



JIGSAW-GEO (1.0): locally orthogonal staggered unstructured grid generation for general circulation modelling on the sphere

Darren Engwirda^{1,2}

¹Department of Earth, Atmospheric and Planetary Sciences, Room 54-1517, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139-4307, USA

²NASA Goddard Institute for Space Studies, 2880 Broadway, New York, NY 10025, USA

Correspondence to: Darren Engwirda (engwirda@mit.edu)

Received: 6 December 2016 – Discussion started: 22 December 2016

Revised: 24 April 2017 – Accepted: 25 April 2017 – Published: 6 June 2017

Abstract. An algorithm for the generation of non-uniform, locally orthogonal staggered unstructured spheroidal grids is described. This technique is designed to generate very high-quality staggered Voronoi–Delaunay meshes appropriate for general circulation modelling on the sphere, including applications to atmospheric simulation, ocean-modelling and numerical weather prediction. Using a recently developed Frontal-Delaunay refinement technique, a method for the construction of high-quality unstructured spheroidal Delaunay triangulations is introduced. A locally orthogonal polygonal grid, derived from the associated Voronoi diagram, is computed as the staggered dual. It is shown that use of the Frontal-Delaunay refinement technique allows for the generation of very high-quality unstructured triangulations, satisfying a priori bounds on element size and shape. Grid quality is further improved through the application of hill-climbing-type optimisation techniques. Overall, the algorithm is shown to produce grids with very high element quality and smooth grading characteristics, while imposing relatively low computational expense. A selection of uniform and non-uniform spheroidal grids appropriate for high-resolution, multi-scale general circulation modelling are presented. These grids are shown to satisfy the geometric constraints associated with contemporary unstructured C-grid-type finite-volume models, including the Model for Prediction Across Scales (MPAS-O). The use of user-defined mesh-spacing functions to generate smoothly graded, non-uniform grids for multi-resolution-type studies is discussed in detail.

1 Introduction

The development of atmospheric and oceanic general circulation models based on unstructured numerical discretisation schemes is an emerging area of research. This trend necessitates the development of unstructured grid-generation algorithms designed to produce very high-resolution, guaranteed-quality unstructured triangular and polygonal meshes that satisfy non-uniform mesh-spacing distributions and embedded geometrical constraints. This study investigates the applicability of a recently developed surface meshing algorithm (Engwirda and Ivers, 2016; Engwirda, 2016) based on restricted Frontal-Delaunay refinement and hill-climbing-type optimisation for this task.

1.1 Semi-structured grids

While simple structured grid types for the sphere can be obtained by assembling a uniform discretisation in spherical coordinates, the resulting *latitude–longitude* (*lat–long*) grid is often inappropriate for numerical simulation, due to the presence of strong *grid singularities* at the two poles. Such features manifest as local distortions in grid quality, consisting of regions of highly distorted quadrilateral grid cells. These low-quality elements can lead to a number of undesirable numerical effects – imposing restrictions on model time step and stability and compromising local spatial accuracy. As a result, a majority of current-generation general circulation models are instead based on *semi-structured* quadrilateral discretisation schemes, including the cubed-sphere (Adcroft et al., 2004; Marshall et al.,

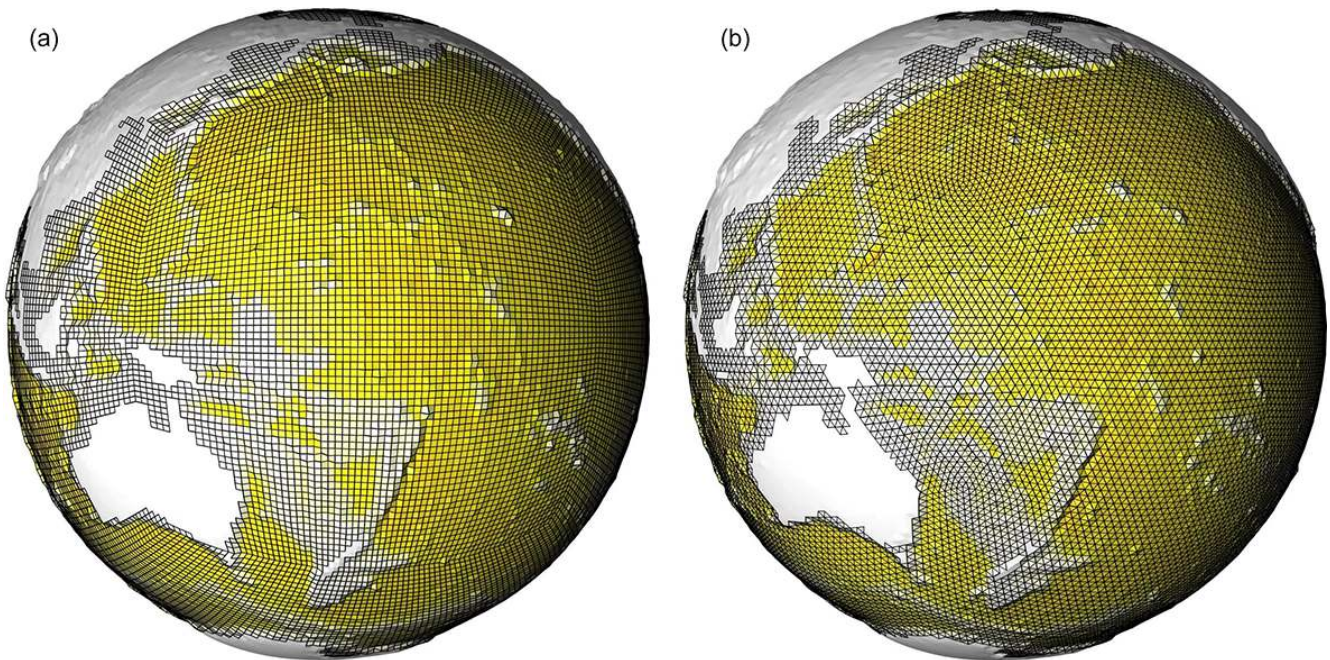


Figure 1. Conventional semi-structured meshes for the sphere, showing a cubed-sphere grid (a) and an icosahedral grid (b). Both grids were generated using equivalent target mean edge lengths and are coloured according to mean topographic height at grid cell centres. Topography is drawn using an exaggerated scale, with elevation from the reference geoid amplified by a factor of 20 in both cases.

1997; Putman and Lin, 2007) and tri-polar configurations (Murray, 1996; Madec et al., 2011; Bleck, 2002).

In the cubed-sphere framework, the spherical surface is decomposed into a cube-like topology, with each of the six quadrilateral faces discretised as a structured curvilinear grid. In such an arrangement, the two strong grid singularities of the lat–long configuration are replaced by eight weak discontinuities at the cube corners, leading to significant improvements in numerical performance. Putman and Lin (2007) presented detailed discussions of techniques for the generation and optimisation of cube-sphere-type grids. A regular *gnomonic-type* cubed-sphere grid is illustrated in Fig. 1a.

In the tri-polar grid, the present-day continental configuration is exploited to *bury* the singularities associated with a three-way polar decomposition of the sphere outside of the ocean mask. The resulting *numerically active* subset of the grid is well conditioned as a result. While such configurations are a popular choice for models designed for present-day Earth-centric ocean studies, the generality of these methods is clearly limited. In this study, we instead pursue the development of more general-purpose techniques, applicable to both oceanic and atmospheric modelling for general planetary and palaeo-Earth environments.

In addition to the standard cubed-sphere and tri-polar configurations, a second class of semi-structured spherical grid can be constructed through *icosahedral-type* decompositions (Heikes and Randall, 1995; Randall et al., 2002). In such cases, the primary grid is defined as a regular spherical trian-

gulation, obtained through recursive bisection of the icosahedron. The associated staggered polygonal *dual* grid, consisting of hexagonal and pentagonal cells, is often used as a basis for finite-volume-type numerical schemes. This *geometric duality* is an example of the locally orthogonal Voronoi–Delaunay-type grid staggering that forms the basis of this paper. Icosahedral-type grids provide a near-perfect tessellation of the sphere – free of topological discontinuities and/or geometrical irregularity. Such methods are applicable to both atmospheric- and oceanic-type simulations. A regular icosahedral-class grid is illustrated in Fig. 1b.

1.2 Unstructured grids

While the semi-structured grids described previously each provide effective frameworks for uniform resolution global simulation, the development of *multi-resolution* modelling environments requires alternative techniques. A range of new general circulation models, including the Finite Element Sea Ice–Ocean Model (FESOM) (Wang et al., 2014), the Finite Volume Community Ocean Model (FVCOM) (Chen et al., 2003, 2007; Lai et al., 2010), the Stanford Unstructured Non-hydrostatic Terrain-following Adaptive Navier–Stokes Simulator (SUNTANS) (Fringier et al., 2006; Vitousek and Fringer, 2014), and the Second-generation Louvain-la-Neuve Ice–ocean Model (SLIM) (Bernard et al., 2007; Comblen et al., 2009), are based on semi-structured triangular grids, with the horizontal directions discretised according to an unstructured spherical triangulation, and the ver-

tical direction represented as a stack of locally structured layers. The Model for Prediction Across Scales (MPAS) (Skamarock et al., 2012; Ringler et al., 2013, 2008) adopts a similar arrangement, except that a *locally orthogonal* unstructured discretisation is adopted, consisting of both a Spherical Voronoi Tessellation (SVT) and its dual Delaunay triangulation. The use of fully unstructured representations, based on general tetrahedral and/or polyhedral grids, are also under investigation in the Fluidity framework (Ford et al., 2004a, b; Pain et al., 2005; Piggott et al., 2008). Such models all impose different requirements on the *quality* of the underlying unstructured grids, with some models, including FE-SOM, SLIM, and Fluidity, offering additionally flexibility. In all cases though, the performance of the numerical simulation can be expected to improve with increased grid quality – encouraging the search for optimised grid-generation algorithms. Further discussion of grid-quality constraints for general circulation modelling is presented in Sect. 2.

Existing approaches for unstructured grid generation on the sphere have focused on a number of techniques, including (i) the use of iterative, optimisation-type algorithms designed to construct Spherical Centroidal Voronoi Tessellations (SCVTs) (Jacobsen et al., 2013), and (ii) the adaptation of anisotropic two-dimensional meshing techniques (Lambrechts et al., 2008) that build grids in associated parametric spaces.

The MPI-SCVT algorithm (Jacobsen et al., 2013) is a massively parallel implementation of iterative Lloyd-type smoothing (Du et al., 1999) for the construction of SCVTs for use in the MPAS framework. In this approach, a set of vertices are distributed over the spherical surface and iteratively *smoothed* until a high-quality Voronoi tessellation is obtained. Specifically, each iteration repositions vertices to the centroids of their associated Voronoi cells and updates the topology of the underlying spherical Delaunay triangulation. While such an approach typically leads to the generation of high-quality *centroidal* Voronoi tessellations, the algorithm does not provide theoretical guarantees on minimum element quality, and often requires significant computational effort to achieve convergence. Additionally, current implementations of the MPI-SCVT algorithm do not provide a mechanism to constrain the grid to embedded features, such as coastal boundaries.

Lambrechts et al. (2008) presented an unstructured spherical triangulation framework using the general-purpose grid-generation package Gmsh (Geuzaine and Remacle, 2009). In this work, unstructured spherical triangulations are generated for the world ocean using a parametric meshing approach. Specifically, a triangulation of the spherical surface is generated by *mapping* the full domain (including coastlines) on to an associated two-dimensional parametric space via stereographic projection. Importantly, as a result of the projection, the grids constructed in parametric space must be highly anisotropic, such that a well-shaped, isotropic triangulation is induced on the sphere. A range of existing two-dimensional

anisotropic meshing algorithms are investigated, including Delaunay-refinement, advancing-front, and adaptation-type approaches. In particular, the algorithm is designed to ensure a faithful representation of complex coastal boundary conditions. While a detailed model of such constraints is often neglected in global simulations, resolution of these features is a key factor for regional and coastal models. Several additional algorithms support the generation of unstructured grids for such two-dimensional domains, including the ADmesh package (Conroy et al., 2012) and the Stomel library (Holleman et al., 2013).

The current study describes a new algorithm for the generation of guaranteed-quality spheroidal Delaunay triangulations and Voronoi tessellations – appropriate for a range of unstructured general circulation models. In this work, meshes are generated on the spheroidal surface directly, without need for local parameterisation or projection. Such an approach will be shown to exhibit significant flexibility – immune to issues of coordinate singularity and/or continental configuration. The applicability of this approach to grid generation for imperfect spheres, including oblate spheroids and general ellipsoids is also discussed. Significant effort is invested to develop techniques designed to produce very high-quality multi-resolution grids appropriate for contemporary unstructured C-grid models, including the MPAS.

The paper is organised as follows: an overview of grid generation for general circulation modelling is presented in Sect. 2, outlining various constraints and requirements on minimum grid quality. A description of the Frontal-Delaunay refinement and hill-climbing-type optimisation algorithms is given in Sects. 3 and 4. A set of uniform and non-uniform grids appropriate for high-resolution, multi-scale general circulation modelling are presented Sect. 5, alongside an analysis of computational performance and optimality. Avenues for future work are outlined in Sect. 6.

2 Grid generation for general circulation modelling

Numerical formulations for atmospheric and/or oceanic general circulation modelling are typically based on a *staggered* grid configuration, with quantities such as fluid pressure, geopotential height, and density discretised using a primary control volume, and the fluid velocity field and vorticity distribution represented at secondary, spatially distinct grid points. Various structured staggered grid types are described by the well-known Arakawa schemes (Arakawa and Lamb, 1977).

The development of models based on unstructured grid types is an emerging area of research, and, as a result, a variety of numerical formulations are currently under investigation. In this study, the use of locally orthogonal grids appropriate for staggered unstructured C-grid schemes is pursued, as these methods represent a logical extension of conventional structured Arakawa-type techniques to the unstruc-

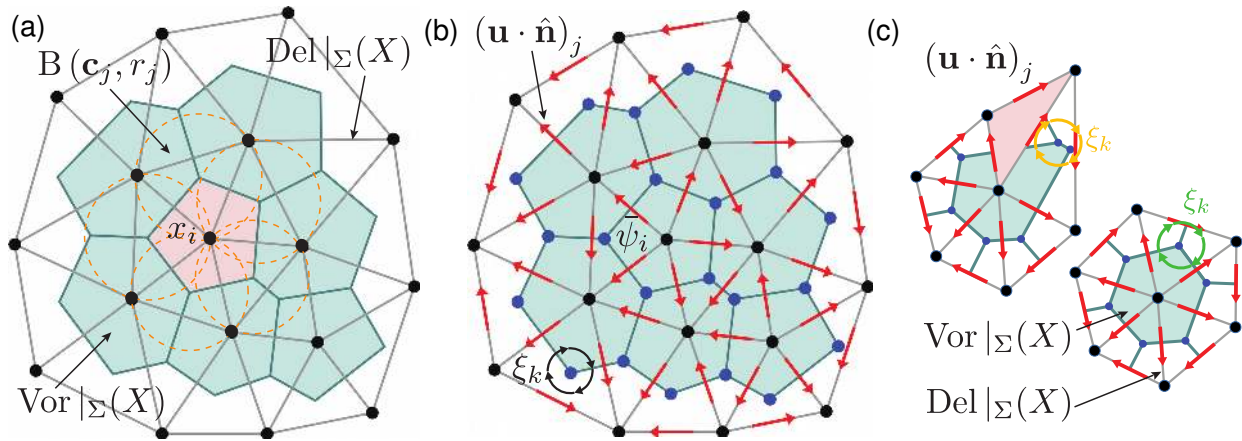


Figure 2. An anatomy of staggered unstructured grids for general circulation modelling, illustrating (a) locally orthogonal Voronoi–Delaunay grid staggering, (b) an unstructured C-grid-type numerical formulation, and (c) a comparison of *well-centred* and *poorly staggered* configurations. In panel (a), Voronoi polygons are formed by joining the centres of circumscribing balls $B(c_j, r_j)$ associated with adjacent Delaunay triangles. In panel (b), the C-grid scheme consists of conservative *cell-centred* tracer quantities ψ_i , *edge-centred* normal velocity components $(\mathbf{u} \cdot \hat{\mathbf{n}})_j$, and auxiliary *vertex-centred* vorticity variables ξ_k . In panel (c), the well-centred (lower) and poorly staggered (upper) configurations differ in the associativity of Voronoi vertices. In the well-centred configuration, Voronoi vertices are *interior* to their parent triangles. Reconstruction of the vertex-centred vorticity distribution can thus be achieved by integrating over the Delaunay triangles. In the poorly staggered configuration, a single Voronoi vertex is exterior to its parent triangle (shaded), leading to a breakdown in the reconstruction. Note that adjacent Voronoi–Delaunay edges do not intersect in the poorly staggered configuration.

tured setting. Such formulations require that the grids satisfy a local orthogonality constraint, with adjacent grid cell edges in the primary and secondary control volumes required to be mutually perpendicular. In the unstructured setting, it is known that the Delaunay triangulation and the Voronoi tessellation constitute a locally orthogonal staggered dual, leading to a natural framework for the construction of such meshes.

