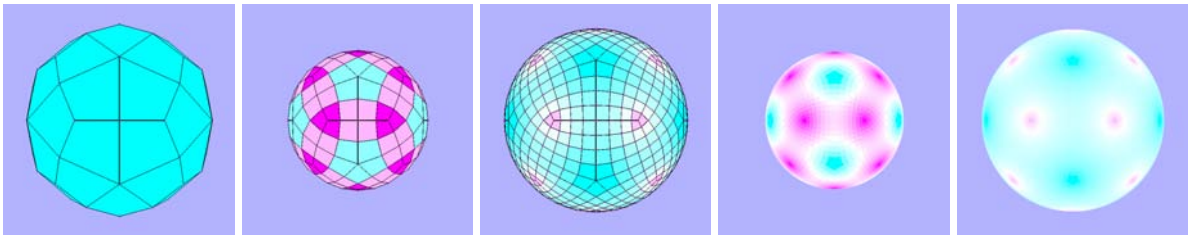

(a) Level 1 (b) Level 2 (c) Level 3 (d) Level 4 (e) Level 5

Figure 6.19: Area Highlighting - Catmull-Clark Split

As a consequence of this lack of global area preservation, the distribution of areal face error varies, as illustrated in the graph of Figure 6.20. The face errors are well distributed, ranging from great reductions to great increases.

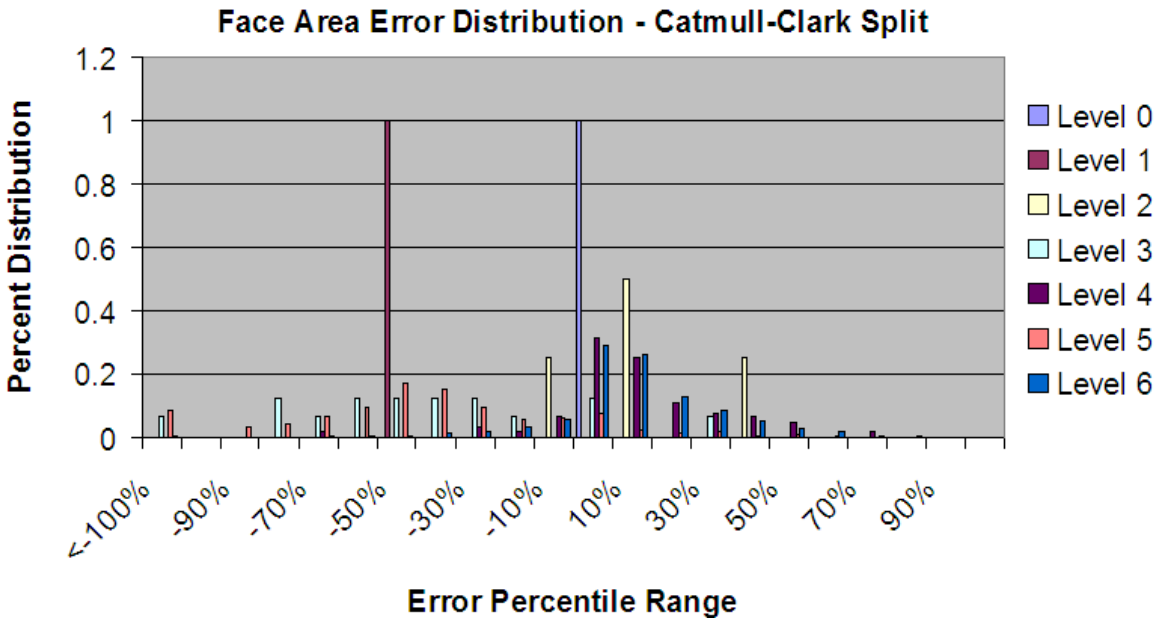

Figure 6.20: Face Area Error Distribution - Catmull-Clark Split

6.3.7 Split (1:3)

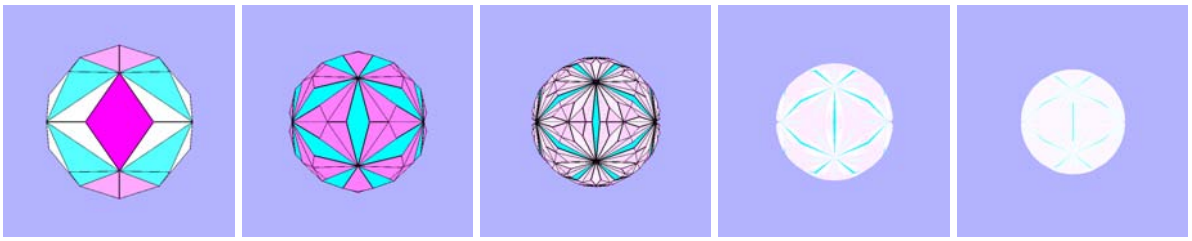

(a) Level 1 (b) Level 2 (c) Level 3 (d) Level 4 (e) Level 5

Figure 6.21: Area Highlighting - Split (1:3)

