# Optimal Any-Angle Pathfinding on a Sphere 

Volodymyr Rospotniuk<br>VOLODYMYR.ROSPOTNIUK@90POE.IO Rupert Small<br>RUPERT.SMALL@90POE.IO

Ninety Percent of Everything, Portman House, 2 Portman Street, London, United Kingdom, W1H 6DU


#### Abstract

Pathfinding in Euclidean space is a common problem faced in robotics and computer games. For long-distance navigation on the surface of the earth or in outer space however, approximating the geometry as Euclidean can be insufficient for real-world applications such as the navigation of spacecraft, aeroplanes, drones and ships. This article describes an any-angle pathfinding algorithm for calculating the shortest path between point pairs over the surface of a sphere. Introducing several novel adaptations, it is shown that Anya as described by Harabor \& Grastien for Euclidean space can be extended to Spherical geometry. There, where the shortest-distance line between coordinates is defined instead by a great-circle path, the optimal solution is typically a curved line in Euclidean space. In addition the turning points for optimal paths in Spherical geometry are not necessarily corner points as they are in Euclidean space, as will be shown, making further substantial adaptations to Anya necessary. Spherical Anya returns the optimal path on the sphere, given these different properties of world maps defined in Spherical geometry. It preserves all primary benefits of Anya in Euclidean geometry, namely the Spherical Anya algorithm always returns an optimal path on a sphere and does so entirely on-line, without any preprocessing or large memory overheads. Performance benchmarks are provided for several game maps including Starcraft and Warcraft III as well as for sea navigation on Earth using the NOAA bathymetric dataset. Always returning the shorter path compared with the Euclidean approximation yielded by Anya, Spherical Anya is shown to be faster than Anya for the majority of sea routes and slower for Game Maps and Random Maps.


## 1. Introduction

Pathfinding on a two-dimensional Euclidean grid is a well-studied mathematical problem with a wide range of practical applications, from logistics to gaming. A fundamental limitation of many early path-finding algorithms in this domain is their inability to travel along any angle on the grid due to the search path being constrained to 8 degrees of freedom: a left, right, up or down step to the nearest neighbouring node, or some pairwise combination of these resulting in a diagonal step.

Any-angle pathfinding refers to all pathfinding methodologies which remove this artificial constraint to the points of the grid by permitting movement in arbitrary directions, many of these being variations on A* (Hart, Nilsson, \& Raphael, 1968), a breakthrough graph search algorithm which uses a heuristic function to determine the most promising directions to explore the graph and thereby to reach the target point. These include Field D* (Ferguson \& Stentz, 2005), Theta* (Nash et al., 2007), Block A* (Yap et al., 2011), Hierarchical Pathfinding A* (HPA*) (Botea, Muller, \& Schaeffer, 2004), and Accelerated A* (Sišlák, Volf, \& Pěchouček, 2009). Amongst other strategies, these algorithms take various novel
approaches to partitioning the search space or post-processing A* routes by removing nodes satisfying certain line-of-sight or collision avoidance criteria in order to selectively reduce the total route distance compared with A*. In so doing the resulting paths are shorter than the paths produced by graph-constrained approaches such as A* and so return routes which are closer to optimal. Still further modifications involve machine learning, for example whereby the measured and heuristic cost functions of A* are learned using a Recurrent Neural Network (RNN) (Wang et al., 2019) and myriad others which build not on A* but instead apply agent-based models with gradient descent optimising a loss function, such as Deep Q-learning (Panov, Yakovlev, \& Suvorov, 2018) or those utilising LSTMs on a training set, for instance one generated using A* (Sartoretti et al., 2019; Bency et al., 2019) or GANbased approaches (Soboleva \& Yakovlev, 2019). These Machine Learning approaches forfeit accuracy for speed and do not guarantee optimality.

Existing any-angle pathfinding algorithms for Euclidean space which guarantee optimality include Visibility Graphs (Lozano-Pérez \& Wesley, 1979), Tangent Graphs (Liu \& Arimoto, 1992), Continuous Dijkstra based approaches (Mitchell, Mount, \& Papadimitriou, 1987; Hershberger \& Suri, 1999) and of course Anya (Harabor \& Grastien, 2013). Anya capitalises on the fact that all turning points on an optimal route in Euclidean space will coincide with the corner points of obstacles, enabling the optimal path to be obtained via online calculation of root point-interval pairs. These are projected to form visibility cones, permitting Anya to search wide intervals of space in aggregate, rather than evaluating visibility between neighbouring point-pairs at each grid point of the world map. This mechanism of projecting into space using line-of-sight intervals is what makes Anya superior in speed to all known path-finding algorithms with an optimality guarantee and indeed, for some maps, faster than many without this guarantee (Uras \& Koenig, 2015).

For real-world applications demanding high accuracy over long-distance voyages it becomes imperative to extend the approach of Anya to spherical coordinates. This opens up an abundance of applications such as space, flight and sea voyage planning, all unsuited to being treated as an approximation of a route in Euclidean space. Transforming the world map representation with a Gnomonic projection whence all straight lines are great circle lines and vice versa, at first glance permits Anya to be implemented without modification; great-circle lines being straight lines in the Gnomonic representation of a world map. The reason this approach falls short of achieving optimality is discussed in Appendix A. Instead, using Anya as a foundation it is shown that a new algorithm, Spherical Anya, can be defined which inherits the beneficial properties of its progenitor and generalises the algorithm to routes in spherical geometry. To the extent possible, the fundamental concepts and definitions of Anya are preserved. The definition of a Grid in Spherical Anya remains almost the same, as do the definitions of intervals, intermediate points and intersections. Flat projections also remain virtually identical to flat projections in Anya, whereas search nodes, cone projections, successors and observability in Spherical Anya must be defined differently due to the great-circle geodesic of Spherical geometry. Most significantly, the foundational theorem on which Anya depends - that of all turning points on an optimal any-angle path being corner points - no longer holds in Spherical geometry. A new type of non-observable cone projection must be introduced as a result, without which not all intervals in the Spherical world map would be explored. This ensures that the optimality, completeness and termination guarantees established for Anya extend to Spherical Anya.

## 2. Spherical Anya Components and Definitions

Anya pathfinding on a flat Euclidean grid takes a distinctly different approach to A-star varieties of pathfinding which restrict travel to discrete pairwise-visible points. Instead, the approach taken by Anya capitalises on the fact that all turning points in optimal Euclidean paths are corner points (Mitchell et al., 1987) making the fundamental unit of abstraction an interval of points $[a, b]$, being either traversable or non-traversable, as opposed to a sequence of discrete nodes. It is in this way, by defining the world map width $W$ height $H$, in terms of traversable and non-traversable intervals and subsequently generating cone- and flat- projections along these, that Anya will select the optimal traversable path through world map intervals without manually inspecting all rational points contained in the Farey Sequence of order $n=\min (W, H)$. The result is an optimal pathfinding algorithm in Euclidean space with superior speed.

To exhibit the algorithm defining Spherical Anya it is useful to revisit and transport the fundamental terminologies and components of Anya, these being used either directly or adapted for Spherical Anya, where the geodesic is a great-circle line. Wherever possible the foundational definitions (Harabor \& Grastien, 2013; Harabor et al., 2016) are maintained and adaptations and extensions are introduced only as necessitated by the novel properties of Spherical geometry versus Euclidean geometry.

For Spherical Anya the grid remains a planar subdivision of $W \times H$ cells. Rows are lines of constant longitude and vertical edges are lines of constant latitude, as shown in Figure 1. This means each cell in Spherical Anya is neither square in the Euclidean sense nor the Spherical sense; the vertical edges of the cell are not great-circle lines, albeit the horizontal edges of the cell are, and therefore constitute the shortest path between points on the same horizontal line, or row. Rows thus defined, obstacle edges - facing either the north or south poles of the sphere - are not great-circle lines with the exception of edges on the equator. Hence a great-circle line passing through the vertically aligned corner points of an obstacle defined on the grid will in general either pass exclusively within the non-traversable bulk of the obstacle or exclusively external to it, with the exception of the corner points themselves, as illustrated in Figure 2.

Of particular note is that the existence of a line of constant latitude between two points, in addition passing adjacent only to traversable cells, does not guarantee that these two points are visible, as illustrated for example in Figure 2(a). Instead the great-circle line, rather than a straight line in Euclidean space, underlies the definition of visibility in Spherical geometry.

Definition 1. Two points are said to be visible if they can be connected by a great-circle line which does not pass through an obstacle or the intersection formed by two diagonallyadjacent obstacles.