Given a Delaunay triangulation $\text{Del}|_{\Sigma}(X)$ of the surface Σ , the associated Voronoi diagram $\text{Vor}|_{\Sigma}(X)$ can be constructed by exploiting this underlying geometric duality. Each edge in the Delaunay triangulation is associated with a *dual edge* in the Voronoi diagram, spanning between the centres of the circumscribing balls $B(c_j, r_j)$ associated with the two adjacent Delaunay triangles. The full Voronoi diagram is assembled from these edges – consisting of a set of (convex) polygonal grid cells centred on each vertex in the underlying surface triangulation. See Fig. 2a for an example of this unstructured Voronoi–Delaunay-type grid staggering.

In Fig. 2b, an unstructured variant of the widely used Arakawa C-grid scheme is described, with fluid pressure, geopotential height, and density degrees of freedom placed within the primary Voronoi control volumes, and a set of orthogonal velocity vectors positioned along Delaunay edges. Additional vorticity degrees of freedom are located at the vertices of the Voronoi grid cells. Such an arrangement facilitates the construction of a standard conservative finite-volume-type scheme for the transport of cell-centred fluid properties, and a *mimetic-class* (Lipnikov et al., 2014; Bochev and Hyman, 2006)

finite-difference formulation for the evolution of velocity components. Overall, this scheme is known to possess a variety of desirable conservation properties, conserving mass, potential vorticity and enstrophy, and preserving geostrophic balance (Ringler et al., 2010). This unstructured C-grid scheme is currently employed in the MPAS for both atmospheric and oceanic modelling (Skamarock et al., 2012; Ringler et al., 2013, 2008).

While numerically elegant, such schemes are not applicable to general unstructured Voronoi–Delaunay grids – requiring that a number of auxiliary geometrical constraints also be satisfied. Specifically, such schemes necessitate the use of grids that are not only locally orthogonal but also well-centred and mutually centroidal. These additional conditions are constraints on the geometry of the underlying Delaunay triangles and Voronoi polygons.

A *well-centred* Delaunay grid is one in which all dual Voronoi vertices lie within the interior of their associated Delaunay triangles. Such a condition guarantees that adjacent Delaunay and Voronoi edges intersect – inducing a dual structure that is *nice*ly staggered. In the context of the unstructured C-grid scheme described previously, this configuration guarantees that a consistent stencil exists for reconstruction of the discrete vorticity variable. In practice, the generation of well-centred grids is known to be a particularly onerous task (VanderZee et al., 2008; Vanderzee et al., 2010), requiring that the triangulation consist of *all acute* elements. Degradation of the vorticity reconstruction for poorly staggered grids is described in Fig. 2c.

Delaunay grids are centroidal when their primary and secondary vertices lie at the centres of mass of their associated dual grid cells, with the vertices of the Voronoi polygons lying at the centroids of the Delaunay triangles and visa versa. Such a condition is effectively an implicit constraint on the *regularity* of the grid, with strongly centroidal meshes typically associated with improved grid cell shapes. In the limit, a perfectly centroidal and well-centred tessellation can be assembled from equilateral triangles and regular hexagons. Centroidal grids typically lead to very high-quality numerical discretisations, with grids containing near-perfect element configurations achieving *optimal* convergence rates. The unstructured C-grid scheme described previously is known to achieve second-order accurate convergence when applied to centroidal Voronoi grids (Ringler et al., 2010).

The generation of locally orthogonal, well-centred, and centroidal unstructured grids is a difficult task, imposing a heavy burden on the underlying mesh generation algorithm. While a number of techniques for unstructured grid generation currently exist, as per Sect. 1, the author is not aware of any that are successful in generating the very high-quality grids required by unstructured C-grid-type general circulation models. As a result, throughout the remainder of this paper, the development of new methods for the generation of such tessellations is pursued in detail.

3 A restricted Frontal-Delaunay refinement algorithm

The task is to generate very high-resolution, guaranteed-quality unstructured Delaunay triangulations for planetary atmospheres and/or oceans. These grids will form a baseline for the hill-climbing mesh-optimisation techniques presented in subsequent sections. Such grids are required to satisfy a number of constraints, including bounds on minimum element quality and adherence to user-defined mesh-spacing distributions. In this work, the applicability of a recently developed Frontal-Delaunay surface meshing algorithm (Engwirda and Ivers, 2016; Engwirda, 2016) is investigated for this task.

An unstructured Delaunay triangulation of the reference spheroid associated with a general planetary geometry is sought. In a general form, this reference surface can be expressed as an axis-aligned triaxial ellipsoid

$$\sum_{i=1}^3 \left(\frac{x_i}{r_i} \right)^2 = 1, \quad (1)$$

where the x_i 's are the Cartesian coordinates in a locally aligned coordinate system, and the scalars $r_i > 0$ are its principal radii. Such a definition can be used to represent ellipsoidal surfaces in general position, based on the application of additional rigid-body translations and rotations. Note that while grid generation for global climate modelling is often restricted to spherical surfaces (e.g. $r_{1,2,3} =$

6371 km) this formulation facilitates mesh generation for general spheroidal and ellipsoidal domains.

3.1 Preliminaries

Before describing the Frontal-Delaunay refinement algorithm in full, several key properties of the underlying Delaunay triangulation and Voronoi tessellation are briefly reviewed. A full account of these structures is not presented here; instead, the reader is referred to the detailed theoretical exposition presented in, for example, Cheng et al. (2013).

The Delaunay triangulation $\text{Del}(X)$ associated with a set of points $X \in \mathbb{R}^d$ is characterised by the *empty-circle* criterion – requiring that the set of circumscribing spheres $B(\mathbf{c}_i, r_i)$ associated with each Delaunay triangle $\tau_i \in \text{Del}(X)$ be *empty* of all points other than its own vertices. It is well known that for tessellations restricted to two-dimensional manifolds, the Delaunay triangulation leads to a maximisation of the minimum enclosed angle in the grid (Cheng et al., 2013). Such behaviour is clearly beneficial when seeking to construct high-quality triangular meshes.

The Voronoi tessellation $\text{Vor}(X)$ is the geometric dual associated with the Delaunay triangulation, consisting of a set of convex polygonal cells formed by connecting the centres of adjacent circumscribing balls – the element circumcentres \mathbf{c}_i 's. The Voronoi tessellation represents a *closest-point map* for the points in X , with each Voronoi cell $v_c \in \text{Vor}(X)$ defining the convex region adjacent to a given vertex $\mathbf{x}_i \in X$ for which \mathbf{x}_i is the nearest point. Importantly, the Voronoi–Delaunay grid staggering, defines a locally orthogonal arrangement, in which grid cell edges in the Voronoi tessellation are orthogonal to adjacent edges in the underlying Delaunay triangulation.

3.2 Restricted Delaunay triangulation

In this study, grid generation is carried out directly on the spheroidal geometry by making use of *restricted* Delaunay mesh generation techniques. Specifically, given a reference surface Σ , grid generation proceeds to discretise the surface into a mesh of triangles. In the restricted Delaunay framework, a full-dimensional Delaunay tessellation $\text{Del}(X)$ (i.e. a tetrahedral tessellation) is maintained, with the surface triangulation represented as a subset of the tetrahedral faces. The restricted Delaunay surface triangulation $\text{Del}|_{\Sigma}(X)$ is said to be *embedded* in $\text{Del}(X)$ as a result. Use of this fully three-dimensional approach elides any reliance on local parametric projections.

Definition 1 (restricted Delaunay tessellation). Let Σ be a smooth surface embedded in \mathbb{R}^3 . Let $\text{Del}(X)$ be a full-dimensional Delaunay tetrahedralisation of a point-wise sample $X \subseteq \Sigma$ and $\text{Vor}(X)$ be the associated Voronoi tessellation. The restricted Delaunay surface triangulation $\text{Del}|_{\Sigma}(X)$ is a sub-complex of $\text{Del}(X)$ including any triangle

$f_i \in \text{Del}(X)$ associated with an intersecting Voronoi edge $v_e \in \text{Vor}(X)$ such that $v_e \cap \Sigma \neq \emptyset$.

The development of restricted Delaunay techniques for general mesh generation applications has been the subject of previous research, and a detailed discussion of such concepts is presented elsewhere. The reader is referred to the original work of Edelsbrunner and Shah (1997) or the detailed reviews presented in Cheng et al. (2013) for additional details and mathematical background.

In the context of this work, it is sufficient to note that the restricted Delaunay framework provides a convenient mechanism for the construction of grids on the spheroidal surface Σ , without requiring the introduction of parametric mappings or anisotropic metrics. Implementation of the restricted Delaunay scheme requires the definition of a single *geometric predicate* – a local comparison used to determine which faces of the underlying tetrahedral complex lie on the spheroidal surface. Specifically, intersections between the edges of the full-dimensional Voronoi complex and the underlying surface Σ are computed. Triangles associated with non-empty intersections $\text{Vor}|_f(X) \cap \Sigma \neq \emptyset$ form part of the surface mesh. In this study, these comparisons are computed analytically, following standard spheroidal trigonometric manipulations, as detailed in Appendix A. See Fig. 3 for a detailed schematic.

3.3 Mesh-spacing functions

The local *density* of the mesh can be controlled via a user-specified mesh-spacing function $\bar{h}(\mathbf{x}) : \mathbb{R}^3 \rightarrow \mathbb{R}^+$, where $\bar{h}(\mathbf{x})$ defines the *target* edge lengths at all points on the surface Σ . In this work, mesh-spacing functions are specified as a discrete set of target values $\bar{h}_{i,j}$, defined on a simple background lat–long grid \mathcal{G} . The continuous mesh-spacing function $\bar{h}(\mathbf{x})$ is reconstructed using bilinear interpolation. As will be illustrated in subsequent sections, such an approach provides support for a wide range of mesh-spacing definitions, including those derived from high-resolution topographic data (Amante and Eakins, 2009) or solution-adaptive metrics.

In order to generate high-quality grids, it is necessary to ensure that the imposed mesh-spacing function is *sufficiently smooth*. Rather than requiring the user to accommodate such constraints, a Lipschitz smoothing process is adopted here. Following the work of Persson (2006), a *gradient-limited* mesh-spacing function $\bar{h}'(\mathbf{x})$ is constructed by constraining the allowable spatial fluctuation over each element in the background grid \mathcal{G} . In this study, a scalar smoothing parameter $g \in \mathbb{R}^+$ is used to limit variation, such that

$$\bar{h}'(\mathbf{x}_i) \leq \bar{h}'(\mathbf{x}_j) + g \cdot \text{dist}(\mathbf{x}_i, \mathbf{x}_j), \quad (2)$$

for all adjacent vertex pairs $\mathbf{x}_i, \mathbf{x}_j$ in the underlying grid \mathcal{G} . The gradient-limited mesh-spacing function $\bar{h}'(\mathbf{x})$ becomes

more uniform as $g \rightarrow 0$. In this work, gradient limiting is implemented following a *fast-marching* method, as described in Persson (2006).

3.4 Grid-quality metrics

The quality of a mesh can be evaluated using various grid-quality metrics. Here, a number of standard measures are introduced.

Definition 2 (radius-edge ratio). Given a surface triangle $f_i \in \text{Del}|_\Sigma(X)$, its radius-edge ratio, $\rho(f_i)$, is given by

$$\rho(f_i) = \frac{r_i}{\|e_{\min}\|}, \quad (3)$$

where r_i is the radius of the circumscribing ball associated with f_i and $\|e_{\min}\|$ is the length of its shortest edge.

The radius-edge ratio is a measure of element shape quality, achieving a minimum of $\frac{1}{\sqrt{3}}$ for equilateral triangles and increasing as elements tend toward degeneracy. It is directly related to the minimum-angle θ_{\min} subtended by triangle edges, such that $\theta_{\min} = \arcsin(\frac{1}{2\rho})$. Due to the summation of angles in a triangle, this relationship extends to the maximum angle, with $\theta_{\max} \leq \pi - 2\theta_{\min}$.

Definition 3 (area-length ratio). Given a surface triangle $f_i \in \text{Del}|_\Sigma(X)$, its area-length ratio, $a(f_i)$, is given by

$$a(f_i) = \frac{4\sqrt{3}}{3} \frac{A_f}{\|e_{\text{rms}}\|^2}, \quad (4)$$

where A_f is the signed area of f_i and $\|e_{\text{rms}}\|$ is the root-mean-square edge length.

The area-length ratio is a robust, scalar measure of element shape quality, and is typically normalised to achieve a score of +1 for ideal elements. The area-length ratio decreases with increasing distortion, achieving a score of +0 for degenerate elements and –1 for tangled elements with reversed orientation.

Definition 4 (relative-edge length). Given an edge in the surface tessellation $e_j \in \text{Del}|_\Sigma(X)$, its relative-edge length, $h_r(e_j)$, is given by

$$h_r(e_j) = \frac{\|e_j\|}{\bar{h}(\mathbf{x}_m)}, \quad (5)$$

where $\|e_j\|$ is the length of the j -th edge and $\bar{h}(\mathbf{x}_m)$ is the value of the mesh-spacing function sampled at the edge midpoint.

The relative length $h_r(e_j)$ is a measure of mesh-spacing conformance, expressing the ratio of actual-to-desired edge-length. A value of $h_r(e_j) = 1$ indicates perfect mesh-spacing conformance.