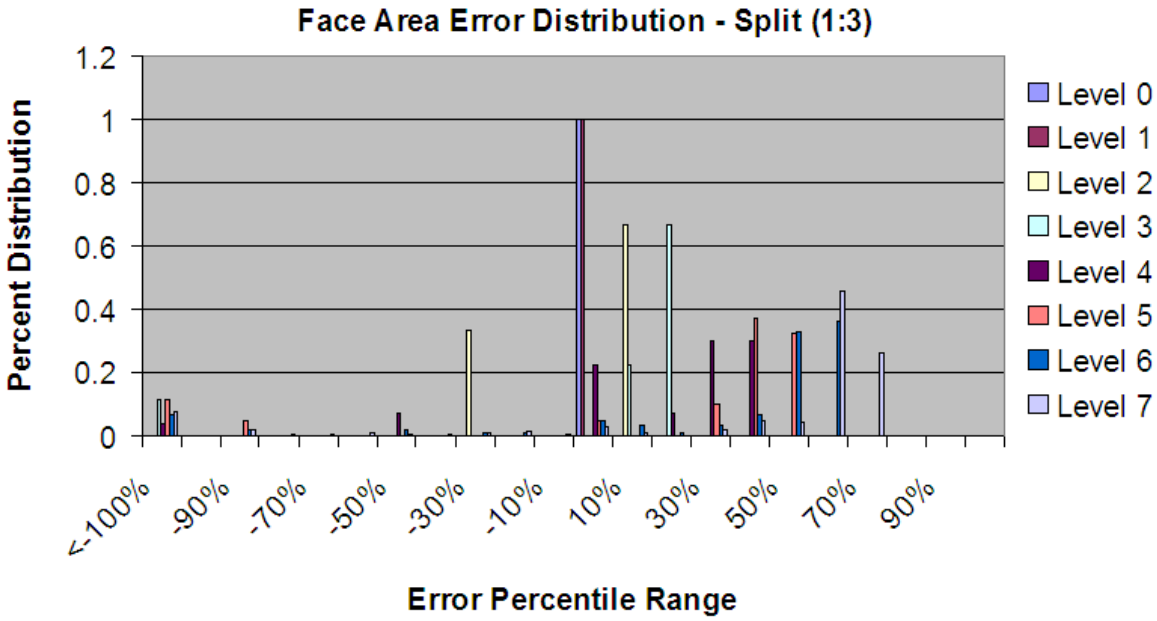

Figure 6.22: Face Area Error Distribution - Split (1:3)

Split (1:3), at first glance, appears to perform well at preserving local area as visualized in the uniformity of face colouration in Figure 6.21. This is supported by the maximal face error of 0.00% from Table 6.6. Unfortunately, what is taking place is that the blue faces, especially in subsequent subdivision levels, are constructing concave attributes. From the construction of the subdivision, visualized in Figure 6.6, all vertices from the previous subdivision level remain in place, with edges unbroken. As subsequent interior vertices and edges are introduced, the

original edges remain locked in place. Thus, as the new vertices more closely approach a sphere's surface, the original edges appear to cave in towards the center. These original edges influence the adjacent faces, which are the resulting blue. Consequently the maximal error consistently increases as a proportion of the faces, as illustrated in Table 6.6.

As a result, the distribution of error across the various percentiles is wide and varied (Figure 6.22). The faces quickly range from high to low error. The visualization in Figure 6.21, however, displays white, due to the overwhelming maximal error from the blue faces, and its impact on the resulting colouration scale.

6.3.8 Split (1:6)

The distribution of face error for Split (1:6) is presented in Figure 6.23. Notice for Level 1, how faces are either white, or a very dark shade of purple or blue. This corresponds to the maximal face error reported in Table 6.6, namely 0.00%. The symmetric construction of faces from the original equilateral triangle results in a nominal deviation in area between faces. It is only on subsequent subdivision when symmetric faces are unable to be created. The distribution of error is seemingly patterned, with a decrease in area along the interior midlines of what would be the original face, and an increase within the corners (Figure 6.23(e)).

(a) Level 1 (b) Level 2 (c) Level 3 (d) Level 4 (e) Level 4 Face

Figure 6.23: Area Highlighting - Split (1:6)

Recall that the Split (1:6) has a maximal face error has less than 10% areal deviation, through all levels of subdivision. The resulting distribution consequently resides around the 0.0% percentile, as confirmed in Figure 6.24.

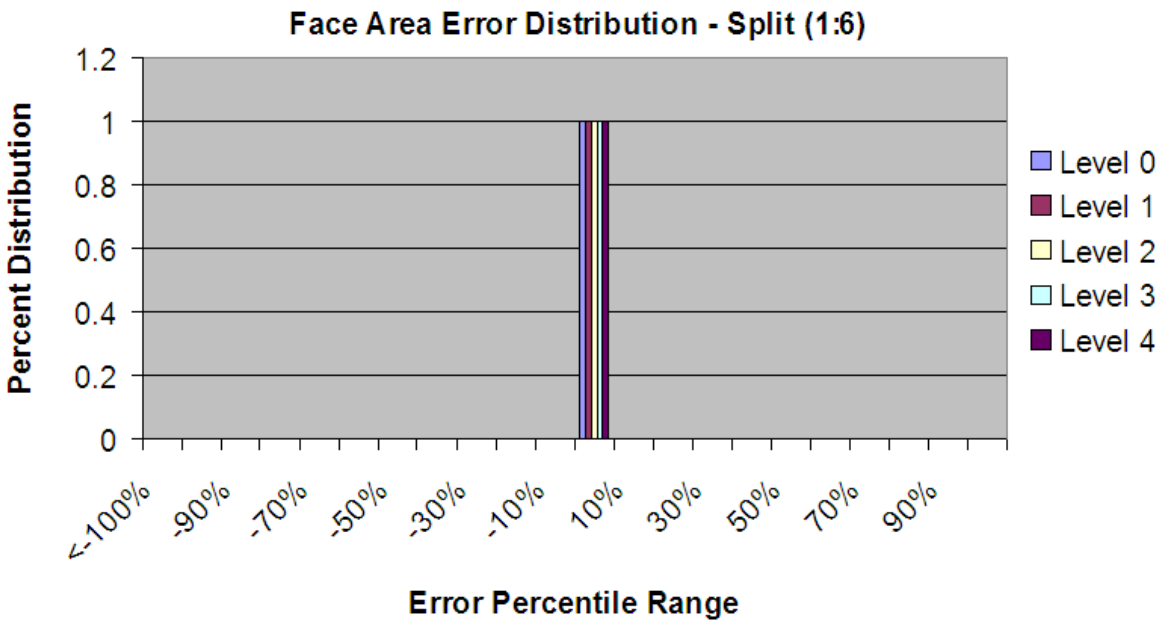

Figure 6.24: Face Area Error Distribution - Split (1:6)

It is recognized that with this approach, there is only one degree of the freedom: the mid point. If the midpoint is not placed precisely in the middle, the discrepancy between cells increases. Consequently, this does not actually offer a mechanism within which face area can be adjusted for local preservation.