By fixing rows as lines of constant longitude, the definition of a grid interval in Spherical Anya can remain virtually identical as for Anya.

Definition 2. A grid interval $I$ is a set of contiguous and pairwise visible points drawn from any discrete row of constant longitude on the grid. Each interval is defined in terms of its endpoints $a$ and $b$. With the possible exception of $a$ and $b$, each interval contains only intermediate non-corner points.


Figure 1: An Equirectangular projection map with $x$-axis defining lines of constant latitude and grid rows of the $y$-axis giving great-circle lines of constant longitude. Defining a constant $y$-value as lines of longitude ensures flat projections remain virtually identical in Spherical Anya as in Anya. Longitudinal values define rows and pairs of latitudinal values, in combination with a row, define intervals.

The grid, intervals and visibility thus defined, the next building block of Anya is the search node. For Spherical Anya it becomes necessary to define not one but two categories of search node, standard and adjoint, the curvature of Spherical geometry making the second category necessary to guarantee completeness of the search space and thereby optimality of the resulting path.

Definition 3. A standard search node $(I, r)$ is a tuple where $r \notin I$ is a point called the root and $I$ is an interval such that each point $p \in I$ is visible from $r$.

A standard node here is much the same as a node in Anya albeit with the altered definition of visibility, being defined by the existence of a traversable great-circle path between the

(a) A great circle line between two points $a$ and $b$ of equal latitude in the southern hemisphere will always curve southwards, intersecting an obstacle with north-most edge upon this line. The symmetric claim holds likewise for a south-facing edge in the northern hemisphere.

(b) A great circle line between two points $a$ and $b$ of equal latitude in the northern hemisphere will always curve northwards, avoiding an obstacle with north-most edge upon this line. The symmetric claim holds likewise for a south-facing edge in the southern hemisphere.

Figure 2: Great-circle lines passing through two vertical corners of an obstacle will either intersect only with points internal to the bulk of the obstacle or only with points external to it, with the exception of the corner points $a$ and $b$ themselves.
root $r$ and all points in the interval $I$ as illustrated in Figure 3, with $I:=[a, b]$. The interval may be closed, as in this example, or alternatively open, or half-open. The kind of interval used will depend on later considerations (see Definitions 5 and 6). The second category of search node, which exists only in Spherical Anya, arises by virtue of the vertical line of an obstacle edge generally not being the shortest path between the two corner points on the vertical edge, as illustrated in Figure 4.

Definition 4. An adjoint search node $(I, r)$ is a tuple where $r \notin I$ is a point called the root, which is located on the vertical edge of an obstacle, and I is a traversable interval such that each point $p \in I$ is invisible from $r$ and furthermore the great-circle line from $r$ to $p$ intersects the vertical edge twice, and only the vertical edge, including once at $r$.

This second kind of search node, illustrated in Figure 4, is unique to Spherical Anya and is not required in Euclidean space where vertical edges describe straight lines and the turning points of an optimal route will always be corner points. Neither is true in Spherical geometry (see Section 3), where by convention the rows of the world map grid have been fixed as great-circle lines of constant longitude and the vertical edges of obstacles are lines of constant latitude.

The existence of adjoint nodes in Spherical geometry marks the most significant departure between the underlying mathematical structure of Spherical Anya compared to Anya. Thus far the differences have been limited; the path of shortest distance merely being substituted for the great-circle path of Spherical geometry, the orientation of the grid fixed explicitly to ensure rows are shortest-path lines of longitude, and the definition of


Figure 3: The standard node in Spherical Anya $[a, b], r$ is an interval $I:=[a, b]$ and a root point $r$ with $r \notin[a, b]$, where a traversible great-circle path exists between all points $p \in I$. Standard nodes form the first primary abstraction of the search space, all successors along a path from source to target destination consisting of a sequence of root points which are either root points of a standard node (definition 3) or an adjoint node (definition 4).


Figure 4: An adjoint node exists only in Spherical Anya due to the relative curvature between the vertical constant-latitude edge of obstacles and a great-circle line drawn between the two vertically aligned corner points of such an edge. An adjoint node is a node $[a, b), r$ where all cells adjacent to $[a, b)$ are traversable and furthermore all great circle lines through the root $r$ and any point $p \in[a, b)$ intersect the vertical edge of the obstacle twice. The opposite scenario exists in Spherical geometry too, where no great-circle line between the vertically aligned corners of an obstacle edge intersects the edge, excepting at the corner points.
visibility similarly adapted to fit the new context. However, with the addition of adjoint nodes the divergence becomes more fundamental. This divergence is formalised in Section 3 , after which it finally becomes practical to provide the concluding definitions required to implement Spherical Anya, those of successors and projections.

## 3. Optimality

Lemma 1. Not all turning points on an optimal path in spherical geometry are corner points.
Proof. By construction. It is assumed the obstacle edge is located in the northern hemisphere and facing southwards as illustrated in Figure 5. All results follow identically for the southern hemisphere and without loss of generality. For longitude, latitude $(\theta, \phi)$ and Gnomonic projection centred on the north pole ( $0, \frac{\pi}{2}$ ) we have projected coordinates

$$
\begin{align*}
x^{\prime} & =\frac{\cos (\phi) \sin (\theta)}{\sin (\phi)} \\
y^{\prime} & =-\frac{\cos (\phi) \cos (\theta)}{\sin (\phi)} \tag{1}
\end{align*}
$$

As the shortest path between any pair of points under the Gnomonic projection is a straight line we can define the last point on the obstacle which remains visible to $\mathbf{p}$ as $\mathbf{x}$, defined by coordinates $\left(\theta_{1}, \phi_{1}\right)$. Target point $\mathbf{p}$ takes coordinates $\left(\theta_{2}, \phi_{2}\right)$. The last point on the obstacle with direct line-of-sight to $\mathbf{p}$ must satisfy $\mathbf{x} \cdot(\mathbf{p}-\mathbf{x})=0$. Plugging into (1) yields $\mathbf{x} \cdot \mathbf{p}=\frac{\cos \left(\phi_{1}\right) \cos \left(\phi_{2}\right)}{\sin \left(\phi_{1}\right) \sin \left(\phi_{2}\right)} \cos \left(\theta_{2}-\theta_{1}\right)$ and $\|\mathbf{x}\|^{2}=\frac{\cos ^{2}\left(\phi_{1}\right)}{\sin ^{2}\left(\phi_{1}\right)}$. Solving for $\theta_{1}$ gives

$$
\begin{equation*}
\theta_{1}=\theta_{2}-\cos ^{-1}\left[\frac{\tan \left(\phi_{2}\right)}{\tan \left(\phi_{1}\right)}\right] \tag{2}
\end{equation*}
$$

This defines the point of departure from the obstacle edge. Now we show that the result which follows from applying string-pulling holds, namely that the shortest path from the root point $\mathbf{r}$ to $\mathbf{p}$ consists of the line segment $L(\mathbf{r}, \mathbf{x})$ defining the boundary of the obstacle between $\mathbf{r}$ and $\mathbf{x}$, plus the great circle line $C(\mathbf{x}, \mathbf{p})$ between $\mathbf{x}$ and $\mathbf{p}$. The central angle between $\mathbf{p}$ and any point $\mathbf{x}^{\prime}$ on the line $\theta=\theta_{1}$ is given by

$$
\begin{equation*}
\Delta \sigma=\cos ^{-1}\left(\sin \left(\phi^{\prime}\right) \sin \left(\phi_{2}\right)+\cos \left(\phi^{\prime}\right) \cos \left(\phi_{2}\right) \cos \left(\theta_{2}-\theta_{1}\right)\right) \tag{3}
\end{equation*}
$$

From identity $\frac{\partial}{\partial x} \cos ^{-1}(x)=\frac{-1}{\sqrt{1-x^{2}}}$ we have that as $\phi^{\prime}$ decreases away from $\phi_{1}$ (recall the equator is at $\phi=0$ ) the central angle $\Delta \sigma$ and therefore the great-circle distance changes by $-\frac{\partial}{\partial \phi^{\prime}}$ which has the sign of $\operatorname{sgn}\left[\cos \left(\phi^{\prime}\right) \sin \left(\phi_{2}\right)-\sin \left(\phi^{\prime}\right) \cos \left(\phi_{2}\right) \cos \left(\theta_{2}-\theta_{1}\right)\right]$. Substituting in $\theta_{1}$ from (2) gives $\operatorname{sgn}\left[\tan \left(\phi_{1}\right)-\tan \left(\phi^{\prime}\right)\right]$ and since $\phi^{\prime}<\phi_{1}$ this is always positive

$$
\begin{equation*}
\operatorname{sgn}\left[-\frac{\partial}{\partial \phi^{\prime}} \Delta \sigma\right]=\operatorname{sgn}\left[\tan \left(\phi_{1}\right)-\tan \left(\phi^{\prime}\right)\right]=1 . \tag{4}
\end{equation*}
$$