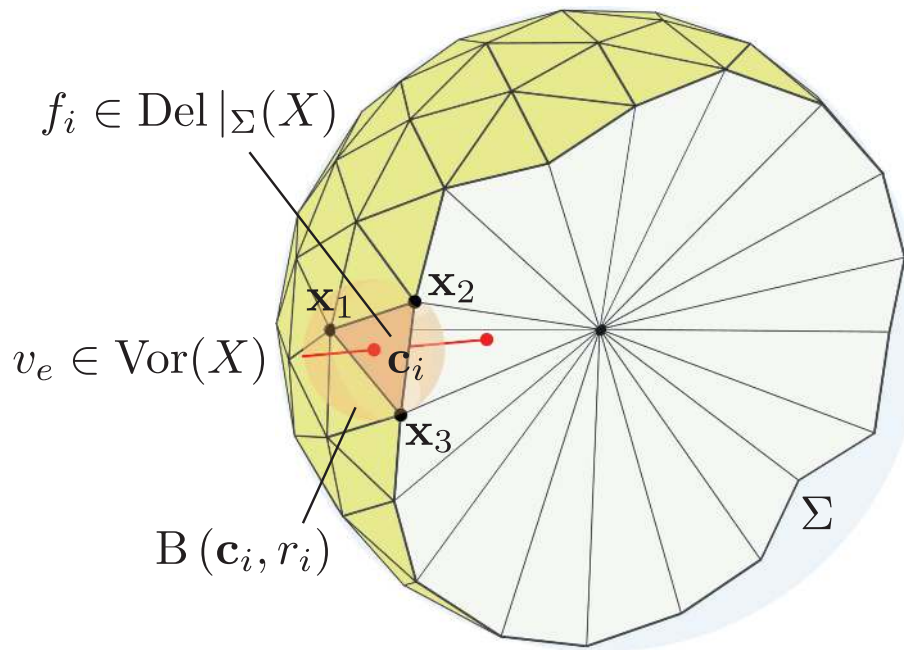


Figure 3. Illustration of the geometrical predicates used to define the restricted Delaunay surface triangles $f_i \in \text{Del}|_{\Sigma}(X)$, showing details of the intersecting Voronoi edge $v_e \in \text{Vor}(X)$ associated with the surface triangle $f_i \in \text{Del}|_{\Sigma}(X)$ and its associated surface ball $B(c_i, r_i)$. Note that the full three-dimensional Delaunay tessellation $\text{Del}(X)$ is a tetrahedral complex that fills the interior of the spheroid.

3.5 Restricted Frontal-Delaunay refinement

A high-quality triangular surface mesh is generated on the spheroidal reference surface Σ using a Frontal-Delaunay variant of the conventional restricted Delaunay refinement algorithm (Boissonnat and Oudot, 2003, 2005; Jamin et al., 2015; Cheng et al., 2007, 2010). This technique is described by the author in detail in Engwirda and Ivers (2016); Engwirda (2016) and differs from standard Delaunay refinement approaches in terms of the strategies used for the placement of new vertices. Specifically, the Frontal-Delaunay algorithm employs a generalisation of various *off-centre* point-placement techniques (Rebay, 1993; Erten and Üngör, 2009), designed to position vertices such that element-quality and mesh-size constraints are satisfied in a *locally optimal* fashion. Previous studies have shown that such an approach typically leads to substantial improvements in mean element quality and mesh smoothness.

While a variety of solutions for the surface meshing problem have been described previously (e.g. Frey and George, 2000; Ryppl and Krysl, 1997; Schreiner et al., 2006; Löhner, 1996), the restricted Frontal-Delaunay algorithm presented here has been found to offer a unique combination of characteristics – combining the smooth, high-quality grid-generation capabilities of an advancing-front approach, with the theoretical robustness and provable guarantees associated with conventional Delaunay refinement techniques. Specifically, the Frontal-Delaunay algorithm described here is known to generate grids with very high mean element qual-

ity, bounded minimum and maximum angles, tight conformance to grid-spacing constraints, and provable guarantees on topological consistency and convergence. A full description of the algorithm, including detailed discussions of its theoretical foundations and proofs of worst-case grid-quality bounds, can be found in Engwirda and Ivers (2016); Engwirda (2016).

Given a user-defined mesh-spacing function $\bar{h}(x)$ and an upper bound on the element *radius-edge* ratios $\bar{\rho} \geq 1$, the Frontal-Delaunay algorithm proceeds to sample the spheroidal surface Σ by refining any surface triangle that violates either the mesh-spacing or element-quality constraints. The full algorithm is described in Algorithm 1, and detailed snap shots of the refinement process are shown in Fig. 4.

Initially, the surface is seeded with a set of 12 vertices to form a standard icosahedron. Refinement then proceeds to insert a new vertex at the off-centre refinement point associated with each element to be *eliminated*. Refinement continues until all constraints are satisfied. The refinement process is priority scheduled, with triangles $f_i \in \text{Del}|_{\Sigma}(X)$ ordered according to their radius-edge ratios $\rho(f_i)$. This ordering ensures that the element with the *worst* ratio is refined at each iteration. Additionally, triangles are subject to a *frontal* filtering – requiring that a low-quality element be adjacent to a *converged* triangle before being considered for refinement. This logic helps the algorithm mimic the behaviour of *advancing-front-type* techniques, with new vertices and elements expanding from initial seeds. Upon termination, the resulting surface triangulation is guaranteed to contain nicely

Algorithm 1 Restricted frontal Delaunay refinement

```

1: function DELFRONT( $\Sigma, \bar{\rho}, \bar{h}(\mathbf{x}), \text{Del}|_{\Sigma}(X)$ )
2:   Place 12 vertices on the spheroidal surface  $\Sigma$  to
   form a regular icosahedron. Initialise the full
   Delaunay complex  $\text{Del}(X)$  and the restricted
   surface triangulation  $\text{Del}|_{\Sigma}(X)$ .
3:   while ( $\exists \text{BADTRIANGLE}(\text{Del}|_{\Sigma}(X))$ ) do
4:     Find the worst bad frontal triangle in the
     surface triangulation  $f_i \leftarrow \text{Del}|_{\Sigma}(X)$ .
5:     Call  $\mathbf{x}_i \leftarrow \text{OFFCENTRE}(f_i)$  to form the new
     off-centre vertex  $\mathbf{x}_i$ , designed to eliminate
     the bad triangle  $f_i$ .
6:     Push  $\text{Del}(X) \leftarrow \mathbf{x}_i$  and update the re-
     stricted surface triangulation  $\text{Del}|_{\Sigma}(X)$ .
7:   end while
8:   return surface triangulation object  $\text{Del}|_{\Sigma}(X)$ .
9: end function
1: function BADTRIANGLE( $f$ )
   if ( $h(f) > \bar{h}(\mathbf{x}_f)$ )           return TRUE
2:   if ( $\rho_2(f) > \bar{\rho}$ )             return TRUE
   else return FALSE
3: end function
1: function OFFCENTRE( $f, \mathbf{x}$ )
2:   Form the restricted triangle circumcentre
    $\mathbf{x}_c \leftarrow \text{Vor}|_f(X) \cap \Sigma$  and the associated sur-
   face ball  $B = \text{SDB}(f)$ .
3:   Find the minimum length edge  $e$  in the triangle
    $f$ . Edge  $e$  is the frontal segment.
4:   Place the size-optimal point  $\mathbf{x}_h$  along the arc
    $\text{Vor}|_e(X) \cap \Sigma$  to satisfy local spacing con-
   straints  $\bar{h}(\mathbf{x})$ .
5:   Place the shape-optimal point  $\mathbf{x}_\theta$  along the arc
    $\text{Vor}|_e(X) \cap \Sigma$  to satisfy minimum angle con-
   straints  $\rho(f) \leq \bar{\rho}$ .
6:   Compute the distances  $d_h = \text{dist}(\mathbf{x}_e, \mathbf{x}_h)$  and
    $d_\theta \leftarrow \text{dist}(\mathbf{x}_e, \mathbf{x}_\theta)$ , where  $\mathbf{x}_e$  is the midpoint of
   the short edge  $e$ .
7:   Set  $\mathbf{x}$  to the point  $\{\mathbf{x}_h, \mathbf{x}_\theta\}$  of minimum dis-
   tance  $d_j$ , such that  $d_j \geq \frac{1}{2}\|e\|$ . Fall-back to
    $\mathbf{x} \leftarrow \mathbf{x}_c$  if  $\{\mathbf{x}_h, \mathbf{x}_\theta\}$  do not satisfy constraints.
8:   return off-centre vertex  $\mathbf{x}$ .
9: end function

```

shaped elements, satisfying both the radius-edge constraints $\rho(f_i) \leq \bar{\rho}$ and mesh-spacing bounds $h(f_i) \leq \bar{h}(\mathbf{x}_f)$. Setting $\bar{\rho} = 1$ guarantees that element angles are bounded above 30° , ensuring that the grid does not contain any highly distorted triangles.

3.6 Off-centre point placement

The performance of the Frontal-Delaunay algorithm hinges on the use of off-centre refinement rules – locally optimal point-placement strategies designed to create very high-quality vertex distributions. A set of candidate off-centre points are considered at each new vertex insertion. Type I vertices, \mathbf{x}_c , are equivalent to conventional element circumcentres (positioned at the centre of the associated element circumballs), and are used to preserve global convergence. Type II vertices, \mathbf{x}_h , are *size-optimal* points, and are designed to satisfy grid-spacing constraints in a locally optimal fashion. Type III vertices, \mathbf{x}_θ , are *shape-optimal* points, and are designed to ensure minimum-angle bounds are satisfied in a worst-first manner. The type II and type III strategies employed here can be seen as a generalisation of the two-dimensional off-centre techniques presented by Rebay (1993) and Erten and Üngör (2009), respectively. See Fig. 5 for additional detail.

Given a low-quality triangle $f_i \in \text{Del}|_{\Sigma}(X)$ to be eliminated, the type II and type III vertices \mathbf{x}_h and \mathbf{x}_θ are positioned along the intersection of an adjacent segment of the

Voronoi complex and the spheroidal surface. This intersection $\text{Vor}|_e(X) \cap \Sigma$, defines a *frontal curve* inscribed on Σ – a *locally optimal* geodesic segment on which to insert new vertices. Here, $\text{Vor}|_e(X)$ is the polygonal face of the Voronoi complex associated with the short edge $e_0 \in f_i$. In the context of conventional advancing-front-type methods, the edge e_0 would denote the current *frontal segment* about which vertex insertion occurs.

The points \mathbf{x}_h and \mathbf{x}_θ are positioned to form new candidate triangles about the frontal edge e_0 , such that local constraints are satisfied optimally. The size-optimal point \mathbf{x}_h is positioned to adhere to local grid-spacing constraints, ensuring that new edges satisfy $\|e_i\| \leq \bar{h}(\mathbf{m}_i)$, where the \mathbf{m}_i 's are the edge midpoints. These constraints can be solved for an associated altitude

$$a_h = \min\left(\tilde{a}_h, \frac{1}{2}\sqrt{3\tilde{h}}\right), \quad (6)$$

where

$$\tilde{a}_h = \sqrt{\tilde{h}^2 - \frac{1}{2}\|e_0\|^2}, \quad \tilde{h} = \frac{\bar{h}(\mathbf{m}_1) + \bar{h}(\mathbf{m}_2)}{2}. \quad (7)$$

Similarly, the shape-optimal point is positioned to adhere to minimum-angle constraints, sliding \mathbf{x}_θ along the inscribed curve $\text{Vor}|_e(X) \cap \Sigma$ so that the new candidate triangle satisfies $\rho \leq \bar{\rho}$. These constraints lead to an associated altitude

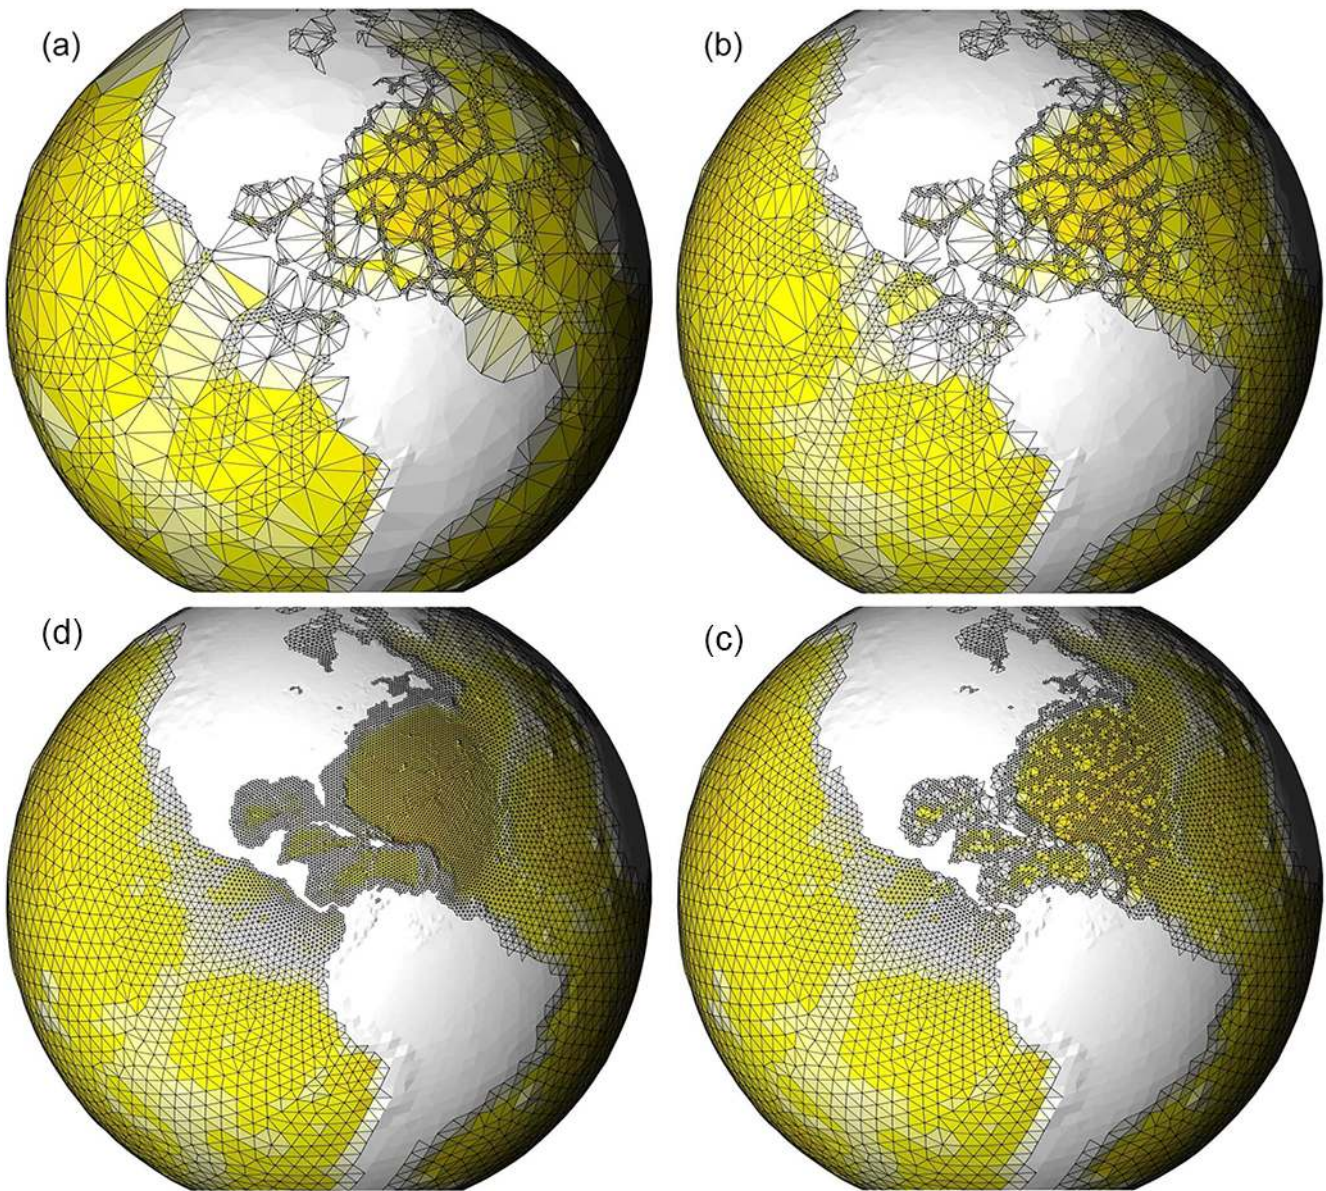


Figure 4. Progress of the Frontal-Delaunay refinement algorithm, clockwise from top left, showing the *space-filling* refinement trajectory.