6.3.9 Split (1:5)

Split (1:5), unlike the other approaches, begins its construction through the use of the dodecahedron. Unlike the traditional Doo-Sabin subdivision, only hexagons are used to replace the vertices. The original faces generally retain their shape (as much can be accomplished through the global preservation adjustment), but are reduced in size, and rotated to accommodate the adjacent elements.

From Table 6.1, we acknowledged the inability of this approach to retain global area preservation. Approximately 20% increase from the desired surface area results. Similar to Catmull-Clark , the resultant faces are not triangles, and consequently non-planar within the sphere. As

a result, the mathematics supporting the global preservation are unable to hold.

This overall increase in size is illustrated by the distribution of error across the faces. Notice how across all subdivision levels presented in Figure 6.25 the majority of faces are a shade of blue. This indicates a general increase in area over the original surface.

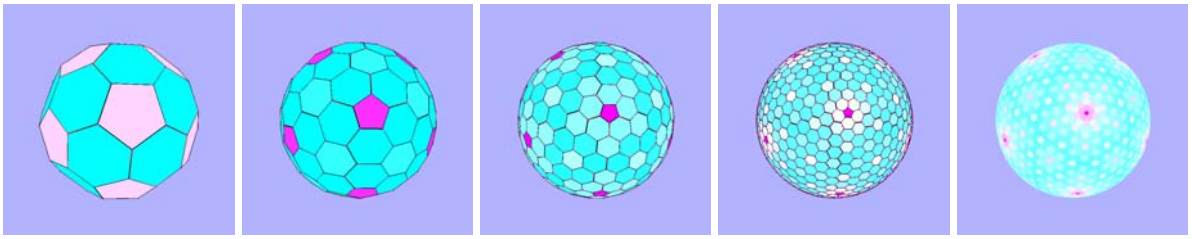

(a) Level 1 (b) Level 2 (c) Level 3 (d) Level 4 (e) Level 5

Figure 6.25: Area Highlighting - Split (1:5)

This deviation in area is visualized in Figure 6.26. Notice how the range of error is not localized to the 0.0%, but instead varies. This corresponds to the maximal face error of 20% as mentioned above, and the visual distribution as presented in Figure 6.25.

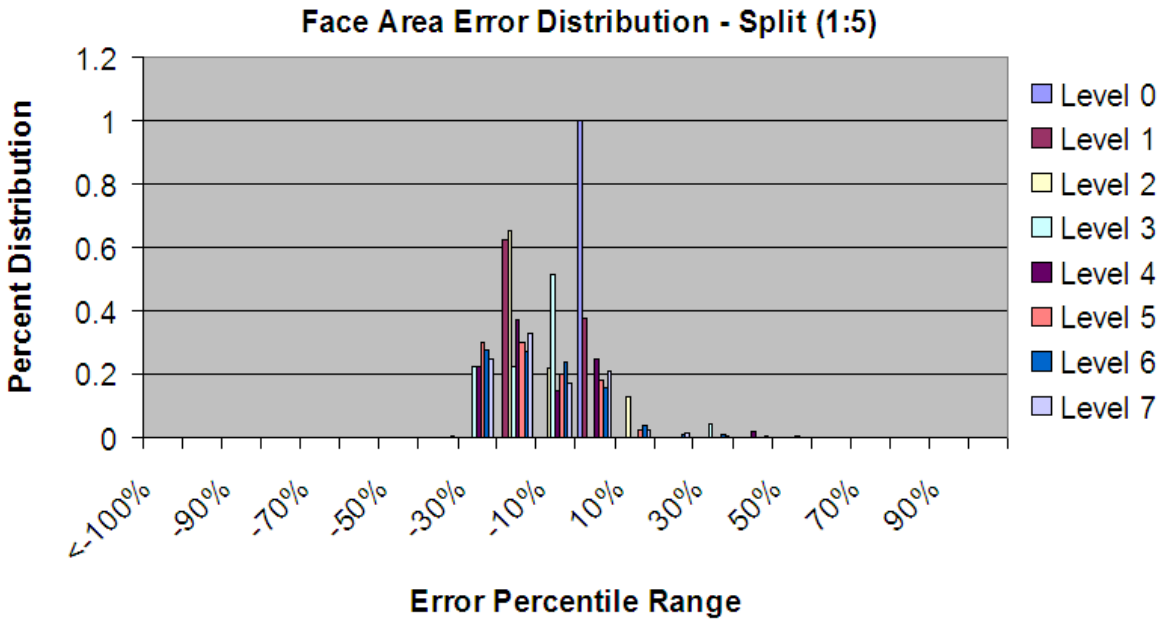

Figure 6.26: Face Area Error Distribution - Split (1:5)

6.3.10 Split (2:6)

Split (2:6) error distribution is presented in Figure 6.27. As with Split (1:6), its first subdivision undergoes a symmetric construction of faces. Thus, with a resulting maximal error of 0.00% (Table 6.6), Figure 6.27(a) merely highlights the nominal deviation between the faces.