It follows that the length of this segment always increases with decreasing $\phi^{\prime}$ and we can conclude that not only is $C(\mathbf{x}, \mathbf{p})$ the shortest traversable path between $\mathbf{x}$ and $\mathbf{p}$ but it is also the shortest traversable path between $\mathbf{p}$ and any point $\mathbf{x}^{\prime}$ on $\theta=\theta_{1}$ conditioned on $\phi^{\prime}<\phi_{1}$. Finally, since $R \cdot \cos \phi_{1}$ is the lower bound on the radius for all paths between longitudes $\theta_{0}$ and $\theta_{1}$ in the northern hemisphere it follows that the shortest distance from $r$ to the line defined by $\theta=\theta_{1}$ is $R \cdot \cos \phi_{1}\left(\theta_{1}-\theta_{0}\right)$, which is the path defining the obstacle boundary between those points. Hence the turning point for the optimal path between $\mathbf{p}$ and $\mathbf{r}$ is $\mathbf{x}$, as would have been concluded by string-pulling in the plane under a Gnomonic projection.

For completeness, the practical consequence of Lemma 1 within the grid as defined for Spherical Anya is also provided explicitly as an additional Lemma.

Lemma 2. All turning points on an optimal path in Spherical geometry are corner points or lie within the interval defined by two vertically aligned corner points of an obstacle.

Proof. The statement is equivalent to the claim that no turning point lies within the horizontal interval of a barrier. True by construction; horizontal lines are great-circle lines.


Figure 5
Lemma's 1 and 2 provide the foundation on which to present the final components required to implement the Spherical Anya algorithm, that of a successor, which is again adapted as minimally as possible from that of a successor in Anya, and three categories of projection; flat projections, standard cone projections and, uniquely to Spherical Anya, adjoint cone projections.

Whereas the optimal path from source to target in Anya is a sequence of pairwise visible root points, in Spherical Anya the optimal path is a sequence of turning points which, with the exception of turning points lying on the vertical edge of an obstacle, are also pairwisevisible root points, by Lemma's 1 and 2. In spite of this and the significant adaptations made, the algorithm for calculating this sequence of optimal turning points is analogous to Anya, where successors are constructed at each node, the mechanism by which this occurs underlying the guarantee of completeness, optimality and termination identically in Spherical Anya as in Anya.

## 4. Successors and Projections

Successors in Spherical Anya are separated into two categories, each according to the kind of Anya node associated with it; standard or adjoint.

Definition 5. A standard successor of a search node $(I, r)$ is a search node $\left(I^{\prime}, r^{\prime}\right)$ such that

1. $\left(I^{\prime}, r^{\prime}\right)$ is a standard node.
2. For all points $p^{\prime} \in I^{\prime}$ there exists a point $p \in I$ such that the great-circle segments $C\left(r, p, p^{\prime}\right)$ form a locally optimal route from $r$ to $p^{\prime}$.
3. $r^{\prime}$ is the unique point closest to $p^{\prime}$ shared by all great-circle segments $C\left(r, p, p^{\prime}\right)$.
4. $I^{\prime}$ is maximal given the points above and the definition of a standard node.

This is the analog of a successor as defined in Anya. Due to the existence of adjoint nodes in Spherical geometry, it is also necessary to define an adjoint successor. For all intents and purposes these are to be treated as a standard successor in the implementation of the Spherical Anya algorithm; the fact that no point in the node interval is visible from the root point is relevant only insofar as local optimality of the path from adjoint node root $r$ through $x$ to any point in $I$ is in those cases accorded via Lemma 1 , as is the corresponding addition to the route $g$-value. Hence

$$
\begin{align*}
f-\text { value } & =g-\text { value }+ \text { Anya heuristic } \\
& =\left|C\left(s, r_{1}, \ldots, r_{\text {current }}\right)\right|+\left|C\left(r_{\text {current }}, t\right)\right| \tag{5}
\end{align*}
$$

is throughout given by the sum of the shortest-path line segments formed by the sequence of root points $r_{1}, \ldots, r_{\text {current }}$ within each node from the source $s$, plus the Anya heuristic, being the great-circle distance between the current root and the target $t$ through the associated interval. The shortest-path line segments $C$ are great-circles in all cases except for adjoint nodes, where the path and therefore the corresponding addition to the $g$-value is described by Lemma 1 .

Definition 6. An adjoint successor of a search node $(I, r)=((a, b), r)$ is a search node $\left(I^{\prime}, r^{\prime}\right)=\left(\left(a^{\prime}, b^{\prime}\right), r^{\prime}\right)$ such that

1. $\left(I^{\prime}, r^{\prime}\right)$ is an adjoint node.
2. $I^{\prime}$ is maximal given the definition of an adjoint node.
3. Either $r^{\prime}=a$ or $r^{\prime}=b$.

The mechanism by which successors are generated is via projections, analogously to Anya, these being (i) observable (visible) cone projections, (ii) non-observable (invisible) standard cone projections, (iii) non-observable adjoint cone projections or (iv) flat projections. Some of these projections may themselves be broken into subcategories dependent on context. It is useful to give an overview of these in Anya, as the terminology and concepts are foundational to both algorithms.

### 4.1 Projections in Anya

An observable cone projection is illustrated in Figure 6(a). This is the line-of-sight straight line projection in Euclidean space which extends the visibility cone described by the search node $[a, b], r$ from an initial row to the row above, resulting in successor node $[e, f], r$. Figure $6(\mathrm{~b})$ provides an example of a non-observable cone projection which results in successor $(f, g], b$. Here the Euclidean straight line segment $L(b, f)$ defines an extremal line of the visibility cone for the successor of Figure 6(a) and the interval endpoint $g$ is defined such that the length of the interval $(f, g]$ is maximal, as per the definition of a successor node. The final manifestations of cone projections are those as illustrated in Figures 6 (c) and (d), the former being an intermediate non-observable cone projection and the latter a (nonintermediate) non-observable cone projection.

Flat projections, those extended from flat nodes whereby the root point lies on the same row as the interval, are typically sub-divided into two types, being intermediate flat projections and non-intermediate flat projections, these being referred to simply as flat projections; see Figures 6 (e) and (f) for illustrations. An intermediate flat projection gives rise to the successor search node ( $b, c], a$ as in Figure 6(c). Due to this interval being adjacent to a non-traversable obstacle, it is known in advance that this search node will either be completely traversed or ignored, as any turning point along an optimal route cannot lie within the interval $(b, c)$. Hence one arrives at Figure 6(d), illustrating the converse which is a non-intermediate flat projection. Here, the successor node ( $c, d], a$ may not merely be traversed but instead may be the site of a turning point, leading to successive cone projections. Note that an interval may be a member of multiple successor nodes, as is the case for intervals $(b, c]$ and $(c, d]$.

### 4.2 Projections in Spherical Anya

Projections in Spherical Anya are analogous to those in Anya but for the fact that line-ofsight checking considers the great-circle line passing through point pairs. This differentiation is further exposed by the appearance of adjoint successors in Spherical Anya.

Flat projections remain the same, visually, due to the convention taking rows as lines of longitude. Where the extremal great-circle line segments $C(r, a)$ and $C(r, b)$ of a node $[a, b], r$ do not intersect an obstacle, a great-circle observable cone projection yields the standard successor $[e, f], r$ much as per Anya, save for the visibility cone defining the maximal interval $I=[e, f]$ now being calculated as the intercept of a great circle line through $r$ and $a$ with the corresponding row of the grid, as illustrated in Figure 7(a). Similarly for the (greatcircle) non-observable cone projection $(f, g], b$ of $7(\mathrm{~b})$, defining a standard successor node whereby $f$ is the last point in the cone projection $[e, f], r$ visible to $r$ and the interval $(f, g]$ is calculated to be maximal as per the definition of a successor node. Figures 7(c) and (d) illustrate non-observable cone projections, these being visually identical to Anya. Only the underlying metric for the geodesic between points on the same row has changed. Finally adjoint cone projections, which are non-observable cone projections unique to Spherical Anya, arise when a successor is an adjoint successor, as illustrated in Figure 7(e).

(a) An observable cone projection. In this instance the visibility cone is projected up one line to yield successor node $[e, f], r$.