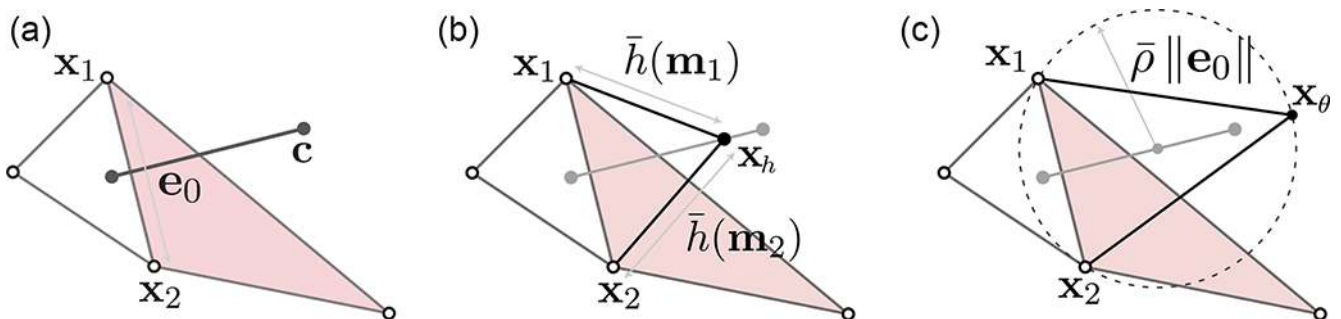


Figure 5. Two-dimensional representation of the off-center point-placement rules, illustrating (a) a low-quality triangle f_i (shaded) to be eliminated, highlighting the *frontal* segment of the Voronoi diagram associated with the short edge $e_0 \in f_i$, (b) placement of the size-optimal point x_h to satisfy local size constraints $\bar{h}(x_f)$, and (c) placement of a shape-optimal point x_θ to satisfy local radius-edge ratios $\bar{\rho}$.

$$a_\theta = \frac{\|e_0\|}{\tilde{\beta}_\theta}, \quad (8)$$

where

$$\tilde{\beta}_\theta = 2 \tan\left(\frac{\tilde{\theta}}{2}\right), \quad \tilde{\theta} = \arcsin\left(\frac{1}{2\bar{\rho}}\right). \quad (9)$$

Given a_h and a_θ , the position of the points \mathbf{x}_h and \mathbf{x}_θ is calculated by computing the intersection of balls of radius a_h and a_θ , centred at the midpoint of the frontal edge e_0 and the frontal curve $\text{Vor}|_e(X) \cap \Sigma$. This approach ensures that new vertices are positioned by advancing a specified distance along the surface Σ in the frontal direction. For non-uniform $\bar{h}(\mathbf{x})$, expressions for the position of the point \mathbf{x}_h are non-linear, with the altitude a_h depending on an evaluation of the mesh-size function at the edge midpoints $\bar{h}(\mathbf{m}_i)$ and visa versa. In practice, since $\bar{h}(\mathbf{x})$ is guaranteed to be smooth, a simple iterative predictor–corrector procedure is sufficient to solve these expressions approximately.

Finally, given the set of candidate off-centre vertices $\{\mathbf{x}_c, \mathbf{x}_h, \mathbf{x}_\theta\}$, the position of the refinement point \mathbf{x} for the triangle f_i is selected. A *worst-first* strategy is adopted, choosing the point that satisfies local constraints in a greedy fashion. Specifically, the closest point lying on the adjacent Voronoi segment $\text{Vor}|_e(X)$ and outside the neighbourhood of the frontal edge e_0 is selected, with

$$\mathbf{x} = \begin{cases} \mathbf{x}_j, & \text{if } \left(d_j \leq d_c \text{ and } d_j \geq \frac{1}{2}\|e_0\|\right), \\ \mathbf{x}_c, & \text{otherwise} \end{cases}, \quad (10)$$

where $j = \text{argmin}(d_h, d_\theta)$, and $d_j = \text{dist}(\mathbf{x}_e, \mathbf{x}_j)$ are distances from the midpoint of the frontal edge e_0 to the size- and shape-optimal points \mathbf{x}_h and \mathbf{x}_θ , respectively. This cascading selection criteria seeks a balance between local optimality and global convergence – smoothly degenerating to a conventional circumcentre-based Delaunay refinement strategy in limiting cases, while using locally optimal points where possible. Specifically, these constraints guarantee that refinement points always lie within a local *safe* region on the Voronoi complex – being positioned on an adjacent Voronoi segment and bound between the circumcentre of the element itself and the diametric ball of the associated frontal edge. This condition ensures that new points are never positioned too close to an existing vertex, leading to provable guarantees on the performance of the algorithm. See previous work by the author (Engwirda and Ivers, 2016; Engwirda, 2016) for additional detail.

3.7 Additional remarks

As a *restricted* Delaunay refinement approach, a full three-dimensional Delaunay tetrahedralisation $\text{Del}(X)$ is incrementally maintained throughout the surface meshing phase.

As per definition 1, the set of restricted surface triangles $\text{Del}|_\Sigma(X)$ that lie on the spheroid Σ , are expressed as a subset of the faces of this tetrahedral complex, such that $\text{Del}|_\Sigma(X) \subseteq \text{Del}(X)$. In an effort to minimise the expense of maintaining this full-dimensional tessellation, an additional *scaffolding* vertex \mathbf{x}_s is initially inserted at the centre of the spheroid. This has the effect of simplifying the topological structure of $\text{Del}(X)$, with the resulting tetrahedral elements forming a simple *wheel-like* configuration – emanating radially outward from the scaffolding vertex \mathbf{x}_s . See Fig. 3 for additional detail.

As shown in Fig. 4, one interesting characteristic of the Frontal-Delaunay algorithm described here relates to the *refinement trajectory* that is followed when inserting new vertices and triangles. Unlike standard Delaunay refinement or advancing-front-type methods, it can be seen that the algorithm adopts a *space-filling curve-type* pattern, covering the surface in a fractal-like configuration, before recursively filling in the gaps. Note that no explicit space-filling curve constraint has been implemented here – this behaviour is simply an emergent property of the algorithm itself, and is due to interactions between the greedy priority schedule, the frontal filtering and the off-centre point-placement strategies. In practice, it has been found that this space-filling-type behaviour leads to the construction of very high-quality triangulations, typically exceeding the performance of standard advancing-front or Delaunay refinement-type schemes.

4 Hill-climbing mesh optimisation

While the Delaunay grids produced by the refinement algorithm described previously are guaranteed to be of very high-quality, producing triangulations with angles bounded between $\theta_f = 30^\circ$ and $\theta_f = 120^\circ$, these tessellations can often be further improved through subsequent *mesh-optimisation* operations. Recalling that the construction of well-centred grids appropriate for unstructured C-grid schemes *require* that maximum angles be bounded below $\theta_f = 90^\circ$, the application of such optimisation procedures can in fact be seen as a *necessary* component of the grid-generation work flow explored here.

In the present work, grid optimisation is realised as a coupled geometrical and topological optimisation task – seeking to reposition vertices and update grid topology to maximise a given element-wise mesh-quality metric $\mathcal{Q}_f(\mathbf{x})$. A *hill-climbing-type* optimisation strategy is pursued, in which a locally optimal solution is sought based on an initial grid configuration. In this study, the grid is optimised according to the *area-length*-quality metric (see definition 3), a robust scalar measure of grid-quality that achieves a score of +1 for *perfect* elements – decreasing toward zero with increasing levels of element distortion. Optimisation predicates are implemented following a hill-climbing strategy (Freitag and Ollivier-Gooch, 1997; Klingner and Shewchuk, 2008),

with modifications to the grid accepted only if the local mesh-quality metrics are sufficiently improved. Specifically, a worst-first strategy is adopted here, in which each given operation is required to improve the *worst-case* quality associated with elements in the local subset acted upon. Such a philosophy ensures that global mesh quality is increased monotonically as optimisation proceeds. Note that such behaviour is designed to maximise the quality of the worst elements in the grid, rather than improving a mean measure. This represents an important distinction when compared to other iterative mesh-optimisation algorithms, such as Centroidal Voronoi Tessellation (CVT)-type schemes (Jacobsen et al., 2013), in which all entities in the grid are adjusted simultaneously until a global convergence criterion is satisfied.

4.1 “Spring”-based mesh smoothing

Considering firstly the geometric optimality of the grid, a *mesh-smoothing* procedure is undertaken – seeking to reposition the nodes of the grid to improve element quality and mesh-spacing conformance. Following the work of Persson and Strang (2004), a *spring-based* approach is pursued, in which edges in the Delaunay triangulation are treated as *elastic rods* with a prescribed natural length. Nodes are iteratively repositioned until a local equilibrium configuration is reached. See Algorithm 2 for full details. In the original work of Persson, nodal positions are adjusted via a local time-stepping loop, with all nodes updated concurrently under the action of explicit spring forces. In the current study, a non-iterative variant is employed, in which each node is repositioned one by one, such that constraints in each local neighbourhood are satisfied directly. Specifically, a given node \mathbf{x}_i is repositioned as a weighted sum of contributions from incident edges

$$\mathbf{x}_i^{n+1} = \frac{\sum w_k (\mathbf{x}_i^n + \Delta_k \mathbf{v}_k)}{\sum w_k}, \quad (11)$$

where

$$\mathbf{v}_k = \mathbf{x}_i^n - \mathbf{x}_j^n \quad \text{and} \quad \Delta_k = \frac{\bar{h}(\mathbf{x}_k^n) - l_k}{l_k}. \quad (12)$$

Here, $\mathbf{x}_i^n, \mathbf{x}_j^n$ are the current positions of the two nodes associated with the k -th edge, l_k is the edge length and $\bar{h}(\mathbf{x}_k)$ is the value of the mesh-spacing function evaluated, at the edge midpoint. Δ_k is the relative spring *extension* required to achieve equilibrium in the k -th edge. The scalars $w_k \in \mathbb{R}^+$ are edge weights. Setting $w_k = 1$ results in an unweighted scheme, consisting of simple *linear* springs. In this study, the use of non-linear weights, defined by setting $w_k = \Delta_k^2$, was found to offer superior performance.

Noting that application of the spring-based operator (Eq. 11) may move nodes away from the underlying spheroidal surface Σ , an additional *projection* operator is introduced to ensure that the grid conforms to the surface

geometry exactly. Following the application of each spring-based adjustment, nodes are moved back onto the geometry via a closest-point projection.

Consistent with the hill-climbing paradigm described previously, each nodal adjustment is required to be *validated* before being committed to the updated grid configuration. Specifically, nodal adjustments are accepted only if there is sufficient improvement in the mesh-quality metrics associated with the set of adjacent elements. A sorted comparison of quality metrics before and after nodal repositioning is performed, with nodal adjustments successful if grid quality is improved in a worst-first manner. This *lexicographical*-quality comparison is consistent with the methodology employed in Klingner and Shewchuk (2008).

4.2 Gradient-based mesh smoothing

While the spring-based smoothing operator is effective in adjusting a grid to satisfy mesh-spacing constraints, and tends to improve the quality of the grid on average, it is not guaranteed to improve element quality in all cases. Therefore, an additional *steepest ascent* strategy is pursued (Freitag and Ollivier-Gooch, 1997), in which nodal positions are adjusted based on the local gradients of incident element-quality functions. See Algorithm 2 for full details. Specifically, a given node \mathbf{x}_i is repositioned along a local *search vector* chosen to improve the quality of the worst incident element

$$\mathbf{x}_i^{n+1} = \mathbf{x}_i^n + \Delta_f \hat{\mathbf{v}}_f, \quad (13)$$

where

$$\mathbf{v}_f = \frac{\partial}{\partial \mathbf{x}} (\mathcal{Q}_f(\mathbf{x}_i)), \quad f = \arg \min_j (\mathcal{Q}_j(\mathbf{x}_i)). \quad (14)$$

Here, the index j is taken as a loop over the Delaunay triangles $f_j \in \text{Del}|_{\Sigma}(X)$ incident to the node \mathbf{x}_i . The scalar step length $\Delta_f \in \mathbb{R}^+$ is computed via a line search along the gradient ascent vector \mathbf{v}_f and, in this study, is taken as the first value that leads to an improvement in the worst-case incident-quality metric $\mathcal{Q}_f(\mathbf{x}_i)$. This length is computed by a simple bisection strategy, iteratively testing $\Delta_f^p = \left(\frac{1}{2}\right)^p \alpha$ until a successful nodal adjustment is found. Here p denotes the local line-search iteration, while the scalar length $\alpha \in \mathbb{R}^+$ is determined as a solution to the first-order Taylor expansion

$$\mathcal{Q}_f(\mathbf{x}_i) + \alpha \hat{\mathbf{v}}_f \cdot \frac{\partial}{\partial \mathbf{x}} (\mathcal{Q}_f(\mathbf{x}_i)) \leq \tilde{\mathcal{Q}}_j(\mathbf{x}_i). \quad (15)$$

The quantity $\tilde{\mathcal{Q}}_j(\mathbf{x}_i)$ is the *second-lowest* grid-quality score associated with the set of triangles $f_j \in \text{Del}|_{\Sigma}(X)$ incident to the central node \mathbf{x}_i . This selection strategy (Freitag and Ollivier-Gooch, 1997) is designed to compute an initial displacement α that will improve the worst element in the set until its quality is equal that of its next best neighbour. Noting that such an expansion is only first-order accurate, the

Algorithm 2 Spring and gradient-based node-smoothing

```

1: function NODESMOOTH( $\mathbf{x}$ , Del $_{|\Sigma}(X)$ ,  $\bar{Q}$ )
2:   if ( $\min(Q(\mathbf{x})) \geq \bar{Q}$ ) then
3:     Call  $\{\hat{\mathbf{v}}, \Delta\} \leftarrow \text{SPRINGVEC}(\mathbf{x}, \text{Del}_{|\Sigma}(X))$ 
       to form the spring-based search vector and
       step-length.
4:   else
5:     Call  $\{\hat{\mathbf{v}}, \Delta\} \leftarrow \text{ASCENTVEC}(\mathbf{x}, \text{Del}_{|\Sigma}(X))$ 
       to form the gradient-based search vector
       and step-length.
6:   end if
7:   while ( $i \leq M$ ) do
8:     Set  $\mathbf{x}' \leftarrow \text{proj}_{|\Sigma}(\mathbf{x} + \Delta^i \hat{\mathbf{v}})$  to move the ver-
       tex  $\mathbf{x}$  along the search vector  $\hat{\mathbf{v}}$  and project
       onto the surface  $\Sigma$ .
9:     if ( $\text{BETTER}(Q(\mathbf{x}'), Q(\mathbf{x}))$ ) then
10:      Set  $\mathbf{x} \leftarrow \mathbf{x}'$  and break.
11:    end if
12:    Step-length bisection  $\Delta^{i+1} \leftarrow \frac{1}{2} \Delta^i$ .
13:  end while
14: end function

1: function SPRINGVEC( $\mathbf{x}$ , Del $_{|\Sigma}(X)$ ,  $\hat{\mathbf{v}}, \Delta$ )
2:   Scan edges adj. to  $\mathbf{x}$ , calc.  $\mathbf{x}' \leftarrow \frac{\sum w_e(\mathbf{x} + \Delta_e \mathbf{e})}{\sum w_e}$ ,
       where  $w_e \leftarrow \Delta_e^2$  and  $\Delta_e \leftarrow \frac{h(\mathbf{x}_e) - \|\mathbf{e}\|}{\|\mathbf{e}\|}$ .
3:   Set  $\mathbf{v} \leftarrow \mathbf{x}' - \mathbf{x}$ ,  $\hat{\mathbf{v}} \leftarrow \frac{\mathbf{v}}{\|\mathbf{v}\|}$  and  $\Delta \leftarrow \|\mathbf{v}\|$ .
4:   return search vector  $\hat{\mathbf{v}}$  and step-size  $\Delta$ .
5: end function

1: function ASCENTVEC( $\mathbf{x}$ , Del $_{|\Sigma}(X)$ ,  $\hat{\mathbf{v}}, \Delta$ )
2:   Scan faces adj. to  $\mathbf{x}$  and return the worst
       adj. element  $f \leftarrow \text{argmin}(Q_f(\mathbf{x}))$ .
3:   Compute the local gradient-ascent search vec-
       tor  $\mathbf{v} \leftarrow \frac{\partial}{\partial \mathbf{x}}(Q_f(\mathbf{x}))$ .
4:   Solve  $Q_f(\mathbf{x}) + \Delta \hat{\mathbf{v}} \cdot \frac{\partial}{\partial \mathbf{x}}(Q_f(\mathbf{x})) \leq \tilde{Q}_j(\mathbf{x})$  for the
       initial step-size  $\Delta$ , where  $\tilde{Q}_j(\mathbf{x})$  is the quality
       of the second-worst triangle adj. to  $\mathbf{x}$ .
5:   return search vector  $\hat{\mathbf{v}}$  and step-size  $\Delta$ .
6: end function

1: function BETTER( $Q', Q$ )
2:   Sort the mesh-quality vectors  $Q'$  and  $Q$ .
3:   return FALSE if any  $Q'_i < Q_i$ , else TRUE.
4: end function

```

step length is iteratively decreased using the bisection strategy described previously. A limited line search is employed here, testing iterations $p \leq 5$ until a successful step is found. Consistent with the spring-based procedure described previously, a geometry-projection operator is implicitly incorporated within each update (Eq. 13), ensuring that nodes remain constrained to the spheroidal surface.