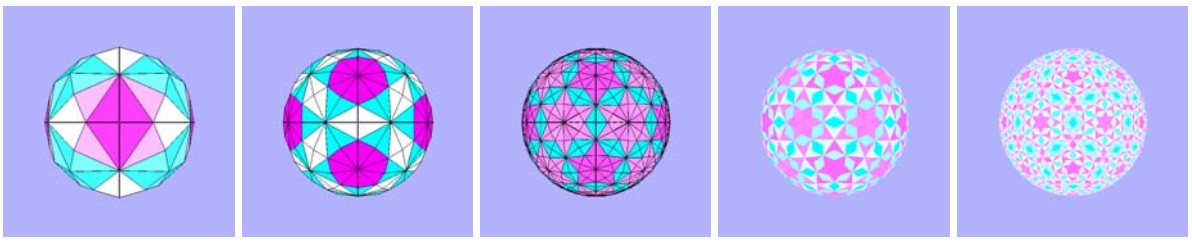

(a) Level 1 (b) Level 2 (c) Level 3 (d) Level 4 (e) Level 5

Figure 6.27: Area Highlighting - Split (2:6)

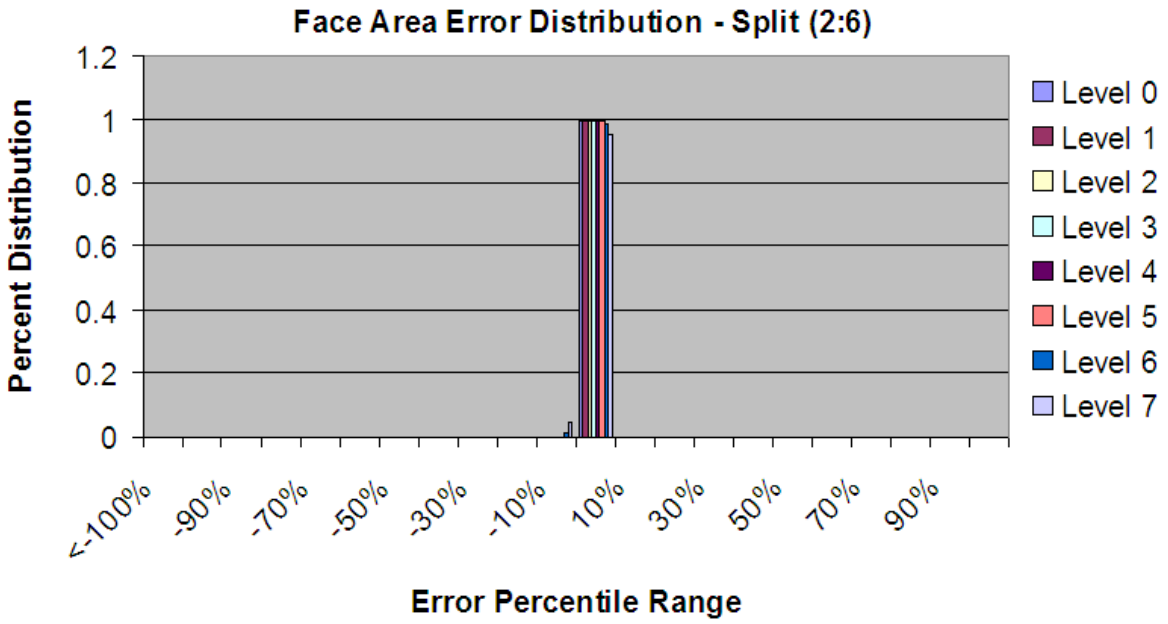

Figure 6.28: Face Area Error Distribution - Split (2:6)

Due to the subdivision process, the distribution of error across the faces, while patterned, appears to be more evenly distributed than the other approaches. The fifth level of subdivision in Figure 6.27(e) indeed has epicentres of increase and decrease, but these are surrounded by

nominal deviations in area.

The distribution of face error is presented in Figure 6.28. As with prior approaches, the majority of the faces undergo nominal deviation from the 0.0% percentile. Only at Levels 6 and 7 do we observe a slight introduction of a deviation outside of this percentile. This corresponds to the maximal face error exceeding 10% during subdivision levels 6 and 7, as presented in Table 6.6.

6.4 Discussion

Based on the above results, a number of conclusions may be drawn. An evaluation of global area preservation is presented. For the local area preservation, each of the respective subdivision approaches are explored and evaluated.

6.4.1 Global Area Preservation

From Table 6.1, we observed that the surface area of Split (1:4), Split (1:9), Split (1:3), Split (1:6), and Split (2:6) all maintained approximately zero deviation from the original surface. This validates the approach demonstrated in Section 5.1.2, illustrating the simultaneous reliance and independence of the angular definitions of face's area.

Split (1:5) and the Catmull-Clark Split, however, were unable to maintain this area preservation property, deviating greatly from the original surface area. Since these approaches construct non-triangular faces, this impacts the global preservation approach two fold. The first, in conjunction with the convergence to a sphere, results in non-planar faces. Consequently, the calculation of area differs, and may even be indefinable. Secondly, the non-triangular nature of the faces impacts the formula for area calculation. Thus, even if it were to be planar in nature, it may no longer have the simplicity of formulae presented in Section 5.1.2. As such, this global area preservation approach may be applied only to triangular meshes.

6.4.2 Local Area Preservation

Of the approaches discussed, only a few present viable options for preserving local area through the employment of the global area preservation in conjunction with the subdivision process. Those that are unable to retain global area are consequently unable to preserve the local area, while those that retain global area vary depending on the methodology of the subdivision construction.