(c) An intermediate non-observable cone projection. The successor node $(b, c], b$ to $[a, b], r$ is said to be intermediate because it itself has only one successor, namely $(c, d], c$ in Figure 6(d).

(e) An intermediate flat projection. Here the search node $(a, b], a$ has been projected horizontally along the same row, generating successor $(b, c], a$, an intermediate successor as per Figure $6(\mathrm{c})$.

(b) A non-observable cone projection. In this case, $b$ being a successor root node, the cone projection of $[a, b], r$ is found to be $(f, g], b$. No point within $(f, g]$ is visible from $r$.

(d) A non-observable cone projection. In this illustration the non-observable successor $(c, d], b$ is non-intermediate because it may have multiple successors, including successors generated via cone projections.

(f) A non-intermediate flat projection. Here the search node ( $a, b]$, $a$ has been projected horizontally along the same row, generating successor $(c, d], a$. This is a non-intermediate successor as per Figure 6(d).

Figure 6: Anya projections on a Euclidean world map

(a) An observable cone projection

(c) An intermediate non-observable cone projection

(b) A non-observable cone projection

(d) A non-intermediate non-observable cone projection

(e) A non-observable adjoint cone projection. Here the adjoint successor of the search node $[a, b], r$ is the adjoint node $[c, d), a$, whereby all points in [c, d) are invisible from $a$ and the interval endpoint $d$ is determined such that $d$ is visible from $a$.

Figure 7: Spherical Anya projections on a world map in Spherical geometry.

## 5. Spherical Anya Algorithm

The recipe for Spherical Anya can now be presented using the foundational components and principles above. Globally, Spherical Anya is defined as per Anya, the difference being in the successors generated and the underlying geodesic and line-of-sight considerations. In the case of the former we distinguish corner points of the standard form from those which are additionally root points of adjoint nodes by denoting these as adjoint corner points. Adjoint corner points then are bound by the geometric considerations of Lemma 1, whereas (non-adjoint) corner points are not. For easy reference, the Anya algorithm as per (Harabor et al., 2016) is reproduced below.

```
Algorithm 1: Anya and Spherical Anya
Require: Grid, source location s, target location t, initial root point ro
    open }\leftarrow{(I=[s],\mp@subsup{r}{0}{})}\quad\triangleright\mathrm{ root point }\mp@subsup{r}{0}{}\mathrm{ located off the grid
    while open is not empty do
        (I,r)}\leftarrowpop(open
        if }t\inI\mathrm{ then
            return path_to(I)
        end if
        for all (I', r') \in successors ( }I,r\mathrm{ ) do
            if }\neg\mathrm{ should_prune ( }\mp@subsup{I}{}{\prime},\mp@subsup{r}{}{\prime})\mathrm{ then
                open }\leftarrow\mathrm{ open }\cup{(\mp@subsup{I}{}{\prime},\mp@subsup{r}{}{\prime})}\quad\triangleright\mathrm{ Successor pruning
            end if
        end for
    end while
    return null
```

The algorithm for calculating successors in Spherical Anya differs in form from that of Anya due to the existence of adjoint nodes in Spherical geometry (see Algorithms 2 and 3). In addition, the operations generate-cone-successors and generate-flat-successors from (Harabor et al., 2016) must be replaced by their great-circle analogs from Spherical Anya utilising great-circle geodesic and line-of-sight considerations. Furthermore, while the pruning strategies for Anya as described in (Harabor et al., 2016) carry across to Spherical Anya, additional pruning strategies may be applied to adjoint successors (see Algorithm 4).

```
Algorithm 2: Successors in Spherical Anya
function \(\operatorname{SUCCESSORS}(n=(I, r)) \quad \triangleright\) Input is current node
    if \(n\) is the start node \(s\) then
        return generate-start-successors \((I=[s])\)
    end if
    successors \(\leftarrow \emptyset\)
```

```
            if n}\mathrm{ is a flat node then
            p\leftarrow endpoint of I farthest from r }\triangleright\mathrm{ Successor interval starts from p
    successors }\leftarrow\mathrm{ generate-flat-successors( }p,r)\quad\triangleright\mathrm{ Observable successors
    if p is a corner turning point on a locally optimal path beginning at r then
        successors \leftarrow successors \cup generate-cone-successors ( }p,p,r)\quad
    Non-observable successors
    end if
    if p is an adjoint turning point on a locally optimal path beginning at r then
        successors }\leftarrow\mathrm{ successors U generate-adjoint-successors( }p,r
    Non-observable successors
        end if
        else if n it is a cone then
            a\leftarrow left endpoint of I
            b\leftarrow right endpoint of I
            successors }\leftarrow\mathrm{ generate-cone-successors(a,b,r) }\triangleright\mathrm{ Observable successors
            if }a\mathrm{ is a corner turning point on a locally optimal path beginning at r then
                successors}\leftarrow\mathrm{ successors }\cup\mathrm{ generate-flat-successors(a,r) D
Non-observable successors
                successors }\leftarrow\mathrm{ successors }\cup\mathrm{ generate-cone-successors (a,a,r)}
Non-observable successors
    end if
    if }a\mathrm{ is an adjoint turning point on a locally optimal path beginning at r then
        successors }\leftarrow\mathrm{ successors }\cup\mathrm{ generate-adjoint-successors (a,r) D
    Non-observable successors
        end if
        if b is a corner turning point on a locally optimal path beginning at r then
        successors \leftarrow successors \cup generate-flat-successors(b,r) \triangleright
    Non-observable successors
        successors \leftarrow successors \cup generate-cone-successors (b,b,r) \triangleright
    Non-observable successors
        end if
        if b is an adjoint turning point on a locally optimal path beginning at r then
        successors}\leftarrow\mathrm{ successors U generate-adjoint-successors }(b,r)\quad
    Non-observable successors
        end if
        else }\triangleright\mathrm{ If node is neither flat or cone, it is an adjoint
            p\leftarrow endpoint of I closest to r
        successors }\leftarrow\mathrm{ successors U generate-adjoint-successors ( }p,r
Non-observable successors
        successors }\leftarrow\mathrm{ generate-cone-successors (p, p,r) }\triangleright\mathrm{ Non-observable successors
        end if
end function
```


### 5.1 Adjoint Successors

Adjoint successors occur only adjacent to the vertical edge of an obstacle, this by convention being defined as a line of constant latitude. Adjoint nodes in the northern hemisphere exist adjacent to the south-facing edge of an obstacle and symmetrically in the southern hemisphere exist adjacent to the north-facing edge of an obstacle (recall Figure 2). Note also that an adjoint node for the same obstacle edge can share a root point with an adjoint node adjacent to the same edge, the difference being only the row on which each of the respective intervals associated with the adjoint nodes lie. The following algorithm provides the recipe with which to calculate adjoint successors.

Figure 8 provides an illustration of the procedure, whereby all adjoint successors of the search node $[a, b], r$ are calculated. The first adjoint successor of this search node is $[c, d), a$, the interval $[c, d)$ calculated as per the definition of an adjoint node, such that the width of the interval is maximal and hence such that $d$ is visible from the new root point $a$ but all points $p \in[c, d)$ are invisible to it. This adjoint node, having only a single possible successor, namely $[e, f), a$, is an intermediate adjoint node, as is $[e, f), a$ itself. Here, the extremal point $f$ is simply the intersection of the great circle line through $a$ and $d$ with the subsequent row of the grid. Finally, projecting the great-circle lines to the next row yields the last adjoint successor $[g, h), a$ associated with this obstacle edge. Search node $[g, h), a$ is not intermediate, having more than one successor and also does not yield further adjoint node successors.

```
Algorithm 3: Computing a set of adjoint successors
function Generate-Adjoint-Successors(an interval endpoint \(p\), a root point \(r\) )
    successors \(\leftarrow \emptyset\)
    if \(p\) is a corner point then
        \(r^{\prime} \leftarrow p\)
    else
        \(r^{\prime} \leftarrow r\)
    end if
    \(p^{\prime} \leftarrow\) a point from adjacent row, farthest and non-observable one from \(r^{\prime}\)
    \(I_{\text {max }} \leftarrow\) maximum closed interval beginning at \(p^{\prime}\) and non-observable from \(r^{\prime}\)
    for all \(I \in\left\{\right.\) split \(I_{\max }\) at each corner point \(\}\) do
        \(n^{\prime} \leftarrow\) a new search node with interval \(I\) and root point \(r^{\prime}\)
        successors \(\leftarrow\) successors \(\cup I\)
    end for
    return successors
end function
```

Algorithm 4: Intermediate node pruning


Figure 8: A series of intermediate adjoint nodes which are calculated in the process of generating the final non-intermediate adjoint node $[g, h), a$.

```
function Is-Intermediate( }n=(I,r)
    if n is a flat node then
        p\leftarrow endpoint of I furthest from r
        if p is a corner turning point for a locally optimal path with prefix }\langler,p\rangle\mathrm{ then
            return false }\trianglerightn\mathrm{ has at least one non-observable successor; it cannot be
intermediate
        end if
    else if n is an adjoint node then
        p
        if p
            return false \trianglerightn has at least one non-observable successor; it cannot be
intermediate
        end if
        else
            if I has a closed endpoint that is also a corner point then
                return false }\trianglerightn\mathrm{ has at least one non-observable successor; it cannot be
intermediate
    end if
    I'}\leftarrow\mathrm{ interval after projecting r through }
    if I' contains any corner points then
```

```
18:
return false \(\quad n\) has at least one non-observable successor; it cannot be
intermediate
            end if
        end if
        return true
end function
```


### 5.2 Spherical Anya Algorithm: Examples

We provide examples which illustrate how Spherical Anya works in practice by illustrating all successors for a given root node as would be necessary in an implementation of the Spherical Anya algorithm, search nodes and their corresponding successors comprising the fundamental units of the search algorithm.