4.3 Topological “flips”

In addition to purely geometrical operations, grid optimisation also requires that adjustments be made to the underlying mesh topology, such that the surface triangulation remains a valid Delaunay structure. While it is possible to simply recompute the full Delaunay tessellation after each adjustment, such an approach would impose significant computational costs, especially when considering that a majority of updates involve small perturbations. In this work, an alternative strategy is pursued, based on local element-wise transformations, known as topological *flips*.

For any given pair of adjacent surface triangles $\{f_i, f_j\} \in \text{Del}_{|\Sigma}(X)$, a local re-triangulation can be achieved by *flipping* the local connectivity about the shared edge $\{x_i, x_j\}$, forming a new edge between the opposing vertices $\{x_a, x_b\}$. Such an operation results in the deletion of the existing triangles $\{f_i, f_j\}$ and the creation of a new pair $\{f'_i, f'_j\}$. This operation is illustrated in Fig. 6a. In the present study, the it-

erative application of such *edge-flipping* operations is used to adjust the topology of the surface triangulation, ensuring that it remains Delaunay throughout the optimisation phase. Specifically, given a general, possibly non-Delaunay, surface triangulation $\text{Tri}_{|\Sigma}(X)$, a cascade of edge flips are used to reach a valid restricted Delaunay surface tessellation $\text{Del}_{|\Sigma}(X)$. For each adjacent triangle pair $\{f_i, f_j\} \in \text{Tri}_{|\Sigma}(X)$ an edge flip is undertaken if a local violation of the Delaunay criterion is detected. New elements created by successful edge flips are iteratively re-examined until no further modifications are necessary. This approach follows the standard flip-based algorithms described in, for instance, Lawson (1977); Cheng et al. (2013).

Given a triangle $f_i \in \text{Tri}_{|\Sigma}(X)$, the local Delaunay criterion is violated if there exists a node $x_q \notin f_i$ interior to the associated circumscribing ball. In this work, such violations are detected by considering the restricted circumballs $B(\mathbf{c}_i, r_i)$ associated with each triangle $f_i \in \text{Tri}_{|\Sigma}(X)$, where the centre of the balls \mathbf{c}_i are a projection of the *planar* element circumcentres onto the spheroidal surface Σ . Such constructions are designed to account for the curvature of the underlying surface. Given an adjacent triangle pair $\{f_i, f_j\} \in \text{Tri}_{|\Sigma}(X)$ an edge flip is undertaken if either opposing vertex x_a, x_b lies within the circumball associated with the adjacent triangle. To prevent issues associated with exact floating-point comparisons, a small relative tolerance

is incorporated. Specifically, nodes are required to penetrate the opposing circumball by a distance ϵ before an edge flip is undertaken, with $\epsilon = \frac{1}{2}(r_i + r_j)\bar{\epsilon}$ and $\bar{\epsilon} = 1 \times 10^{-10}$ in the current double-precision implementation.

4.4 Edge contraction

In some cases, grid-quality and mesh-spacing conformance can be improved through the application of *edge-contraction* operations, whereby nodes are removed from the grid by collapsing certain edges. Given an edge $e_k \in \text{Tri}|_{\Sigma}(X)$, a re-triangulation of the local *cavity* $\mathcal{C}_i \subseteq \text{Tri}|_{\Sigma}(X)$, formed by the set of triangles incident to the nodes $\{x_i, x_j\} \in e_k$, can be achieved by *merging* the nodes x_i and x_j at some local mean position. In addition to collapsing the edge e_k , edge contraction also removes the two surface triangles $\{f_i, f_j\} \in \text{Tri}|_{\Sigma}(X)$ adjacent to e_k , resulting in a re-triangulation of the local cavity $\mathcal{C}'_i \subseteq \text{Tri}|_{\Sigma}(X)$. See Fig. 6b for illustration. In the present work, nodes are merged to a mean position \mathbf{x}_n – taken as the average of adjacent element circumcentres, such that $\mathbf{x}_n = \frac{1}{|\mathcal{C}_i|} \sum c_j$. Here, the c_j 's are the centres of the circumballs associated with the triangles $f_j \in \mathcal{C}_i$. As per previous discussions, the mean position \mathbf{x}_n is projected onto the spheroidal surface Σ . While such an approach is slightly more computationally intensive than strategies based on collapsing to edge midpoints, this weighted circumcentre-based scheme was found to be substantially more effective in practice. Consistent with the hill-climbing philosophy pursued throughout this study, edge contraction operations are only utilised if there is sufficient improvement in local mesh-quality metrics. Specifically, edge contraction is undertaken if the quality of the new cavity \mathcal{C}'_i exceeds that of the initial configuration \mathcal{C}_i in a worst-first fashion.

4.5 Edge refinement

Fulfilling the opposite role to edge contraction, *edge-refinement* operations seek to improve the grid through the addition of new nodes and elements. In the present study, a simplified refinement operation is utilised, in which a given edge $e_k \in \text{Tri}|_{\Sigma}(X)$ is refined by placing a new node x_n at the centre of the restricted circumball $B(c_i, r_i)$ associated with the lower-quality adjacent triangle $f_i \in \text{Tri}|_{\Sigma}(X)$. Insertion of the new node x_n induces a re-triangulation of the local cavity $\mathcal{C}_i \in \text{Tri}|_{\Sigma}(X)$ – constructed by expanding about x_n in a local greedy fashion. Starting from the initial cavity $\mathcal{C}_i = \{f_i, f_j\}$, where $\{f_i, f_j\} \in \text{Tri}|_{\Sigma}(X)$ are the triangles adjacent to the edge e_k , additional elements are added in a *breadth-first* manner, with a new, unvisited neighbour f_k added to the cavity \mathcal{C}_i if doing so will improve the worst-case element-quality metric. The final cavity \mathcal{C}_i is therefore a locally optimal configuration. In practice, this iterative deepening of \mathcal{C}_i typically converges in one or two iterations. See Fig. 6c for illustration. As per the edge-contraction and node-smoothing operations described previously, edge refinement

is implemented according to a hill-climbing-type philosophy, with operations successful only if there is sufficient improvement in local grid quality. Consistent with previous discussions, a lexicographical comparison of the grid-quality metrics associated with elements in the initial and final states \mathcal{C}_i and \mathcal{C}'_i is used to determine success.

4.6 Optimisation schedule

The full grid-optimisation procedure is realised as a combination of the various geometrical and topological operations described previously, organised into a particular iterative optimisation *schedule*. See Algorithm 3 for full details. Each outer iteration consists of a fixed set of operations: four node-smoothing passes, a single pass of edge refinement/contraction operations, and, finally, iterative edge flipping to restore the Delaunay criterion. In this study, 16 outer iterations are employed. Each node-smoothing pass is a composite operation, with the spring-based technique used to adjust nodes adjacent to high-quality elements, and the gradient-ascent method used otherwise. Specifically, spring-based smoothing is used to adjust nodes adjacent to elements with a minimum-quality score of $\bar{Q}_f \geq 0.9375$. Such thresholding ensures that the expensive gradient-ascent-type iteration is reserved for the worst elements in the grid. The optimisation schedule employed here is not based on any rigorous theoretical derivation, but is simply a set of heuristic choices that have proven to be effective in practice. The application of multiple node-smoothing passes within an outer iteration containing subsequent topological, contraction and refinement-type operations is consistent with the methodologies employed in, for instance, Freitag and Ollivier-Gooch (1997); Klingner and Shewchuk (2008).

5 Results and discussions

The performance of the Frontal-Delaunay refinement and hill-climbing optimisation algorithms presented in Sects. 3 and 4 was investigated experimentally, with the methods used to mesh a series of benchmark problems. The algorithm was implemented in C++ and compiled as a 64 bit executable. The full algorithm has been implemented as a specialised variant of the general-purpose JIGSAW meshing package, denoted JIGSAW-GEO, and is currently available online (Engwirda, 2017) or by request from the author. All tests were completed on a Linux platform using a single core of an Intel i7 processor. Visualisation and post-processing was completed using MATLAB.

5.1 Preliminaries

A set of three benchmark problems were considered. The UNIFORM-SPHERE test case is a fixed resolution meshing problem on the sphere, suitable for uniformly resolved atmospheric and/or oceanic simulations. The REGIONAL-

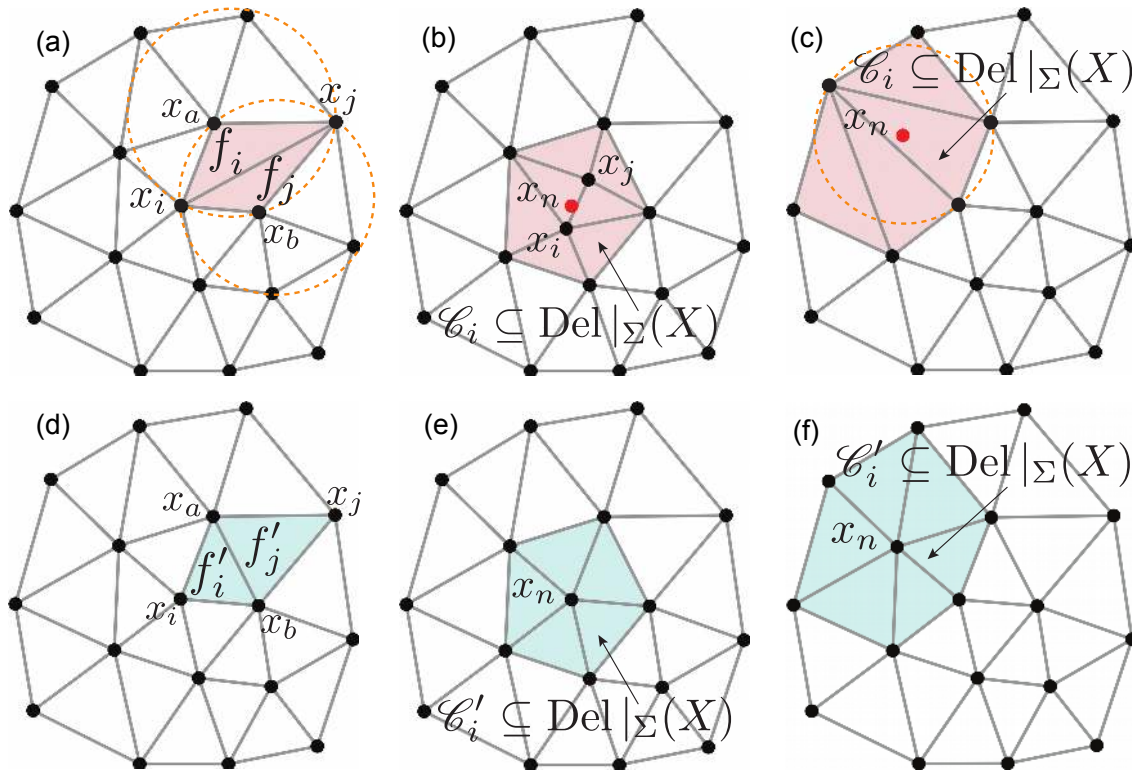


Figure 6. Topological operations for grid optimisation, showing (a, d) an edge flip, (b, e) an edge contraction, and (c, f) an edge-refinement operation. Grid configurations before and after each flip are shown in the upper and lower panels, respectively.

ATLANTIC test case explores simple, regional refinement for ocean modelling, embedding a high-resolution, eddy-permitting representation of the North Atlantic ocean basin in a coarse global grid. The SOUTHERN-OCEAN test case is designed to test the multi-resolution capabilities of the algorithm, defining a complex non-uniform grid-spacing function for the Southern Ocean and Antarctic regions. The mesh-spacing constraints for this problem were designed to incorporate a combination of topographic and regional refinement. The Voronoi–Delaunay grids for these test cases are shown in Figs. 8, 10, 13, and 15, with the associated grid-quality statistics presented in Figs. 9, 11, and 14. The underlying grid-spacing definitions used for the REGIONAL-ATLANTIC and SOUTHERN-OCEAN problems are shown in Fig. 7.

In all test cases, limiting radius-edge ratios were specified, with $\bar{\rho}_f = 1.05$. This constraint ensures that the minimum angles in the Delaunay triangulations are bounded, with $\theta_{\min} \geq 28.4^\circ$. For all test problems, detailed statistics on element quality are presented, including histograms of element area-length ratios a_f , element angles θ_f , and relative-edge lengths h_r . These quality metrics are described in detail in Sect. 3.4. Histograms highlight the minimum, maximum, and mean values of the relevant distributions as appropriate.

5.2 Uniform global grid

The performance of the JIGSAW-GEO algorithm was first assessed using the UNIFORM-SPHERE test case, seeking to build a uniformly resolved global grid for general circulation modelling. Spatially uniform mesh-size constraints were enforced, setting $\bar{h}(x) = 150$ km over the full sphere. The resulting grid is shown in Fig. 8 and contains 83 072 Delaunay triangles and 41 538 Voronoi cells. Grid-quality metrics are presented in Fig. 9, showing distributions before and after the application of the grid-optimisation procedure.

Visual inspection of these grids show that very high levels of geometric quality are achieved, with both the Delaunay triangulation and Voronoi tessellation consisting of highly regular elements, free of obvious distortion and/or areas of over- or under-refinement. Focusing on the element shape-quality statistics, it is noted that very high levels of mesh regularity are achieved, with the vast majority of element area-length scores tightly clustered about $a_f = 1$. Similarly, the distribution of element angles shows strong convergence around $\theta_f = 60^\circ$, revealing most triangles to be near equilateral. Finally, analysis of the relative-length distribution shows that edge lengths follow the imposed mesh-spacing constraints closely, with very tight clustering about $h_r = 1$.