The Catmull-Clark Split, modified to ensure convergence to a circumscribed sphere, is unable to retain global area due to the non-planar, non-triangular faces inherent with the approach. Furthermore, the faces quickly become drastically asymmetric as a consequence of the spherical projection in conjunction with the construction. As visualized in Figure 6.19, the vertices of the original faces result in larger faces than those created within the interior. While this offers a good evaluation of the global and local preservation for a standard subdivision approach, it is unable to meet our requirements.

Similarly, Split (1:5) is unable to retain the global area. Again, this comes as a result of the non-planar, non-triangular faces, as discussed in the description for this approach. The faces, however, once created generally evenly distribute the error of the faces. The exception is the regularly occurring pentagons, and their nearby faces. Figure 6.25(e) illustrates this relatively smooth distribution of error across the face. The maximal deviation occurs within the pentagonal faces. While this approach could offer local area support (with the exception of the pentagons), the inability to preserve global area renders it insufficient.

Within the remaining five approaches, it is particularly important to consider the maximal face area error in its percentage and absolute form. Split (1:9), for example, converges to 1 km of the Earth's surface within three levels of subdivision, attaining a maximal areal error of 9.67%. However, in terms of absolute area, this means that for the face, or faces, having this maximal error, a deviation of 3,380 km^2 is present. Similarly, Split (1:6), converges within four subdivision levels, achieving a maximal error of 8.84%. This value transforms to 1,740

km^2. Figure 6.16 seems to illustrate that the 1:9 approach results in the majority of the faces retaining this maximal error, due to the darkness of the faces. Figure 6.23, on the other hand, seems to more evenly distribute and vary the error across the resultant faces. Ranging from an increase in the interior, to concentrated decreases within the corners of the original face, this Catmull-Clark Split offers a better alternative over the Split (1:9).

Split (1:4) and Split (1:3) approaches not only take longer to converge within 1 km of the circumscribing sphere - or not at all, in the case of Split (1:3) - but also maintain a higher maximum face error. As detailed in Table 6.1, Split (1:4) - adjusted to converge to a sphere - maintains a fairly steady 20% maximal face error. At the sixth level of subdivision, this equates to 1,290 km^2 on a given face. From the visualization of the distribution in Figure 6.14, and again through the concentration of distribution in Figure 6.15, this maximal error is almost uniformly present across the surface of the resulting polyhedron. Alternating between regions of increase and regions of decrease, this is undesirable for a full solution.

Comparably, Split (1:3) has a drastic increase in areal error as the level of subdivision increases. Table 6.1 indicates a maximal 8,525% face error at the seventh level of subdivision. This is a direct consequence of the construction of this otherwise simplistic subdivision. Since edges remain intact, and are not broken up, an original edge acts as a chord through the eventually sphere-approximating surface, as visualized in Figure 6.21. The slivers of blue represent the faces adjacent to this chord, which ends up constructing a concave surface. While this still retains global area preservation, it does not offer a valid approach for preserving the local area. This is even further emphasized by pointing out the extremely non-uniform resultant faces, as visualized in Figure 6.21(c). Furthermore, the maximal displacement between a resultant face and the sphere in which it is aiming to circumscribe is 9.96%. This correlates to the faces which have an original edge along their boundary. The resultant face, interior to the increasingly convex surface, continues to be drastically displaced from the otherwise spherical approximation.

It should be noted that as a result of these constructions, Split (1:3) and Split (1:6) increase the valence of vertices throughout their subdivision process. This property is traditionally not ideal within the field of computer graphics, as it constructs long, skinny triangles challenging to render out to the screen. While Split (2:6) initially increases the valence of the vertices, it does not exceed 12 edges, due to the construction of the approach.

Lastly, Split (2:6), as with Split (1:4), takes longer to converge to within 1 km of the circumscribing sphere. However, its gradual increase in face error still equates to a reduction in absolute error. As illustrated, although Table 6.6 shows an increase in percentage, this is reflected in Table 6.7 as a decrease. In fact, at the seventh subdivision level, the maximal error equates to a mere 970 km^2, or roughly the size of San Diego, California, USA. Since the global area continues to be preserved, it is the occasional face that is absorbing this resultant error. As visualized in Figure 6.27, the error in area is fairly evenly distributed across the surface, with pockmarks of increases, surrounded by regions of neutrality. Only Split (1:6) compares with this distribution visualization. Split (2:6) also prevents drastic valence increase.

Through this, since Split (2:6) does the best at meeting the desired criteria. Not only does it preserve global area, but it also has a limited maximal local error. This limited local area error in turn is nicely distributed across the surface of the sphere. Finally, its valence stabilization makes it additionally preferable for subdivision. Consequently Split (2:6) offers the best option presented for preserving local area across the surface of a sphere-converging surface.

Chapter 7

Future Work

While the work herein has answered a number of questions relating to the computational potential for projection improvement, and the validity of marginal adjustments to existing subdivision for equal area spherical subdivisions, it has induced a range of subsequent research.

The motivation for this project has been the establishment of a multiresolution representation for the Earth's surface. As such, data can then be associated, and due to the hierarchical, and neighbouring information, be easily employed for data analysis. Given the success of Split (2:6), transference to an indexed regime would be the next logical step. Of note, this is currently under exploration through our research group. A supplemental addition - which lends itself further from the current graphic exploration - is the adaptation of this particular structure to a sufficiently adequate database. Within the framework of such an indexing system, and the overall subdivision representation, such a database should be able to store the vast amounts of information available through the requisite sources.

An emphasis on employing multiresolution approaches to the existing framework is of particular interest. The work shown herein demonstrates the ability to subdivide the surface to a sphere-converging representation, while simultaneously preserving area. Given this ability, reversing the approach, akin to other reverse subdivision schemes [43, 44], we could increase the support for data extraction, while reducing internal representation and storage requirements. By applying such techniques, one could quickly transition between a close visualization of, say, a city, and zoom out to a broader context without increasing storage requirements. Every increase in speed, and reduction in storage, means the ability to show more, in a reduced time frame.