### 5.2.1 Successors of a Flat Search Node

In the following scenario the Spherical Anya algorithm is required to generate all successors of the flat search node $(r, a], r \equiv(0,2], r$, of Figure 9 . As the root point of the node is on the same row as the interval, the first projections are flat projections, generating flat successors until the wall of Figure 9(d) interrupts progress. In this instance the wall is created by the existence of two diagonally adjacent non-traversable cells centred on the point $e=(7,2)$ of Figure 9(d). Not all of these successors survive however, those not containing the target point and being intermediate by virtue of their adjacency to a non-traversable cell, as ( $b, c], r$ is in Figure 9(b) and (d,e],r is in Figure 9(d), being discarded. These cannot precede a turning point on any optimal route. The search node ( $d, e], r$ of Figure 9(d) is in addition a cul-de-sac node. It has no successors, and the target does not lie in the interval ( $d, e$ ].

All flat successors being thus generated, the Spherical Anya algorithm generates all flat-to-adjoint projections and flat-to-cone projections. These occur at the farthest corner of each non-traversable interval from the root $r$, along the same row and prior to the edge defined by the wall. In this example, flat-to-adjoint and flat-to-cone projections will be generated by Spherical Anya at each of the new root points $a=(2,2)$ and $c=(5,2)$, as illustrated in figures 9(e)-(h).

### 5.2.2 Successors of a Cone Search Node

In the next example, the search node $[a, b], r$ is now a cone search node, as illustrated in Figure 10, as the root point $r$ no longer lies on the same row as the interval $[a, b]$. The first successor of this is calculated as the great-circle cone projection to the next row of the map, yielding successor $[c, d], r$ of Figure $10(a)$ with the same root point as its parent. Subsequent to the (non-adjoint) cone projection yielding search node $[c, d], r$ all invisible cone projections must be generated, these being invisible from root point $r$. This is achieved by forming the search node $(d, e], b$ of Figure $10(\mathrm{~b})$ by searching right from $d$, the last point on row 3 visible to $r$.

Once there are no further options for expanding a great-circle visibility cone to the third row of the map the search continues horizontally, generating flat, intermediate, flat-

(a) An observable flat projection. Not being adjacent to a non-traversable cell, this successor $[2,4], r$ is added to the list of successors of $[0,2], r$.

(c) An observable flat projection yielding successor $[5,6], r$.

(e) An adjacent flat-to-cone projection. The root point is $r^{\prime}$, invisible to all points in the node interval $I$, because a great-circle line intersects the neighbouring obstacle twice.

(g) An adjacent flat-to-cone projection. The root point is $r^{\prime}$, invisible to all points in the node interval $I$, because a great-circle line intersects the neighbouring obstacle twice.

(b) An intermediate flat projection. Being adjacent to a non-traversable cell this successor node is discarded.

(d) An intermediate flat projection. This interval is adjacent to a cul-de-sac wall formed by two diagonally adjacent obstacles, and hence is discarded.

(f) A non-adjoint standard cone projection provides the final successor with root point $r^{\prime}$ and interval $J$. All points in $J$ are visible with respect to $r^{\prime}$.

(h) A non-adjoint standard cone projection provides the final successor with root point $r^{\prime}$ and interval $J$. All points in $J$ are visible with respect to $r^{\prime}$.

Figure 9: All successors of search node $(r, a], r$ are illustrated to provide an example of successor generation in Spherical Anya.
to-adjoint and flat-to-cone projections. This process continues along row 2 until a wall is reached, or as is the case in this scenario, the edge of the map. The first flat projection is $(b, f], b$, this being the case because the new root point $b$ is selected due to being the extremal edge of the parent interval. Node $(b, f], b$ however, being adjacent to an obstacle, is intermediate and not added to the list of successors. The subsequent successor $(f, g], b$, while not being adjacent to a non-traversable cell in this instance, is also not added to the list of successors for the parent node $[a, b], r$ because $(f, g], b$ is adjacent to the edge of the map and therefore is a successor without successors. Given the target doesn't lie in the interval $(f, g]$ it is a cul-de-sac, and hence dropped as well. The only remaining successors in this scenario are therefore the flat-to-adjoint and flat-to-cone successors of Figures 10(e) and (f).


Figure 10: A series of flat, flat-to-adjoint and flat-to-cone projections define the successors to the illustrated search node $[a, b], r$.

| Benchmark | \#Maps | \#Instances | Spherical Anya vs. Anya ET (All Routes) |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  | \% less | Min | Q1 | Median | Mean | Q3 | Max | StDev |
| Starcraft | 75 | 18,068 | 35.27 | 0.003 | 0.806 | 1.188 | 2.048 | 2.018 | 60.019 | 3.451 |
| Dragon Age 2 | 67 | 14,369 | 31.58 | 0.006 | 0.857 | 1.000 | 1.368 | 1.525 | 56.366 | 1.359 |
| Baldurs Gate II, scaled to $512 \times 512$ | 75 | 17,580 | 31.18 | 0.005 | 0.861 | 1.000 | 1.561 | 1.690 | 53.262 | 1.798 |
| Warcraft III, scaled to $512 \times 512$ | 36 | 7,992 | 38.63 | 0.015 | 0.833 | 1.000 | 2.881 | 1.367 | 60.421 | 6.346 |
| City/street Maps | 90 | 21,515 | 33.53 | 0.014 | 0.804 | 1.261 | 3.424 | 3.103 | 60.442 | 6.160 |
| Random $10 \%$, $256 \times 512$ | 75 | 18,241 | 53.37 | 0.000 | 0.476 | 0.911 | 1.510 | 1.686 | 13.807 | 1.906 |
| Random $20 \%, 256 \times 512$ | 75 | 18,305 | 46.48 | 0.000 | 0.593 | 1.052 | 1.587 | 1.876 | 13.782 | 1.724 |
| Random 30\%, $256 \times 512$ | 75 | 18,175 | 41.67 | 0.001 | 0.725 | 1.112 | 1.419 | 1.757 | 13.724 | 1.209 |
| Random 40\%, $256 \times 512$ | 75 | 16,091 | 31.73 | 0.000 | 0.944 | 1.100 | 1.207 | 1.359 | 13.582 | 0.694 |
| NOAA bathymetry, $10800 \times 21600$ | 1 | 22,060 | 74.20 | 0.003 | 0.796 | 0.905 | 1.630 | 1.000 | 90.596 | 5.085 |

Table 1: Execution time (ET) on a series of standard game maps, random maps and NOAA bathymetry is calculated from randomly sampled source and target points and alternatively considered to be defined over Euclidean and Spherical geometry. The column "\% less" indicates the percentage of routes whereupon Spherical Anya ET is less than Anya ET.