The effect of the grid-optimisation procedure can be assessed by comparing the mesh-quality statistics presented in

Algorithm 3 Hill-climbing grid optimisation

```

1: function OPTIMISE( $\Sigma, \bar{h}(\mathbf{x}), \text{Del}|\Sigma(X)$ )
2:   for ( $i_{\text{outer}} \leq N_{\text{outer}}$ ) do
3:     for ( $i_{\text{inner}} \leq N_{\text{inner}}$ ) do
4:       for ( $\forall$  nodes  $j \in X$ ) do
5:         Call NODESMOOTH( $\mathbf{x}_j, \text{Del}|\Sigma(X)$ )
           to improve geometric distribution of
           triangle vertices.
6:       end for
7:     end for
8:     Call MERGENODE( $\text{Del}|\Sigma(X)$ ) to collapse
           any short edges.
9:     Call SPLITEDGE( $\text{Del}|\Sigma(X)$ ) to refine any
           long edges.
10:    Call EDGEFLIPS( $\text{Del}|\Sigma(X)$ ) to restore grid
           Delaunay-ness (i.e. local-orthogonality).
11:  end for
12:  return surface triangulation object  $\text{Del}|\Sigma(X)$ .
13: end function
1: function EDGEFLIPS( $X, \text{Del}|\Sigma(X)$ )
2:   for ( $\forall$  edges  $e \in \text{Del}|\Sigma(X)$ ) do
3:     Given the adj. triangle-pair  $\mathcal{C} \leftarrow \{f_i, f_j\}$ ,
           flip the common diagonal to re-triangulate
           the cavity  $\mathcal{C}'$ .
4:     if BETTER( $Q(\mathcal{C}'), Q(\mathcal{C})$ ),  $\text{Del}|\Sigma(X) \leftarrow \mathcal{C}'$ .
5:   end for
6: end function
1: function MERGENODE( $X, \text{Del}|\Sigma(X)$ )
2:   for ( $\forall$  edges  $e \in \text{Del}|\Sigma(X)$ ) do
3:     Form a local merge point  $\mathbf{x}'$  about the
           edge  $e$ . Given  $\mathbf{x}_i, \mathbf{x}_j \in e$ , set  $\mathbf{x}' \leftarrow$ 
            $\text{proj}|\Sigma(\frac{1}{|\mathcal{C}|} \sum \mathbf{c}_f)$ , where  $\mathcal{C}$  is the set of tri-
           angles adj. to  $\mathbf{x}_i, \mathbf{x}_j$  and  $\mathbf{c}_f$  are the associ-
           ated circumcentres.
4:     Merge the vertices  $\mathbf{x}_i, \mathbf{x}_j \leftarrow \mathbf{x}'$ , and re-
           triangulate the cavity  $\mathcal{C}'$  as a result.
5:     if BETTER( $Q(\mathcal{C}'), Q(\mathcal{C})$ ),  $\text{Del}|\Sigma(X) \leftarrow \mathcal{C}'$ .
6:   end for
7:   return updated  $X$  and  $\text{Del}|\Sigma(X)$ .
8: end function
1: function SPLITEDGE( $X, \text{Del}|\Sigma(X)$ )
2:   for ( $\forall$  edges  $e \in \text{Del}|\Sigma(X)$ ) do
3:     Form a local refinement point  $\mathbf{x}'$  about the
           edge  $e$ . Given the two adj. triangles,  $\mathcal{C} \leftarrow$ 
            $\{f_i, f_j\}$ , set  $\mathbf{x}'$  equal to the circumcentre of the
           lower quality element.
4:     Append the new vertex  $X \leftarrow \mathbf{x}'$ , and re-
           triangulate the cavity  $\mathcal{C}'$  as a result.
5:     if BETTER( $Q(\mathcal{C}'), Q(\mathcal{C})$ ),  $\text{Del}|\Sigma(X) \leftarrow \mathcal{C}'$ .
6:   end for
7:   return updated  $X$  and  $\text{Del}|\Sigma(X)$ .
8: end function

```

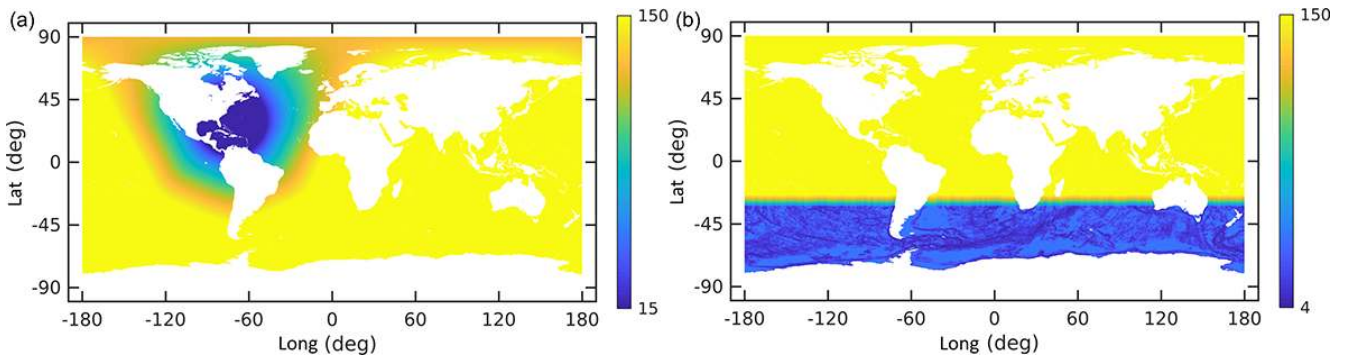


Figure 7. Mesh-spacing functions $\bar{h}(\mathbf{x})$ for the regionally refined North Atlantic and topographically refined Southern Ocean grids. Mesh spacing is shown in kilometres.

Fig. 9. The application of mesh optimisation is seen to be most pronounced at the *tails* of the distributions, showing that, as expected, the hill-climbing-type procedure is effective at improving the worst elements in the grid. Specifically, the minimum area-length metric is improved from $a_f = 0.67$ to $a_f = 0.94$, and the distribution of element-wise angles is

narrowed from $31^\circ \leq \theta_f \leq 112^\circ$ to $44^\circ \leq \theta_f \leq 78^\circ$. A slight broadening of the mean parts of the distributions is also evident, showing that in some cases, higher-quality elements are slightly compromised to facilitate improvements to their lower-quality neighbours. This behaviour is consistent with the worst-first philosophy employed in this study.

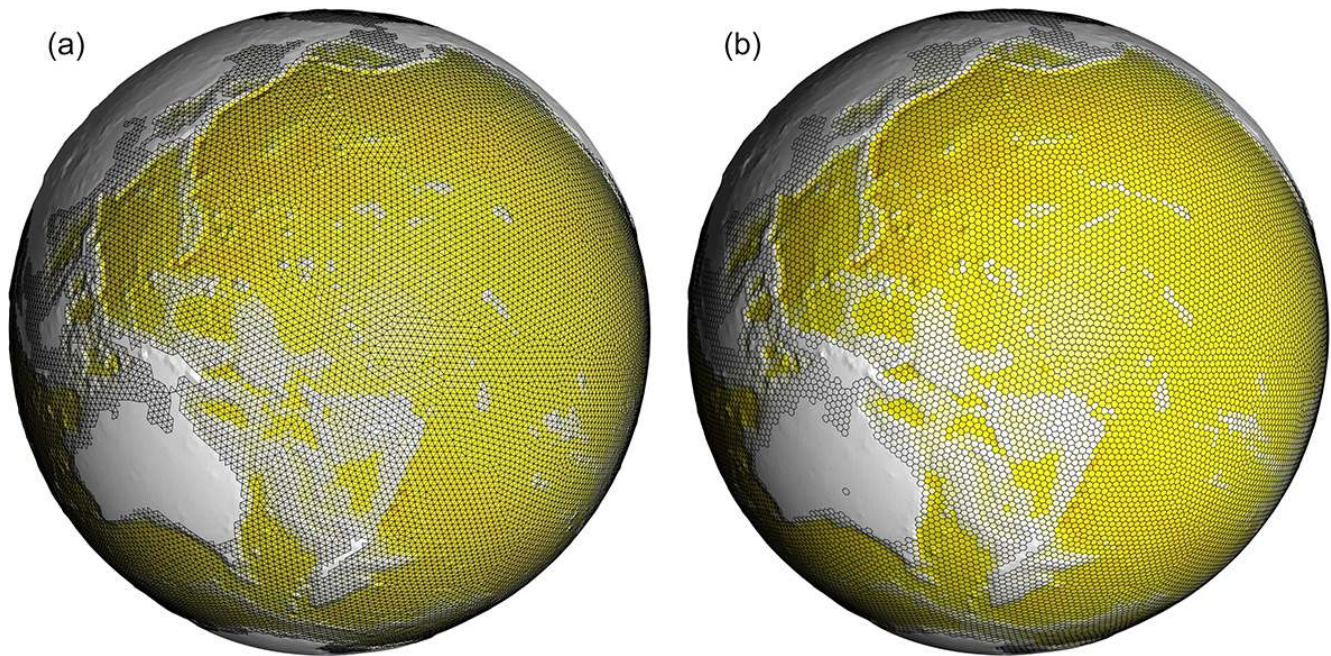


Figure 8. A uniform resolution global grid, showing (a) the underlying spheroidal Delaunay triangulation, and (b) the associated staggered Voronoi dual; 150 km grid spacing was specified globally. Topography is drawn using an exaggerated scale, with elevation from the reference geoid amplified by a factor of 10.

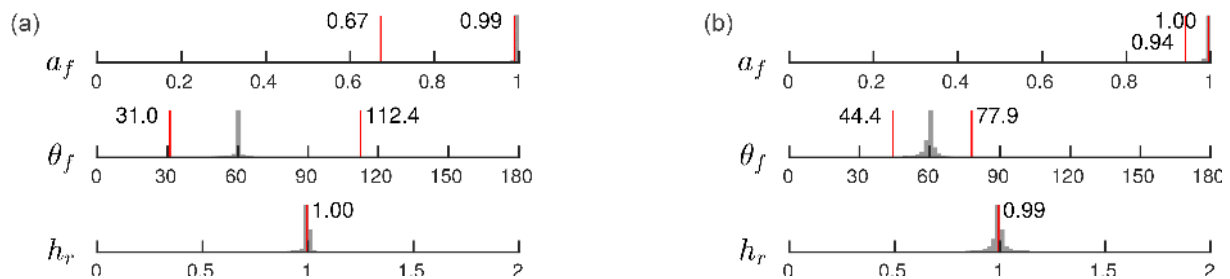


Figure 9. Mesh-quality metrics for the uniform resolution grid, before (a) and after (b) the application of mesh optimisation. Histograms of area-length ratio a_f , enclosed-angle θ_f and relative-length ratio h_r are illustrated, with min., max., and mean values annotated.

Beyond improvements to grid-quality statistics, the impact of mesh optimisation can be further understood by considering the well centredness of the resulting staggered Voronoi–Delaunay grid. Well-centred triangulations are those for which all element circumcentres are located within their parent triangles, ensuring that the associated Voronoi cells are nicely staggered with respect to the underlying triangulation as a result. Such a constraint is equivalent to requiring that all Delaunay triangles are acute, with $\theta_f \leq 90^\circ$. Further details regarding these constraints are presented in Sect. 2.

Well-centred grids are highly desirable from a numerical perspective, allowing for, for instance, the unstructured C-grid scheme employed in the MPAS framework to achieve optimal rates of convergence. Specifically, when a grid is well centred, it is guaranteed that associated edges in the staggered Voronoi–Delaunay cells intersect, ensuring that

evaluation of the cell transport and circulation terms can be accurately computed using compact numerical stencils. In the case of *perfectly centred* grids, such intersections occur at edge midpoints – allowing for a numerical scheme based on local linear interpolants to achieve second-order accuracy.

The construction of well-centred grids is known to be a difficult problem, and the development of algorithms for their generation is an ongoing area of research (VanderZee et al., 2008; Vanderzee et al., 2010). For the uniform-resolution case studied here, it is clear that the JIGSAW-GEO algorithm is successful in generating such a grid, with all angles bounded below 77.9° .

5.3 Regionally refined North Atlantic grid

The multi-resolution capabilities of the JIGSAW-GEO algorithm were investigated in the REGIONAL-ATLANTIC test

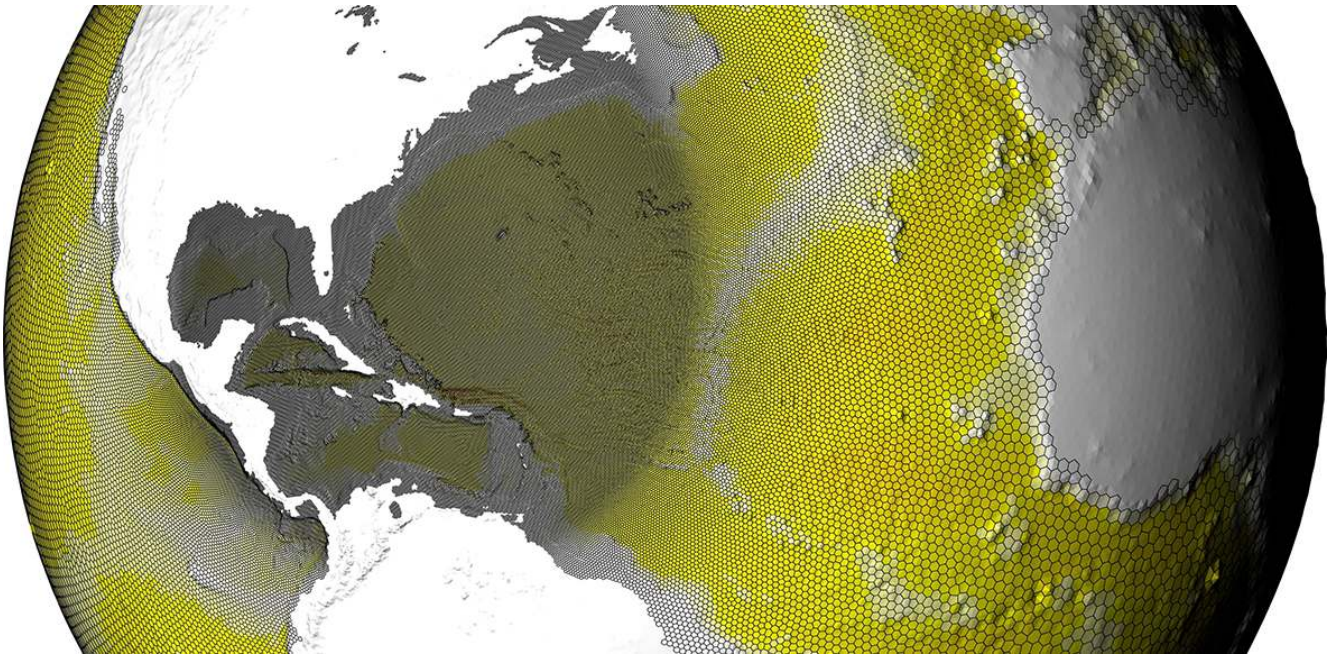


Figure 10. A regionally refined Voronoi-type grid of the North Atlantic region. Global coarse grid resolution is 150 km, with a 15 km eddy-permitting grid-spacing specified over the Atlantic ocean basin. Topography is drawn using an exaggerated scale, with elevation from the reference geoid amplified by a factor of 10.

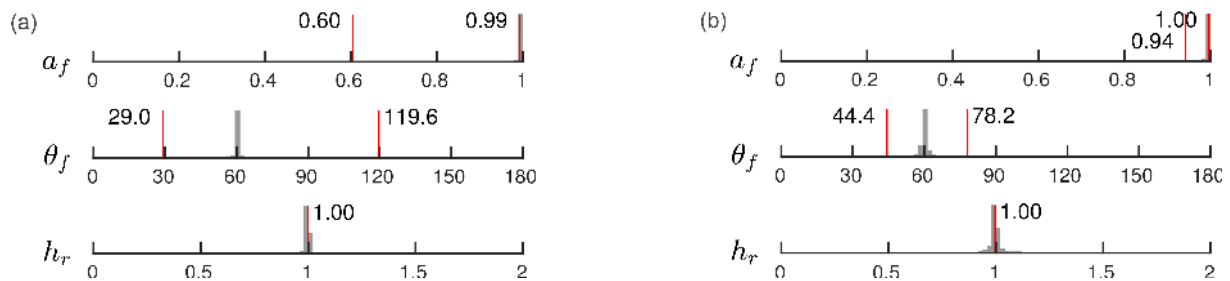


Figure 11. Mesh-quality metrics for the regionally refined grid, before (a) and after (b) the application of mesh optimisation. Histograms of area-length ratio a_f , enclosed-angle θ_f and relative-length ratio h_r are illustrated, with min., max., and mean values annotated.

case, seeking to build a regionally resolved grid for high-resolution modelling of the North Atlantic Ocean basin. Non-uniform mesh-size constraints were enforced, setting $\bar{h}(x) = 150$ km globally, with 15 km *eddy-permitting* mesh spacing specified over the North Atlantic region. The resulting grid is shown in Figs. 10 and 12 and contains 358 064 Delaunay triangles and 179 081 Voronoi cells. Grid-quality metrics are presented in Fig. 11, showing distributions before and after the application of grid optimisation.