Alternative subdivisions may also be considered, and constructed to ensure equi-areal

preservation. Exploring a grid system, inspired by the traditional latitude and longitude coordinate system may be feasible. For a given number of lines of latitude, and a given number of lines of longitude, a polyhedral approximation could be constructed. The placement of the lines of latitude could be adjusted so that the resulting quadrilaterals would have equi-areal properties. It is yet unclear as to whether such an approach would be efficient, or could be sufficiently defined mathematically as to efficiently construct a subdivision for any number of latitude or longitude lines.

It is of worth to point out that the Earth is not a perfect sphere. At its coarsest, the equator bulges out more than it should, and at its most detailed, the mountains, valleys, and other geographic features negate this otherwise spherical model. Cartographers have traditionally acknowledged, and generally proceeded in spite of this fact. However, with the improved computation power, we have the potential to address this issue. Instead of presenting the Earth in its more usual spherical representation, a true visualization becomes increasingly possible. In order to incorporate this within the proposed system, the equal-area preservation must be considered for non-spherical converging surfaces. This is briefly discussed in Section 5.1.3. An expansion on this, and a more concrete support of surfaces will better facilitate this non-spherical representation of the Earth. Again, this added information will improve the quality of data, and reduce the error during analysis.

Chapter 8

Conclusion

With the advent of modern computing, and its increased computation power, there is a push to present data through its interactive and ubiquitous medium. As such, there is the potential to visualize cartographic data in its spherical form, rather than the traditional two dimensional maps they have historically occurred on. In particular, this alternative representation lends itself to a more interactive approach, enabling the rotation, translation and scaling of otherwise static world data. With the employment of a true spherical representation of the Earth's surface, there is the drive to eliminate the distortions that classical map projections have induced.

At the very least, this computational employment can improve the speed at which data is transformed between spherical and planar coordinates as through the traditional projections. This is observed with our optimization of the Inverse Snyder Polyhedral Projection. This particular type of projection is particularly meaningful as it uses a polyhedron, rather than a planar map upon which to project the Earth's data. Due to its improved approximation of the Earth's spherical surface, it lends itself to the reduction of distortions.

Expanding on this polyhedral approximation concept, we present the application of computer graphics subdivision techniques to define a graphical representation of the Earth's surface. Starting with a coarse polyhedral approximation of a sphere, we iteratively increase the number of faces, which in turn improves the representation of a spherical surface. If the subdivision is well construction, neighbourhood and hierarchical information may remain present - beneficial for information traversal. Additionally, for the purposes of our research, defining an equal area subdivision is instrumental in supporting the scientific community at large.

Within this work, a global preservation approach is presented and validated. It demonstrates that for triangular meshes, and triangle-generating subdivisions, the area may be preserved.

The definition also ensures that the original sphere-circumscribing polyhedron will remain inscribed within a sphere - though at an alternative radius to ensure areal preservation. As the faces increase, while spherical circumscription is maintained, this increasingly improves its spherical approximation. From our validating approaches, there exist subdivisions which can approximate the Earth's surface to within 1 km by their third iteration.

This coarse definition of areal equivalence does not lend itself to the true meaning of area preservation, adopted from projection nomenclature. As such, a region bounded on the sphere must have equivalent area to that on the planar representation. Within our approach, we explore various subdivisions and determine their viability at supporting local area preservation.

Applying the aforementioned global area preservation throughout the subdivision iterations, we observe several informative attributes. For instance, surfaces which have non-planar, or non-triangular faces, and consequently are not supported by the global area preservation approach, are not viable for this approach to equal area spherical subdivision. This is due to the inconsistency of areal calculation for a non-planar face.

Additional approaches pointed towards a need for breaking edges to prevent inadvertent creation of a convex surface, as illustrated in Split (1:3). This particular subdivision, along with Split (1:6), rapidly increases the valence of vertices on the constructed surfaces. This does not lend itself to effective support of a hierarchical data structure, nor of an ease of visualization, due to the creation of thinly shaped faces.

Both Split (1:4) and Split (1:9) support and validate the global area preservation, but due to their inflexibility of adjustment, and splitting mechanism, end up becoming very polar in their distribution of error across the faces. Consequently, either regions greatly increase the error, or they greatly decrease the error, with nominal non-influencing regions. The Split (1:4) approach, for example, maintains a maximal face areal error of approximately 20% for the first seven levels of subdivision. While Split (1:9) achieves a maximal face areal error of 9.67%, this corresponds to 3,380 km^2 - a significant distance - due to the early convergence to 1 km of

the Earth's surface.

By far, the most suitable approach explored is Split (2:6), as presented in Section 6.1.7. While the face error - initially kept low due to the symmetric nature of the subdivision - achieves a maximal value of 16.3%, this corresponds to 970 km^2, or a region the size of San Diego, California. The visualization of the distribution (Figure 6.27) illustrates that on the whole, this maximal error occurs uniformly and sporadically throughout the subdivision process, with the remainder of the surface having a predominately negligible areal error.

Through this development in incorporating computer graphics techniques to the visualization and analysis of geographic information, we can continue to improve upon the representation and interaction with positionally-based data. As mentioned in Chapter 7, future work involves the incorporation of an indexing method for associating these constructed surfaces with data. Through modern advancements in data bases, or alternatively, direct storage within the subdivision, data can more effectively be associated to the visual structure. This will assist in the determination of neighbouring data, and due to the subdivision approach, hierarchically-based elements.