## 6. Results

Results are reported on execution time (ET) and number of tiles crossed (CT), a tile being a single unit of the world map grid. Statistics are calculated for 4 game maps plus city/street maps and four random maps, in addition to the Earth maps determined by the NOAA bathymetric dataset at several resolutions. Comparisons and benchmarks to existing Euclidean approaches being already reported in (Harabor et al., 2016), the following results compare exclusively the difference between Anya and Spherical Anya. Both algorithms are implemented in Python and executed on an AWS EC2 Dual Core Intel Xeon Platinum 8000 "R5.Large" Instance with 16GB of RAM. The definition of the benchmark is designed to measure the relative expense in time and route length of achieving optimality in Spherical geometry with the application of Spherical Anya, compared to applying Anya with the assumption that the underlying geometry is Euclidean. We highlight that such a comparison is dependent on programming language and employed data structures, making all benchmark statistics herein indicative in Python only. We do not provide a comparison across programming languages, for which these figures would likely be different. For any world map defined as a list of coordinate pairs $\left[\left(x_{1}, y_{1}\right), \ldots .,\left(x_{m}, y_{n}\right)\right]$ for an $n \times m$ world, each item is assigned a value of zero for a non-traversable cell and one for a traversable cell. For the application of Anya, source and target points are selected uniformly from the set of traversable points, and Anya is applied on the world map directly, considering each cell to be a uniform square of unit length. The resulting route is mapped to Spherical coordinates with the linear mapping, in degrees, being

$$
\begin{align*}
x_{i}^{l a t} & =90\left(\frac{2 x_{i}}{m}-1\right) \\
y_{j}^{\text {lon }} & =180\left(\frac{2 y_{j}}{n}-1\right) . \tag{6}
\end{align*}
$$

Results are shown for two distinct recipes for connecting the root points returned from Anya to form a route in Spherical geometry. Recipe 1 connects root points directly with greatcircle lines. A consequence of this is that Recipe 1 achieves routes whose total length is closer

| Benchmark | \#Maps | \#Instances | Spherical Anya vs. Anya Route Length (Recipe 1) |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  | \% less | Min | Q1 | Median | Mean | Q3 | Max | StDev |
| Starcraft | 75 | 18,068 | 56.98 | 0.132 | 0.990 | 1.000 | 0.977 | 1.000 | 1.509 | 0.069 |
| Dragon Age 2 | 67 | 14,369 | 36.27 | 0.162 | 1.000 | 1.000 | 0.995 | 1.000 | 2.047 | 0.047 |
| Baldurs Gate II, scaled to $512 \times 512$ | 75 | 17,580 | 36.71 | 0.191 | 0.999 | 1.000 | 0.997 | 1.000 | 5.528 | 0.082 |
| Warcraft III, scaled to $512 \times 512$ | 36 | 7,992 | 29.24 | 0.279 | 0.998 | 1.000 | 1.008 | 1.000 | 5.201 | 0.143 |
| City/street Maps | 90 | 21,515 | 60.54 | 0.180 | 0.910 | 0.996 | 0.947 | 1.000 | 5.983 | 0.137 |
| Random $10 \%$, $256 \times 512$ | 75 | 18,241 | 96.75 | 0.049 | 0.738 | 0.931 | 0.834 | 0.988 | 1.003 | 0.206 |
| Random 20\%, $256 \times 512$ | 75 | 18,305 | 95.96 | 0.046 | 0.748 | 0.932 | 0.836 | 0.988 | 1.002 | 0.205 |
| Random 30\%, $256 \times 512$ | 75 | 18,175 | 93.84 | 0.053 | 0.792 | 0.949 | 0.851 | 0.994 | 1.001 | 0.204 |
| Random $40 \%, 256 \times 512$ | 75 | 16,091 | 84.41 | 0.024 | 0.965 | 0.998 | 0.924 | 1.000 | 1.001 | 0.172 |
| NOAA bathymetry, $10800 \times 21600$ | 1 | 22,060 | 37.44 | 0.320 | 1.000 | 1.000 | 0.997 | 1.000 | 1.103 | 0.017 |

Table 2: Route Length (RL) on a series of standard game maps, random maps and NOAA bathymetry is calculated from randomly sampled source and target points and alternatively considered to be defined over Euclidean and Spherical geometry. The table includes both illegal and legal Anya routes generated by interpolating between route points with a great-circle (Recipe 1). The column "\% less" indicates the percentage of routes whereupon Spherical Anya RL is less than Anya RL. The ratio being greater than 1 occurs when nonlegal routes are permitted, i.e. when Anya generates a route which in Spherical geometry intersects at least one non-traversable cell. All routes generated by Spherical Anya are legal.

| Benchmark | \#Maps | \# Instances | Spherical Anya vs. Anya Route Length (Recipe 2) |  |  |  |  |  |  |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
|  |  |  | $\%$ less | Min | Q1 | Median | Mean | Q3 | Max | StDev |
| Starcraft | 75 | 18,068 | 91.22 | 0.131 | 0.969 | 0.993 | 0.965 | 0.999 | 1.000 | 0.072 |
| Dragon Age 2 | 67 | 14,369 | 86.90 | 0.162 | 0.996 | 0.999 | 0.990 | 1.000 | 1.000 | 0.043 |
| Baldurs Gate II, scaled to $512 \times 512$ | 75 | 17,580 | 75.04 | 0.191 | 0.989 | 0.999 | 0.984 | 1.000 | 1.000 | 0.045 |
| Warcraft III, scaled to $512 \times 512$ | 36 | 7,992 | 53.32 | 0.279 | 0.938 | 1.000 | 0.954 | 1.000 | 1.000 | 0.079 |
| City/street Maps | 90 | 21,515 | 85.94 | 0.176 | 0.870 | 0.971 | 0.917 | 0.999 | 1.000 | 0.111 |
| Random $10 \%, 256 \times 512$ | 75 | 18,241 | 99.27 | 0.049 | 0.737 | 0.930 | 0.833 | 0.988 | 1.000 | 0.206 |
| Random $20 \%, 256 \times 512$ | 75 | 18,305 | 99.62 | 0.046 | 0.748 | 0.931 | 0.836 | 0.988 | 1.000 | 0.205 |
| Random $30 \%, 256 \times 512$ | 75 | 18,175 | 99.86 | 0.053 | 0.792 | 0.948 | 0.851 | 0.994 | 1.000 | 0.204 |
| Random $40 \%, 256 \times 512$ | 75 | 16,091 | 99.93 | 0.024 | 0.965 | 0.998 | 0.924 | 1.000 | 1.000 | 0.172 |
| NOAA bathymetry, $10800 \times 21600$ | 1 | 22,060 | 49.79 | 0.320 | 0.999 | 1.000 | 0.993 | 1.000 | 1.000 | 0.024 |

Table 3: Distribution of Route Length Ratios (including illegal routes) where Anya routes are generated by linear enrichment every 1 arc-second between root point pairs for which a great-circle connection would yield an intersection in Spherical geometry (Recipe 2). Compared to the mechanism of Table 2 this simultaneously reduces the number of illegal routes generated by Anya and increases the corresponding route lengths further from those of Recipe 1. The column "\% less" indicates the percentage of routes whereupon Spherical Anya route length is less than Anya route length. The opposite occurs for illegal Anya routes.
to that of the optimal route returned by Spherical Anya, but with a relatively high number of routes intersecting non-traversable cells. We refer to routes as legal whereupon this does not occur. All routes generated by Spherical Anya in Spherical geometry are both legal and optimal whereas this is not the case for Anya - Anya not being designed for application in Spherical geometry and therefore guaranteeing neither legality nor optimality of routes. Recipe 2 enriches the route segments returned by Recipe 1 which intersect non-traversable cells with additional synthetic root points every arc-second along a straight (Euclidean) line between the root points where this occurs. Finally a great-circle line is assumed between

| Benchmark | \#Instances | Spherical Anya vs. Anya Route Length (Exclusive, Recipe 1) |  |  |  |  |  |  |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
|  |  |  | \% less | Min | Q1 | Median | Mean | Q3 | Max |
| StDev |  |  |  |  |  |  |  |  |
| Starcraft | 12,594 | 81.75 | 0.132 | 0.970 | 0.997 | 0.964 | 1.000 | 1.000 | 0.077 |
| Dragon Age 2 | 9,511 | 54.80 | 0.162 | 0.999 | 1.000 | 0.990 | 1.000 | 1.000 | 0.049 |
| Baldurs Gate II, scaled to $512 \times 512$ | 12,454 | 51.81 | 0.191 | 0.996 | 1.000 | 0.987 | 1.000 | 1.000 | 0.045 |
| Warcraft III, scaled to $512 \times 512$ | 6,295 | 37.12 | 0.279 | 0.989 | 1.000 | 0.978 | 1.000 | 1.000 | 0.052 |
| City/street Maps | 17,094 | 76.20 | 0.180 | 0.875 | 0.974 | 0.922 | 1.000 | 1.000 | 0.107 |
| Random $10 \%, 256 \times 512$ | 17,997 | 98.07 | 0.049 | 0.734 | 0.929 | 0.832 | 0.987 | 1.000 | 0.206 |
| Random $20 \%, 256 \times 512$ | 17,860 | 98.35 | 0.046 | 0.740 | 0.927 | 0.832 | 0.986 | 1.000 | 0.206 |
| Random $30 \%, 256 \times 512$ | 17,289 | 98.65 | 0.053 | 0.777 | 0.941 | 0.844 | 0.990 | 1.000 | 0.206 |
| Random $40 \%, 256 \times 512$ | 13,716 | 99.02 | 0.024 | 0.945 | 0.995 | 0.911 | 1.000 | 1.000 | 0.183 |
| NOAA bathymetry, $10800 \times 21600$ | 19,924 | 41.46 | 0.320 | 1.000 | 1.000 | 0.996 | 1.000 | 1.000 | 0.017 |

Table 4: Distribution of Route Length Ratios excluding illegal routes wherein the Anya route is shorter than the Spherical Anya route. This occurs if and only if Anya returns a route which intersects a non-traversable cell at least once.
root-point pairs in the enriched set. Recipe 2 produces fewer illegal routes than Recipe 1, albeit at the expense of generating legal routes with distance greater than those produced by Recipe 1 .