Consistent with the results presented previously, a very high-quality Voronoi–Delaunay grid was generated for the REGIONAL-ATLANTIC problem, with each grid-quality metric tightly clustered about its optimal value, such that $a_f \rightarrow 1$, $\theta_f \rightarrow 60^\circ$, and $h_r \rightarrow 1$. As per the uniform resolution case, the application of mesh optimisation led to significant improvements in worst-case grid quality, with the

minimum area-length metric improved from $a_f = 0.60$ to $a_f = 0.94$, and the distribution of element angles narrowed from $29^\circ \leq \theta_f \leq 120^\circ$ to $44^\circ \leq \theta_f \leq 78^\circ$. The resulting optimised grid is also clearly well centred, with all angles in the Delaunay triangulation less than 78.2° . Overall, grid quality can be seen to achieve essentially the same levels of optimality as the uniform resolution test case, showing that the JIGSAW-GEO algorithm can be used to generate very high-quality spatially adaptive grids without obvious degradation in mesh quality.

5.4 Multi-resolution Southern Ocean grid

The JIGSAW-GEO algorithm was then used to mesh the challenging SOUTHERN-OCEAN test case, allowing its performance for large-scale problems involving rapidly vary-

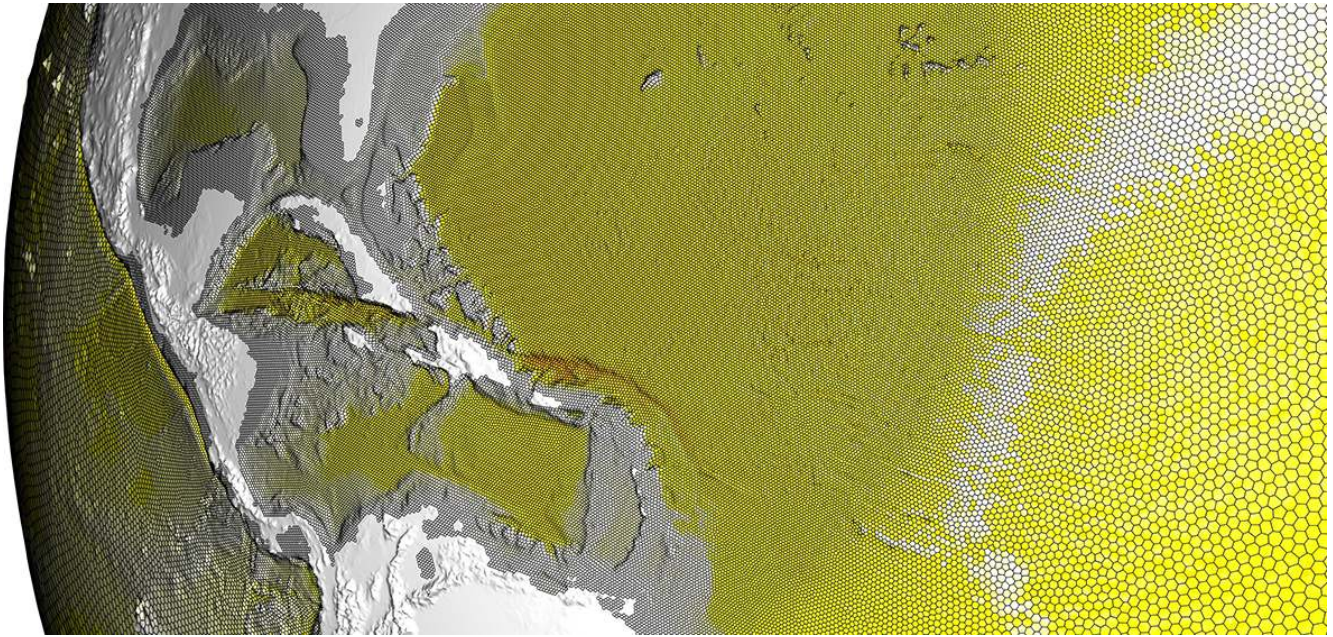


Figure 12. Additional detail of the regionally refined North Atlantic grid shown in Fig. 10.

ing mesh-spacing constraints to be analysed in detail. This test case seeks to build a multi-resolution grid for regionally refined ocean studies, with a particular focus on resolution of the Antarctic Circumpolar Current, and adjacent Antarctic processes. Composite mesh-spacing constraints were enforced, consisting of a coarse global background resolution of 150 km, with an eddy-permitting 15 km grid spacing specified south of 32.5° S. An additional topographic adaptation was also utilised within the southern annulus, with grid resolution increased in regions of large bathymetric gradient. A minimum grid spacing of 4 km was specified. Topographic gradients were computed using the high-resolution ETOPO1 Global Relief dataset (Amante and Eakins, 2009). The resulting grid is shown in Fig. 13, with additional detail shown in Fig. 15. The grid contains 3 119 849 Delaunay triangles and 1 559 927 Voronoi cells. Associated grid-quality metrics are presented in Fig. 14, showing distributions before and after the application of the grid-optimisation procedure.

Consistent with previous results, visual inspection of Figs. 13 and 15 confirm that the JIGSAW-GEO algorithm is capable of generating very high-quality multi-resolution grids, containing a majority of near-perfect Delaunay triangles and Voronoi cells. Additionally, it can be seen that grid resolution varies smoothly, even in regions of rapidly fluctuating mesh-spacing constraints (see Fig. 15 for detail). Analysis of the grid-quality metrics shown in Fig. 14 show that very high levels of mesh regularity are achieved, with element area-length scores tightly clustered about $a_f = 1$ and element angles showing strong convergence around $\theta_f = 60^\circ$. Interestingly, despite the complexity of the imposed mesh-spacing function, the relative-length distribution

still shows relatively tight conformance, with a sharp clustering about $h_r = 1$. Overall, mean grid quality is slightly reduced compared to the uniform resolution case, illustrated by a slight broadening of the grid-quality distributions. Note that such behaviour is expected in the multi-resolution case, with imperfect triangle geometries required to satisfy non-uniform mesh-spacing constraints.

The effect of the grid-optimisation procedure can be assessed by comparing the mesh-quality statistics presented in Fig. 14. As per the uniform resolution case, mesh optimisation appears to be most aggressive at the tails of the distributions, acting to improve the worst elements in the grid. The minimum area-length metric is improved from $a_f = 0.59$ to $a_f = 0.90$, and the distribution of element angles is narrowed from $28^\circ \leq \theta_f \leq 121^\circ$ to $40^\circ \leq \theta_f \leq 80^\circ$. Consistent with previous results, a moderate broadening of the mean components of the distributions can be observed, especially in the enclosed-angle and relative-length metrics. As per previous discussions, this behaviour shows that some high-quality elements are slightly compromised to improve their lower-quality neighbours.

The optimised grid is also clearly well centred, with all angles in the Delaunay triangulation bounded below 80° . This result shows only marginal degradation compared to the simpler examples presented previously – demonstrating the effectiveness of the optimisation strategies described here. These results show that very high-quality, well-centred grids can be generated for complex multi-resolution cases, even when the grid-spacing function incorporates strong local fluctuation. Nonetheless, the construction of well-centred grids remains a challenging task, and it is noted that there

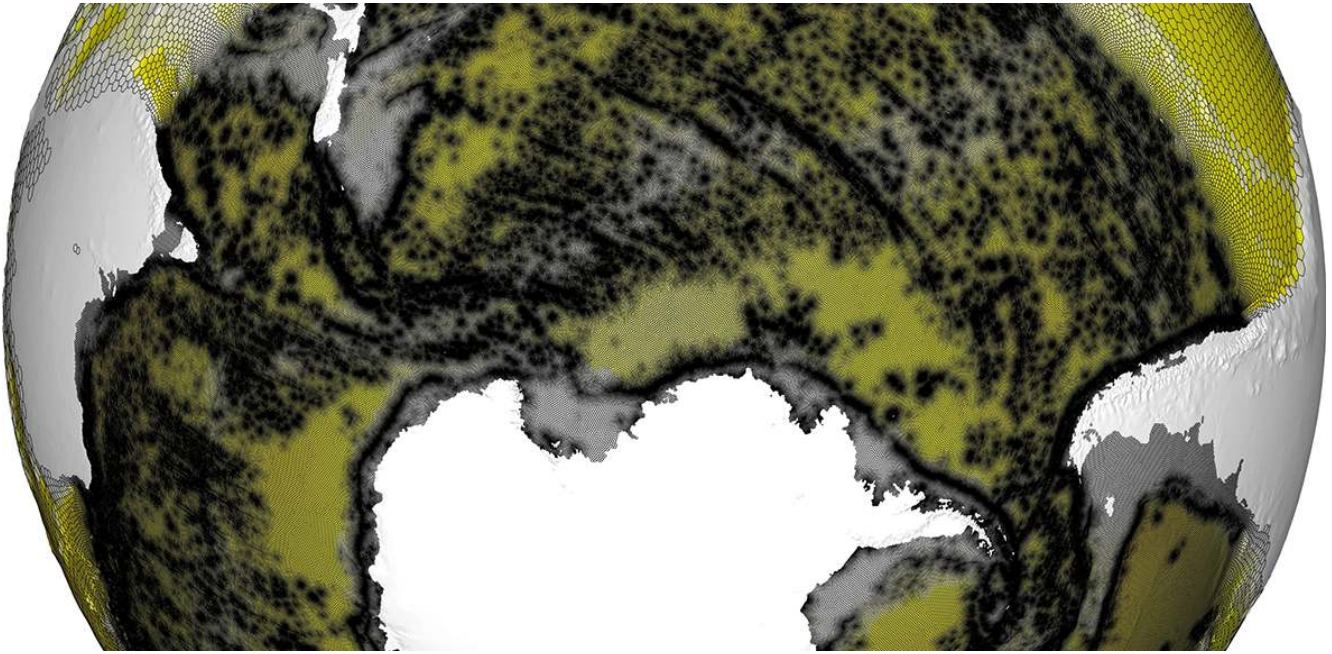


Figure 13. A multi-scale grid for the Southern Ocean. Coarse resolution is 150 km, with 15 km eddy-permitting grid spacing specified south of 32.5° S. Local topographic adaptation is also utilised, with resolution increased in areas of large bathymetric gradient. Minimum grid spacing is 4 km. Topography is drawn using an exaggerated scale, with elevation from the reference geoid amplified by a factor of 10.

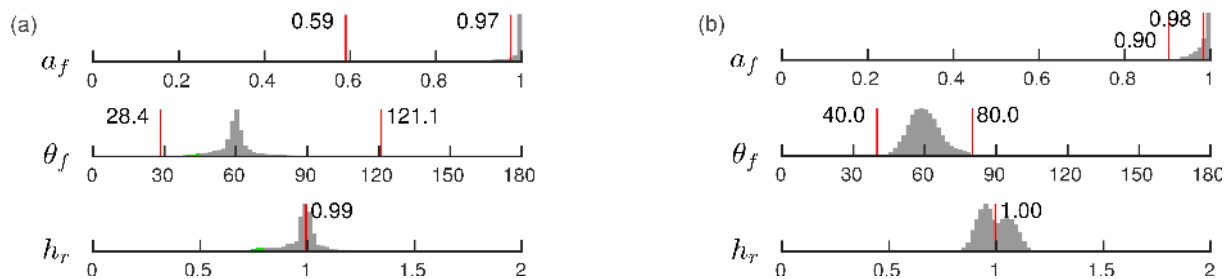


Figure 14. Mesh-quality metrics associated with the multi-resolution grid, before (a) and after (b) the application of mesh optimisation. Histograms of area-length ratio a_f , enclosed-angle θ_f and relative-length ratio h_r are illustrated, with min., max., and mean values annotated.

may exist test cases that defeat the current strategy. Therefore, the pursuit of alternative optimisation techniques, designed to target grid well-centredness directly, is an interesting avenue for future research.

5.5 Computational performance

In addition to the generation of very high-quality grids, the new JIGSAW-GEO algorithm also imposes a relatively moderate computational burden, producing large-scale, multi-resolution grids in a matter of minutes using standard desktop-based computing infrastructure. Specifically, grid-generation for the UNIFORM-SPHERE, REGIONAL-ATLANTIC, and SOUTHERN-OCEAN test cases required approximately 12 s, 1½ min, and 10 min of computation time, respectively, running on a single core of an Intel i7 processor. In all cases, grid optimisation was found to be approxi-

mately 4 times as expensive as the initial Frontal-Delaunay refinement. Compared to the existing iterative MPI-SCVT algorithm (Jacobsen et al., 2013), commonly used to generate grids for the MPAS framework, these results represent a significant increase in productivity, with the MPI-SCVT algorithm often requiring days, or even weeks of distributed computing time.

Additionally, practical experience with the MPI-SCVT algorithm has shown that it cannot always be relied upon to generate an appropriate grid, irrespective of the amount of computational time allowed for convergence to be reached. While always generating locally orthogonal and centroidal tessellations with very high mean grid qualities, the MPI-SCVT algorithm does not provide bounds on the worst-case metrics. In practice, multi-resolution grids generated using the MPI-SCVT algorithm are often observed to contain a mi-