While technologies are always improving, its more recent advances make it increasingly accessible to fields of study which have gone otherwise untouched in generations. To this end, the direction of this, and comparable work, will see technological improvements towards geographic representation. The majority of cartographers construct maps - increasingly with GIS tools - with a specific purpose in mind. As such, they employ previously constructed projections, rather than developing new ones. The employment of computer graphics, and in particular this subdivision approach, offers a novel way of representing the Earth's surface, while still supporting the equal area feature desirable during map development.

Bibliography

[1] Map Projections with Tissot's Indicatrix. http://commons.wikimedia.org/w/index.php?title=Category:Map_projections_with_Tissot%27s_indicatrix.

[2] Short Proceedings of the 1st European Workshop on Reference Grids. 2003.

[3] M. Abramowitz and I. A. Stegun, editors. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Dover, 9 edition, 1972.

[4] R. H. Bartels and F. F. Samavati. Reversing Subdivision Rules: Local Linear Conditions and Observations on Inner Products. *Journal of Computaional and Applied Mathematics*, 119(1-2):29–67, 2000.

[5] W. H. Beyer. *CRC Standard Mathematical Tables*. CRC Press, 28 edition, 1987.

[6] P. Borwein and T. Erdlyi. *Polynomials and Polynomial Inequalities*, pages 8, 29–41. Springer-Verlag, New York, NY, 1995.

[7] A. Bradley. Equal-Area Projection on the Icosahedron. *Geographical Review*, 36(1):101–104, 1946.

[8] M. D. Buhmann. *Radial Basis Functions*. Cambridge University Press, 2003.

[9] J. Burkardt. Data Files. http://people.sc.fsu.edu/~burkardt/data/obj/icosahedron.obj.

[10] F. Canters. *Small-Scale Map Projection Design*. Taylor & Francis, New York, NY, 2002.

[11] J. C. Carr, W. R. Fright, and R. K. Beatson. Surface Interpolation With Radial Basis Functions for Medical Imaging. *IEEE Transactions on Medical Imaging*, 16:96–107, 1997.

[12] E. Catmull and J. Clark. Recursively Generated B-spline Surfaces on Arbitrary Topological Meshes. *Computer-Aided Design*, 10(6):350 – 355, 1978.

[13] Commission for Environmental Cooperation. North America. `http://atlas.nrcan.gc.ca/site/english/maps/archives/various/north_america_cec`, 2004. [Online; accessed 01-May-2011].

[14] C. I. de l'Eclairage (CIE). Guide to Recommended Practice of Daylight Measurement. `http://www.cie.co.at/publ/abst/108-94.html`, 1989.

[15] B. Delaunay. Sur la sphre vide, Izvestia Akademii Nauk SSSR, Otdelenie Matematicheskikh i Estestvennykh Nauk. 7:793–800, 1934.

[16] T. DeRose, M. Kass, and T. Truong. Subdivision Surfaces in Character Animation. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '98, pages 85–94, New York, NY, USA, 1998. ACM.

[17] D. Doo. A Subdivision Algorithm for Smoothing Down Irregularly Shaped Polyhedrons. *Proceedings on Interactive Techniques in Computer Aided Design*, pages 157–165, 1978.

[18] D. Doo and M. Sabin. Behavior of Recursive Division Surfaces Near Extraordinary Points. *Computer-Aided Design*, 10:356–360, 1978.

[19] A. Durer. *Unterweysung des Messung mit dem Zirckel und Richtscheyt*. 1538.

[20] N. Dyn, D. Levin, and J. A. Gregory. A Butterfly Subdivision Scheme for Surface Interpolation with Tension Control. *ACM Transactions on Graphics*, 9:160–169, 1990.

[21] D. R. Forsey and R. H. Bartels. Hierarchical B-Spline Refinement. *Comput. Graph.*, 22(4):205–212, 1988.

[22] A. Girard. *Invention nouvelle en algebra*. 1629.

[23] Google. Google Earth. http://www.google.com/earth/index.html, 2011. [Online; accessed 20-November-2011].

[24] Google - Imagery. Google Maps. http://maps.google.ca, 2010. [Online; accessed 11-December-2010].

[25] A. Gore. The Digital Earth: Understanding our Planet in the 21st Century. California Science Center. Los Angeles, USA, 31, Jan 1998.

[26] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann. HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere. *The Astrophysical Journal*, 622(2):759–771, Apr. 2005.

[27] S. J. Gortler and M. F. Cohen. Hierarchical and Variational Geometric Modeling with Wavelets. *Proceedings of the 1995 Symposium on Interactive 3D graphics, I3D '95*, pages 35–42, 1995.

[28] G. Guidi, F. Remondino, M. Russo, F. Menna, A. Rizzi, and S. Ercoli. A Multi-Resolution Methodology for the 3D Modeling of Large and Complex Archeological Areas. *International Journal of Architectural Computing*, 7:39–55, 2009.

[29] E. Harrison, A. Mahdavi-Amiri, and F. Samavati. Optimization of Snyder's Inverse Polyhedral Projection. In *International Conference on CyberWorlds*, pages 1–8. IEEE, 2011.

[30] L. Kettner. Using Generic Programming for Designing a Data Structure for Polyhedral Surfaces. *Comput. Geom. Theory Appl.*, 13:65–90, 1999.

[31] A. J. Kimerling, K. Sahr, D. White, and L. Song. Comparing Geometrical Properties of Global Grids. *Cartography and Geographic Information Science*, 26(4):271–288, October 1999.

[32] L. Kobbelt, S. Campagna, J. Vorsatz, and H.-P. Seidel. Interactive Multi-Resolution Modeling on Arbitrary Meshes. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '98, pages 105–114, New York, NY, USA, 1998. ACM.