For Spherical Anya the world map is first transformed into Spherical coordinates and the shortest-distance route is subsequently calculated. Under such circumstances it is guaranteed that Spherical Anya will find the optimal legal route whereas Anya may not, the relevant question here being at what cost? What is the cost in execution time (ET) and route length (RL) of achieving optimality in Spherical geometry, compared with the simplification of assuming a Euclidean geodesic on a "flat earth".

Results are provided in Tables 1, 2, 3 and 4. The standard benchmark game maps which have been selected are Starcraft, Dragon Age 2, Baldurs Gate II and Warcraft III (Sturtevant, 2012), with source and target uniformly selected with replacement an \#Instances number of times over a \#Maps number of maps for each game. These game maps have features which are interesting to study in their own right and we do not dwell on them here. For example Real Time Strategy games such as Starcraft contain large, obstacle free, traversable areas with choke points (see Figure 11) whereas Role-playing games such as Baldurs Gate II are maze-like. Figure 12 illustrates the ratio of Spherical Anya ET over Anya ET for these game maps and City/street maps. Also presented are results from random maps (see Figure 13). Note that these maps are originally 2D maps, repurposed here for illustrating the statistical relationship between Anya and Spherical Anya given (i) different recipes for creating routes in Spherical geometry from the root points yielded by Anya (ii) different underlying measures for line-of-sight checking and finally, but significantly (iii) differing world maps. Beyond this statistical comparison it is the sea-voyages based on the NOAA bathymetric relief map of Earth (NOAA, 2009) which additionally bear practical utility in real world applications. Whereas the bathymetric map is assumed to wrap around at the edges in the standard way - the line of longitude $180^{\circ}$ coinciding with the line $-180^{\circ}$ - all other benchmark maps are assumed to have solid borders. This is achieved by padding their borders with a series of consecutive non-traversable points.

These simulations show that Spherical Anya is mostly similar in execution time to Anya, but on average it is slower (see Table 1). Moreover the median ET of Spherical Anya be-
ing faster or slower than Anya is strongly dependent on the features of each world map. The reward is only occasionally speed but always legality and optimality, Spherical Anya producing routes on average over $7.4 \%$ shorter than routes generated by Anya across all maps using Recipe 1 (see Table 4). The statistics for Random maps at $10 \%, 20 \%, 30 \%$, $40 \%$ clearly show the importance of the features of the world map as well as the number of tiles crossed (see Figure 13), the percentage here defining the Bernoulli random variable determining the existence of a non-traversable coordinate in the world map. Random maps produce route length ratios smaller than those of other maps and suffer a high RL ratio standard deviation compared to other world maps. More importantly for practical applications such as the navigation of seafaring vessels, median ET for Spherical Anya is slightly less than Anya and average ET is higher for almost all world map resolutions and route lengths (see Figures 14 and 15). Never generating routes which intersect land, as Anya does in Spherical geometry, Spherical Anya always returns the optimal route. Refer to Appendix B, Figures 18 and 19 for illustrations of instances of ocean voyages. In summary, Spherical Anya being always optimal and producing legal routes in Spherical geometry, is faster on most routes compared with Anya for sea-faring routes, but more often it is slower.

## 7. Conclusion

It has been shown that a fast, efficient and optimal pathfinding algorithm for Spherical geometry exists in the form of Spherical Anya, a new online pathfinding algorithm which modifies and extends the principles of the Anya algorithm in Euclidean geometry. Modifications include the obvious change from the Euclidean geodesic to the great-circle geodesic of Spherical geometry. Further changes are a consequence of this, for instance one can no longer rely upon all turning points on an optimal route being corner points, and the existence of adjoint search nodes and successors which do not exist in Anya. The Spherical Anya algorithm for optimal pathfinding in Spherical geometry is described and illustrated with examples, and the primary mathematical consequence of the transition from Euclidean geometry to Spherical geometry is stated and proven. All remaining terminology, procedures and fundamental principles remaining the same as Anya ensures that Spherical Anya adopts the completeness, optimality and termination guarantees of its progenitor.

Finally, examples are shown applying Anya and Spherical Anya to several game maps, random maps and a map of Earth formed of bathymetric data from NOAA. In each case the map is assumed to be a representation of a world in Spherical geometry, and the application of Anya therefore a first order approximation of the optimal route, whereby a "flat earth" approximation is made for all projections, but the resulting series of root points are mapped to Spherical geometry, route length being calculated as the sum of great-circle arcs between them (Recipe 1) or a sequence of enriched root point pairs joined by great-circle lines (Recipe 2). These simulations illustrate the following primary features and advantages of Spherical Anya. Firstly what is known mathematically in advance; that the result from Spherical Anya is always optimal in Spherical geometry and that Anya is not (and is not intended to be). Secondly, that applying Anya naively to a spherical world map to generate a sequence of root point pairs joined by great-circle lines yields a large proportion of illegal routes which at times cross non-traversable cells. The quantity of illegal routes is a function of the recipe used to convert root points in Euclidean geometry to Spherical geometry, however it is also

(a) A Starcraft map instance, with traversable regions coloured yellow and non-traversable regions coloured purple. This can be considered either a flat map, as was originally intended, or a map in Spherical geometry via the mapping of equation (6).

(c) A third example of a Starcraft map instance. The ratio of non-traversable cells compared to traversable cells is $42.11 \%$. The average length of a route for this map is $10,920.2 \mathrm{~km}$ in Euclidean geometry and $10,829.7 \mathrm{~km}$ in Spherical geometry. To compare, the median distance over all Starcraft map instances in Euclidean and Spherical geometry is $13,822.4 \mathrm{~km}$ and $13,454.9 \mathrm{~km}$ respectively.

(b) A second example of a Starcraft map instance. Here the ratio of non-traversable regions compared to traversable regions is $38.37 \%$.

(d) A final example of a Starcraft map instance. The ratio of non-traversable cells compared to traversable cells is $1.38 \%$. The average length of a route for this map is $10,606.3 \mathrm{~km}$ in Euclidean geometry and $9,577.0 \mathrm{~km}$ in Spherical geometry. The median ratio over all Starcraft map instances is $25.77 \%$.

Figure 11: Starcraft Real Time Strategy game map instance examples.


Figure 12: Spherical Anya execution time (ET) over Anya ET against Number of Tiles Crossed (CT) for Game and City/street maps.


Figure 13: Spherical Anya execution time (ET) over Anya ET against Number of Tiles Crossed (CT) for Random Maps.
important to note that as of yet there is no established method which can guarantee all routes produced by Anya are legal in Spherical geometry. And lastly, that Spherical Anya is more often faster than Anya for sea-voyages but in general and on average it is slower, the relative statistical properties depending not only on the algorithm but on the features of each world map.

## Appendix A. Gnomonic Projection

A Gnomonic projection is a mapping of a coordinate system wherein all great-circle lines are represented by a straight line and vice versa (see Figure 16). Used widely in Astronomy, Seismology, Signals Intelligence, air traffic and maritime navigation, a Gnomonic projection cannot be used to enable an error-free application of Anya directly. The reason is understood with the following considerations.


NOAA Grid Resolution
Figure 14: Spherical Anya execution time (ET) over Anya ET for Earth NOAA bathymetric maps taken at several resolutions. This indicates that for most routes one can calculate optimal voyages for seafaring vessels faster than with a "flat earth" approximation, albeit the execution time is on average slower.


Figure 15: Spherical Anya execution time (ET) over Anya ET for buckets of route lengths at intervals of 500 km using the NOAA Earth bathymetric map at resolution $10,800 \times 21,600$.

The conversion of world map points defined in spherical coordinates with a Gnomonic projection centred on the equator $(0,0)$ is given by

$$
\begin{align*}
& x=\tan (\theta) \\
& y=\frac{\tan (\phi)}{\cos (\theta)} \tag{7}
\end{align*}
$$



Figure 16: On a Gnomonic projection of a spherical world map all straight lines describe a great circle. This would be a suitable reference frame for pathfinding on a sphere, but for the fact errors in estimating obstacle boundaries cannot be removed in this case.
where $\theta$ is the azimuthal angle and $\phi$ is the polar angle. Hence latitudinal and longitudinal values from each hemisphere will map to the same $x$ and $y$ values, necessitating that routes in each hemisphere are treated with distinct coordinate systems. The dependence of $y$ upon both $\phi$ and $\theta$ implies that points on a line in Spherical geometry are not uniformly divided and an expanded basis set is required to define every $y$-axis with respect to a given $x$ (see Figure 17). In practice, this means that any visibility graph defined on the Gnomonic world map requires not an $M \times N$ as is the case in Euclidean and Spherical geometry, but instead an $M \times N^{2}$ matrix - potentially in excess of available memory for a large number for real-world applications. More problematic is that implementing an error-free application of Anya using a Gnomonic projection is mathematically impossible as can be shown using the two non-traversable cells illustrated in Figure 17. The square corners of obstacle $A$, transformed by equation 7, form obstacle $A^{\prime}$ in Gnomonic space and analogously $B$ will transform to $B^{\prime}$. Lines of constant longitude in Spherical geometry will remain straight horizontal lines under the Gnomonic projection, whereas lines of constant latitude will be curved. The corresponding edges of $A$ compared to $B$, equal in Spherical geometry in units of $\phi$ and $\theta$, will no longer be equal under the Gnomonic projection. The right and left edges of each obstacle, being defined by a line of constant latitude, form curved lines under the Gnomonic projection, therefore intersecting the lines of the grid regardless of the world map resolution. In other words, there will always be an error in the


Figure 17: Two non-traversable cells or obstacles $A$ and $B$ in Spherical geometry (left). Under a Gnomonic projection, lines of constant latitude will intersect with the grid ensuring that any grid, regardless of the interval length, will always result in an error in the approximated boundary of a transformed obstacle and this error, being a function of $\theta$ and $\phi$, will be non-uniform.
approximated boundary of the transformed obstacles and this error, being a function of $\theta$ and $\phi$, will be non-uniform. We should note that this statement applies exclusively to the error in representing mathematically ideal (perfectly accurate) obstacles using a Gnomonic projection and that it is specifically this which prohibits a guarantee of optimality being associated with the Gnomonic approach. However in the absence of ideal obstacles, such as representations of real-world objects which are prone to measurement error, both the Gnomonic approach and the approach of Spherical Anya as outlined in this paper may suffer sub-optimality.

## Appendix B. Examples of Routes in NOAA Bathymetric Dataset

Optimal pathfinding for ocean voyages is an important practical problem for the maritime industry and global trade, wherein commercial vessels must choose the most expedient routes in order to ensure fast and efficient delivery of cargo.

To measure the performance of Spherical Anya on a real-world problem, more than 20,000 route instances were calculated using the NOAA bathymetric dataset at resolution $10,800 \times 21,600$. Comparing against Anya the statistics are as per Tables 1, 2, 3 and 4 . Table 2 includes all routes, including non-legal routes and $47.2 \%$ of these correspond to cases where the length of the Anya route (as per Recipe 1) is shorter than the length of the corresponding Spherical Anya route, which occurs due to the geometric feature wherein a straight line in Euclidean geometry which doesn't intersect any barriers, nonetheless crosses them in Spherical geometry. Figure 18 illustrates an example of this for the shortest seapassage between New York, USA and Gothenburg, Sweden where a great-circle connecting Anya root points intersects land in Spherical geometry. An example of the opposite case, illustrated in Figure 19, is of Spherical Anya and Anya routes connecting Bahia Blanca, Argentina and Colombo, Sri Lanka, where neither of the routes intersect land. When


Figure 18: Spherical Anya route (solid purple line) and Anya route (dashed red line) from New York, USA to Gothenburg, Sweden. The sequence of great-circles connecting root points of the Anya route are shown with a green solid line. Some of these great-circles intersect land. The length of the Spherical Anya route is $6,222.888 \mathrm{~km}$, while the total length of all green great-circles is $6,132.063 \mathrm{~km}$. The Spherical Anya route is calculated 4.65 times faster than the Anya route.


Figure 19: Spherical Anya route (solid purple line) and Anya route (dashed red line) from Bahia Blanca, Argentina to Colombo, Sri Lanka. The sequence of great-circles connecting root points of the Anya route are shown with a green solid line. None of these great-circles intersect land. The length of the Spherical Anya route is $14,715.343 \mathrm{~km}$, while total length of all green great-circles is $14,918.216 \mathrm{~km}$. The Spherical Anya route is calculated 311 times faster than the Anya route.
this happens, Spherical Anya always returns the shortest route or at best, the two routes coincide.

## References

Bency, M. J., Qureshi, A. H., \& Yip, M. C. (2019). Neural path planning: Fixed time, near-optimal path generation via oracle imitation. arXiv:1904.11102 [cs.RO].
Botea, A., Muller, M., \& Schaeffer, J. (2004). Near optimal hierarchical path-finding. J. Game Dev, 1(1), 7-28.
Ferguson, D., \& Stentz, A. (2005). Field D*: An interpolation-based path planner and replanner. ISRR, 239-253.
Harabor, D., \& Grastien, A. (2013). An optimal any-angle pathfinding algorithm. In Proceedings of ICAPS-23.
Harabor, D., Grastien, A., Öz, D., \& Aksakalli, V. (2016). Optimal any-angle pathfinding in practice. Journal of Artificial Intelligence Research, 56(89).
Hart, P. E., Nilsson, M. J., \& Raphael, B. (1968). A formal basis for the heuristic determination of minimum cost paths.. pp. 100-107.
Hershberger, J., \& Suri, S. (1999). An optimal algorithm for euclidean shortest paths in the plane. SIAM J. Comput., 28(6), 2215-2256.
Liu, Y.-H., \& Arimoto, S. (1992). Path planning using a tangent graph for mobile robots among polygonal and curved obstacles. Int. J. Robot Research, 11, 376-382.
Lozano-Pérez, T., \& Wesley, M. A. (1979). An algorithm for planning collision-free paths among polyhedral obstacles. Commun. ACM, 22(10), 560-570.
Mitchell, J. S. B., Mount, D. M., \& Papadimitriou, C. H. (1987). The discrete geodesic problem. SIAM J. Comput., 16(4), 647-668.
Nash, A., Daniel, K., Koenig, S., \& Felner, A. (2007). Theta*: Any-angle path planning on grids. AAAI, 1177-1183.
NOAA (2009). ETOPO1 earth global relief model. https://www.ngdc.noaa.gov/mgg/ global/.
Panov, A. I., Yakovlev, K. S., \& Suvorov, R. (2018). Grid path planning with deep reinforcement learning: Preliminary results. In 8th Annual International Conference on Biologically Inspired Cognitive Architectures, BICA, Vol. 123, pp. 347-353.
Sartoretti, G., Kerr, J., Shi, Y., Wagner, G., Kumar, T. K. S., Koenig, S., \& Choset, H. (2019). Primal: Pathfinding via reinforcement and imitation multi-agent learning. IEEE Robotics and Automation Letters, 4(3), 2378-2385.
Sišlák, D., Volf, P., \& Pěchouček, M. (2009). Accelerated A* trajectory planning: Gridbased path planning comparison. In Proceedings of 4 'th Workshop on Planning and Plan Execution for Real-World Systems.
Soboleva, N., \& Yakovlev, K. (2019). KI 2019: Advances in Artificial Intelligence., Vol. 11793. Springer.

Sturtevant, N. (2012). Benchmarks for grid-based pathfinding. Transactions on Computational Intelligence and AI in Games, 4(2), 144-148.

Uras, T., \& Koenig, S. (2015). An empirical comparison of any-angle path-planning algorithms. In Proceedings of the Eighth Annual Symposium on Combinatorial Search.

Wang, J., Wu, N., Zhao, W., Peng, F., \& Lin, X. (2019). Empowering A* search algorithms with neural networks for personalized route recommendation.. pp. 539-547.

Yap, P., Burch, N., Holte, R. C., \& Schaeffer, J. (2011). Block A*: Database-driven search with applications in any-angle path-planning. AAAI.