[33] J. H. Lambert. *Anmerkungen und Zustze zur Entwerfung der Land- und Himmelscharten. Mit 21 Textfiguren*. 1772.

[34] L. Lee. Conformal Projections Based on Elliptic Functions. *Cartographica. Monograph no. 16, supplement no. 1 to Canadian Cartographer*, 13:128, 1976.

[35] C. Li. Application of Catmull-Clark Subdivision Method in Plastic Injection Mould Parting Surface Design. In *Information Visualisation, 2002. Proceedings. Sixth International Conference on*, pages 477 – 482, 2002.

[36] C. Loop. *Smooth Subdivision Surfaces Based on Triangles*. M.S. Mathematics Thesis, University of Utah, 1987.

[37] C. Mandal, B. Vemuri, and H. Qin. Shape Recovery Using Dynamic Subdivision Surfaces. In *Computer Vision, 1998. Sixth International Conference on*, pages 805 –810, Jan 1998.

[38] NASA. Blue Marble. `http://visibleearth.nasa.gov/`, 2000. [Online; accessed 11-December-2010].

[39] I. Newton. Methodus fluxionum et serierum infinitarum, 1664-1671.

[40] PYXIS Innovation Inc. How PYXIS Works. `http://www.pyxisinnovation.com/pyxwiki/index.php?title=How_PYXIS_Works`, 2011. [Online; accessed 01-May-2011].

[41] D. Rosca and G. Plonka. Uniform Spherical Grids via Equal Area Projection from the Cube to the Sphere. *J. Comput. Appl. Math.*, 236:1033–1041, 2011.

[42] J. Rossignac and P. Borrel. Multi-Resolution 3D Approximations for Rendering Complex Scenes. *Geometric Modeling in Computer Graphics*, pages 455–465, 1993.

[43] J. Sadeghi and F. F. Samavati. Smooth Reverse Subdivision. *Computers & Graphics*, 33(3):217–225, 2009. IEEE International Conference on Shape Modelling and Applications 2009.

[44] J. Sadeghi and F. F. Samavati. Smooth reverse Loop and Catmull-Clark Subdivision. *Graphical Models*, 73(5):202 – 217, 2011.

[45] K. Sahr, D. White, and A. J. Kimerling. Geodesic Discrete Global Grid Systems. *Cartography and Geographic Information Science*, 30(2):121–134, 2003.

[46] F. F. Samavati and R. H. Bartels. Multiresolution Curve and Surface Representation: Reversing Subdivision Rules by Least-Squares Data Fitting. *Computer Graphics Forum*, 18(2):97–119, Jun 1999.

[47] J. P. Snyder. *Map Projections - A Working Manual: U.S. Geological Survey Professional Paper 1396*. U.S. Government Printing Office, Washington, DC, 1987.

[48] J. P. Snyder. An Equal-Area Map Projection for Polyhedral Globes. *Cartographica*, 29(1):10–21, Spring 1992.

[49] J. P. Snyder. *Flattening the Earth: Two Thousand Years of Map Projections*. University of Chicago Press, Chicago, Illinois, USA, 1997.

[50] J. P. Snyder and P. M. Voxland. *An Album of Map Projections*. U.S. Geological Survey, 1989.

[51] L. Song, A. J. Kimerling, and K. Sahr. *Discrete Global Grids*, chapter Developing an Equal Area Global by Small Circle Subdivision. National Center for Geographic Information and Anlaysis, Santa Barbara, California, USA, 2002.

[52] E. J. Stollnitz, T. D. Derose, and D. H. Salesin. *Wavelets for Computer Graphics: Theory and Applications, publisher = Morgan Kaufmann Publishers, location = San Francisco, USA, year = 1996*.

[53] N. A. Tissot. *Memoire sur la representation des surfaces et les projections des cartes geographiques*. Gauthier-Villars, 1881.

[54] P. Tregenza. Subdivision of the Sky Hemisphere for Luminance Measurements. *Lighting Research and Technology*, 19.

[55] United States Department of the Interior. Map Projections: From Spherical Earth to Flat Map. `http://www.nationalatlas.gov/articles/mapping/a_projections.html`, 2011. [Online; accessed 01-May-2011].

[56] U.S. Government. U.S. Census. `http://www.census.gov/geo/www/gazetteer/files/Gaz_places_national.txt`.

[57] D. van Leeuwen and D. Strebe. A Slice-and-Dice Approach to Area Equivalence in Polyhedral Map Projections. *Cartography and Geographic Information Science*, 33(4):269–286, 2006.

[58] L. Wanhammar, K. Johansson, and O. Gustafsson. Efficient Sine and Cosine Computation Using a Weighted Sum of Bit-Products. In *Proceedings of the 2005 European Conference on Circuit Theory and Design*, volume 1, pages I/139 – I/142, Aug 2005.

[59] J. Warren and H. Weimer. *Subdivision Methods for Geometric Design: A Constructive Approach*. Morgan Kaufmann, 2001.

[60] J. Werner. Nova translatio primi libri geographiaae C, 1514.

[61] J. Wheeler, P. Hoffman, K. Card, A. Davidson, B. Sanford, A. Okulitch, and W. Roest. Geological Map of Canada: Map 1860A, scale 1:5 000 000, 1996.

[62] D. White, A. Kimerling, and W. Overton. Cartographic and Geometric Components of a Global Sampling Design for Environmental Monitoring. *Cartography and Geographic Information Systems*, 19:5–22, 1992.

[63] D. White, A. J. Kimerling, K. Sahr, and L. Song. Comparing Area and Shape Distortion on Polyhedral-Based Recursive Partitions of the Sphere. *International Journal of Geographical Information Science*.