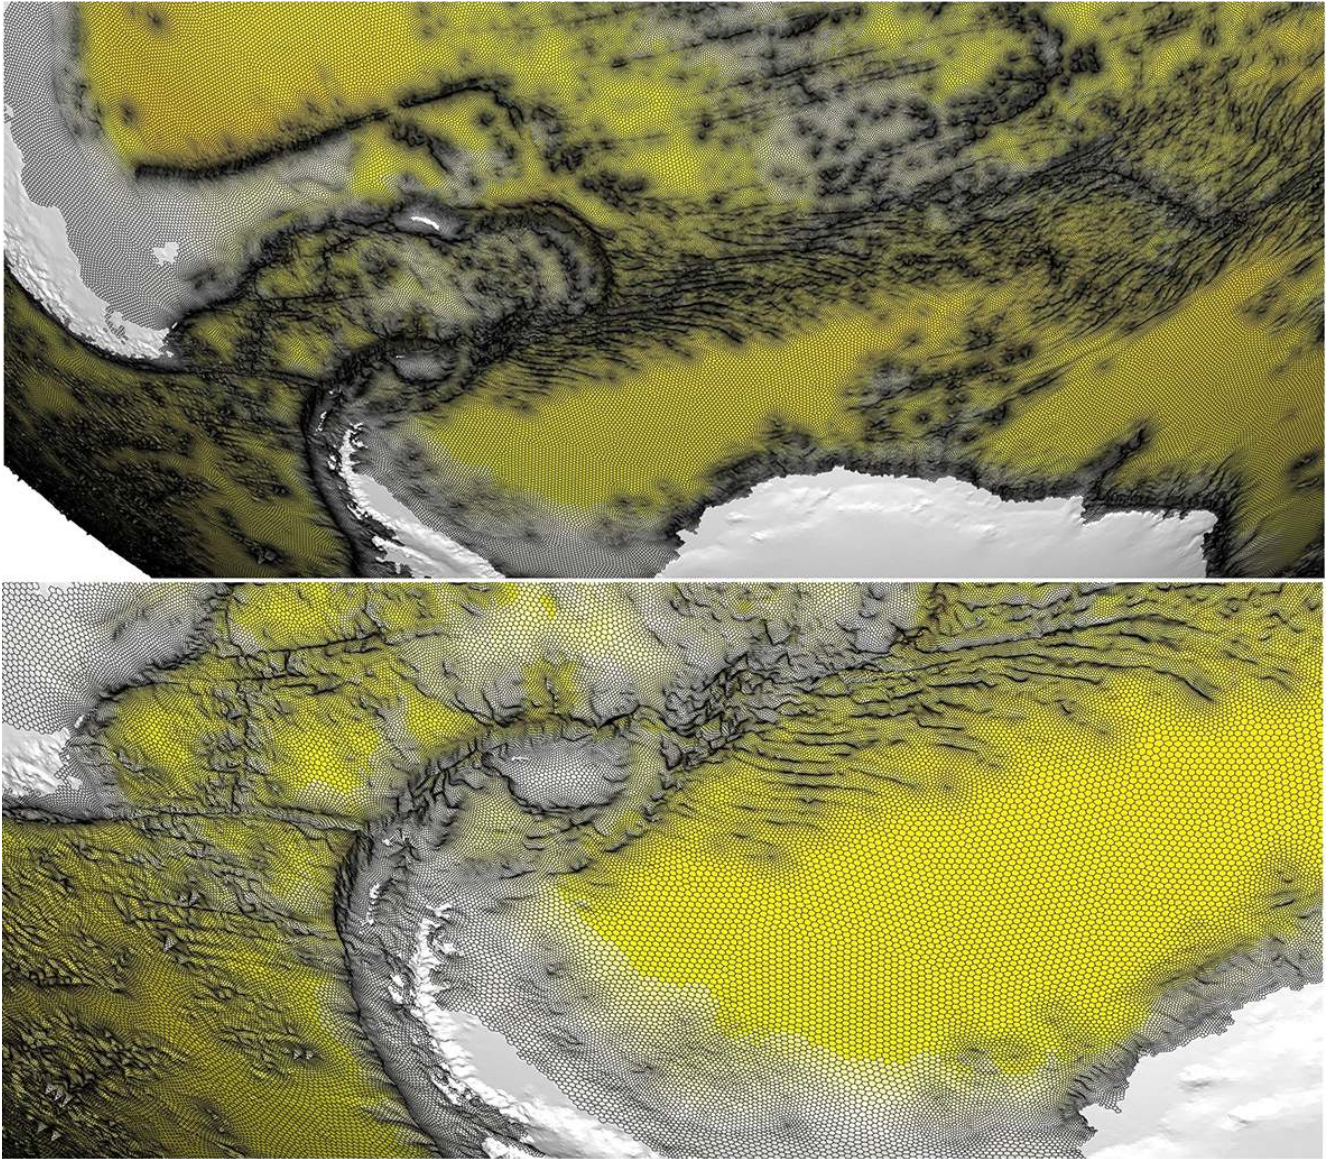


Figure 15. Additional detail of the multi-resolution Southern Ocean grid shown in Fig. 13.

nority of obtuse triangles that violate the well-centred constraint. As per the discussions presented in Sect. 2, such grids are inappropriate for use in an unstructured C-grid model such as MPAS. Currently, grid generation for such models often requires a degree of user-driven *trial and error* as a result, making grid generation a somewhat arduous task for model users. Initial experiments conducted using the JIGSAW-GEO algorithm have shown it to be a useful alternative, reliably generating valid well-centred multi-resolution grids for a wide range of user-defined constraints and configuration settings. Initial evaluation of JIGSAW-GEO for use with the MPAS framework is currently underway.

6 Conclusions and future work

A new algorithm for the generation of multi-resolution staggered unstructured grids for large-scale general circulation modelling on the sphere has been described. Using a combination of Frontal-Delaunay refinement and hill-climbing-type optimisation techniques, it has been shown that very high-quality locally orthogonal, centroidal, and well-centred spheroidal grids appropriate for unstructured C-grid-type general circulation models can be generated. The performance of this new approach has been verified using a number of multi-scale global benchmarks, including difficult problems incorporating highly non-uniform mesh-spacing constraints.

This new algorithm is available as part of the JIGSAW meshing package, providing a simple and easy-to-use tool for the oceanic and atmospheric modelling communities. A number of benchmark problems have been analysed, examining the performance of the new approach. The Frontal-Delaunay refinement algorithm has been shown to generate *guaranteed-quality* spheroidal Delaunay triangulations – satisfying worst-case bounds on element angles and exhibiting smooth grading characteristics. This algorithm has been shown to produce very high-quality multi-resolution triangulations, with a majority of elements exhibiting strong conformance to element-shape and grid-spacing constraints.

The use of a coupled geometrical and topological hill-climbing optimisation procedure was shown to further improve grid quality, especially for the lowest quality elements in each mesh. It was demonstrated that these optimisation techniques allow for grid quality to be improved to the extent that fully well-centred mesh configurations can be achieved, with angles in the Delaunay triangulation bounded below 90° . For the three global test cases presented here, angles were bounded above $\theta_f \geq 40^\circ$ and below $\theta_f \leq 80^\circ$.

The construction of meshes appropriate for a range of contemporary unstructured C-grid-type general circulation mod-

els was also discussed in detail, with a focus on the generation of multi-resolution grids for the MPAS framework. The availability of this new algorithm is expected to significantly reduce the grid-generation burden for MPAS model users.

Future work will focus on a generalisation of the algorithm and improvements to its efficiency, including (i) support for coastal constraints, (ii) improvements to computational performance through parallelism, and (iii) further enhancements to the mesh-optimisation procedures, with a focus on improving the well centredness of the resulting staggered grids. The investigation of *solution-adaptive* multi-scale representations, in which grid resolution is adapted to spatial variability in model state (Sein et al., 2016), is also an obvious direction for future investigation.

Code availability. The JIGSAW-GEO grid generator used in this study is available as a Zenodo archive: <https://doi.org/10.5281/zenodo.556602> (Engwirda, 2017). The JIGSAW-GEO framework is under active development and the latest version can be accessed here: <https://github.com/dengwirda/jigsaw-geo-matlab>.

Appendix A: Spheroidal predicates

Computation of the restricted Delaunay surface tessellation $\text{Del}_\Sigma(X)$ requires the evaluation of a single predicate. Given a spheroidal surface Σ , the task is to compute intersections between edges in the Voronoi tessellation $\text{Vor}(X)$ and the surface Σ .

Restricted surface triangles

Restricted surface triangles $f_i \in \text{Del}_\Sigma(X)$ are defined as those associated with an *intersecting* Voronoi edge $v_e \in \text{Vor}(X)$, where $v_e \cap \Sigma \neq \emptyset$. These triangles provide a good piecewise linear approximation to the surface Σ . For a given triangle f_i , the associated Voronoi edge v_e is defined as the line-segment joining the two *circumcentres* c_i and c_j associated with the pair of tetrahedrons that share the face f_i . The task then is to find intersections between the line-segments v_e and the surface Σ . Let p be a point on a given Voronoi edge-segment v_e

$$p = \bar{c} + t\Delta, \quad -1 \leq t \leq +1, \quad (\text{A1})$$

where

$$\bar{c} = \frac{1}{2}(c_i + c_j) \quad \text{and} \quad \Delta = \frac{1}{2}(c_j - c_i). \quad (\text{A2})$$

Substituting Eq. (A1) into the equation of the spheroidal surface Σ leads to a quadratic expression for the parameter t

$$\sum_{i=1}^3 \left(\frac{\bar{c}_i + t\Delta_i}{r_i} \right)^2 = 1, \quad (\text{A3})$$

$$\sum_{i=1}^3 \frac{\bar{c}_i^2 + 2t\bar{c}_i\Delta_i + t^2\Delta_i^2}{r_i^2} = 1, \quad (\text{A4})$$

$$\sum_{i=1}^3 \left(\frac{\Delta_i^2}{r_i^2} \right) t^2 + \left(\frac{2\bar{c}_i\Delta_i}{r_i^2} \right) t + \left(\frac{\bar{c}_i^2}{r_i^2} - 1 \right) = 0. \quad (\text{A5})$$

Any real solutions $-1 \leq t_\Sigma \leq +1$ correspond to non-trivial intersections $v_e \cap \Sigma \neq \emptyset$. The corresponding point of intersection p_Σ can be found by substituting t_Σ into Eq. (A1).

Competing interests. The authors declare that they have no conflict of interest.

Acknowledgements. This work was carried out at the NASA Goddard Institute for Space Studies, the Massachusetts Institute of Technology, and the University of Sydney with the support of a NASA–MIT cooperative agreement and an Australian Postgraduate Award. The author wishes to thank Todd Ringler, Luke Van Roekel, Mark Petersen, Matthew Hoffman, and Phillip Wolfram for their assistance on grid generation for the MPAS-O environment. John Marshall provided feedback on an earlier version of the manuscript. The author also wishes to thank the anonymous reviewers for their helpful comments and feedback.

Edited by: S. Marras

Reviewed by: two anonymous referees

References

- Adcroft, A., Campin, J.-M., Hill, C., and Marshall, J.: Implementation of an atmosphere-ocean general circulation model on the expanded spherical cube, *Mon. Weather Rev.*, 132, 2845–2863, 2004.
- Amante, C. and Eakins, B. W.: ETOPO1 1 Arc-Minute Global Relief Model: Procedures, Data Sources and Analysis, NOAA technical memorandum, NESDIS NGDC-24, <https://doi.org/10.7289/V5C8276M>, 2009.
- Arakawa, A. and Lamb, V. R.: Computational design of the basic dynamical processes of the UCLA general circulation model, *Methods in computational physics*, 17, 173–265, 1977.
- Bernard, P.-E., Chevaugnon, N., Legat, V., Deleersnijder, E., and Remacle, J.-F.: High-order h-adaptive discontinuous Galerkin methods for ocean modelling, *Ocean Dynam.*, 57, 109–121, 2007.
- Bleck, R.: An oceanic general circulation model framed in hybrid isopycnic-Cartesian coordinates, *Ocean Model.*, 4, 55–88, 2002.
- Bochev, P. B. and Hyman, J. M.: Principles of mimetic discretizations of differential operators, in: *Compatible spatial discretizations*, 89–119, Springer New York, USA, 2006.
- Boissonnat, J. D. and Oudot, S.: Provably Good Surface Sampling and Approximation, in: *ACM International Conference Proceeding Series*, 43, 9–18, 2003.
- Boissonnat, J. D. and Oudot, S.: Provably Good Sampling and Meshing of Surfaces, *Graph. Models*, 67, 405–451, 2005.
- Chen, C., Liu, H., and Beardsley, R. C.: An unstructured grid, finite-volume, three-dimensional, primitive equations ocean model: application to coastal ocean and estuaries, *J. Atmos. Ocean. Tech.*, 20, 159–186, 2003.
- Chen, C., Huang, H., Beardsley, R. C., Liu, H., Xu, Q., and Cowles, G.: A finite volume numerical approach for coastal ocean circulation studies: Comparisons with finite difference models, *J. Geophys. Res.-Oceans*, 112, 2007.
- Cheng, S. W., Dey, T. K., Ramos, E. A., and Ray, T.: Sampling and Meshing a Surface with Guaranteed Topology and Geometry, *SIAM J. Comput.*, 37, 1199–1227, 2007.
- Cheng, S. W., Dey, T. K., and Ramos, E. A.: Delaunay Refinement for Piecewise Smooth Complexes, *Discrete Comput. Geom.*, 43, 121–166, <https://doi.org/10.1007/s00454-008-9109-3>, 2010.
- Cheng, S. W., Dey, T. K., and Shewchuk, J. R.: *Delaunay Mesh Generation*, Taylor & Francis, New York, USA, 2013.
- Comblen, R., Legrand, S., Deleersnijder, E., and Legat, V.: A finite element method for solving the shallow water equations on the sphere, *Ocean Model.*, 28, 12–23, 2009.
- Conroy, C. J., Kubatko, E. J., and West, D. W.: ADMESH: An advanced, automatic unstructured mesh generator for shallow water models, *Ocean Dynam.*, 62, 1503–1517, 2012.
- Du, Q., Faber, V., and Gunzburger, M.: Centroidal Voronoi tessellations: applications and algorithms, *SIAM Rev.*, 41, 637–676, 1999.
- Edelsbrunner, H. and Shah, N. R.: Triangulating Topological Spaces, *Int. J. Comput. Geom. App.*, 7, 365–378, 1997.
- Engwirda, D.: Conforming restricted Delaunay mesh generation for piecewise smooth complexes, *Procedia Eng.*, 163, 84–96, 2016.
- Engwirda, D.: JIGSAW-GEO: Grid-generation for general circulation modelling on the sphere (Version 1.0.1), <https://doi.org/10.5281/zenodo.556602>, 2017.
- Engwirda, D. and Ivers, D.: Off-centre Steiner points for Delaunay-refinement on curved surfaces, 23rd International Meshing Roundtable Special Issue: Advances in Mesh Generation, *Comput. Aided Design*, 72, 157–171, <https://doi.org/10.1016/j.cad.2015.10.007>, 2016.
- Erten, H. and Üngör, A.: Quality Triangulations with Locally Optimal Steiner Points, *SIAM J. Sci. Comp.*, 31, 2103–2130, 2009.
- Ford, R., Pain, C., Piggott, M., Goddard, A., De Oliveira, C., and Umpleby, A.: A nonhydrostatic finite-element model for three-dimensional stratified oceanic flows. Part I: model formulation, *Mon. Weather Rev.*, 132, 2816–2831, 2004a.
- Ford, R., Pain, C., Piggott, M., Goddard, A., de Oliveira, C., and Umpleby, A.: A nonhydrostatic finite-element model for three-dimensional stratified oceanic flows. Part II: model validation, *Mon. Weather Rev.*, 132, 2832–2844, 2004b.
- Freitag, L. A. and Ollivier-Gooch, C.: Tetrahedral mesh improvement using swapping and smoothing, *Int. J. Numer. Meth. Eng.*, 40, 3979–4002, 1997.
- Frey, J. L. and George, P. L.: *Mesh Generation: Applications to Finite Elements*, Hermès, Paris, France, 2000.
- Fringer, O., Gerritsen, M., and Street, R.: An unstructured-grid, finite-volume, nonhydrostatic, parallel coastal ocean simulator, *Ocean Model.*, 14, 139–173, 2006.
- Geuzaine, C. and Remacle, J.-F.: Gmsh: A 3-D finite element mesh generator with built-in pre-and post-processing facilities, *Int. J. Numer. Meth. Eng.*, 79, 1309–1331, 2009.
- Heikes, R. and Randall, D. A.: Numerical integration of the shallow-water equations on a twisted icosahedral grid. Part I: Basic design and results of tests, *Mon. Weather Rev.*, 123, 1862–1880, 1995.
- Holleman, R., Fringer, O., and Stacey, M.: Numerical diffusion for flow-aligned unstructured grids with application to estuarine modeling, *Int. J. Numer. Meth. Fl.*, 72, 1117–1145, 2013.
- Jacobsen, D. W., Gunzburger, M., Ringler, T., Burkardt, J., and Peterson, J.: Parallel algorithms for planar and spherical Delaunay construction with an application to centroidal Voronoi tessellations, *Geosci. Model Dev.*, 6, 1353–1365, <https://doi.org/10.5194/gmd-6-1353-2013>, 2013.

- Jamin, C., Alliez, P., Yvinec, M., and Boissonnat, J.-D.: CGALmesh: A Generic Framework for Delaunay Mesh Generation, *ACM T. Math. Software*, 41, 1–24, <https://doi.org/10.1145/2699463>, 2015.
- Klingner, B. M. and Shewchuk, J. R.: Aggressive Tetrahedral Mesh Improvement, in: *Proceedings of the 16th international meshing roundtable*, edited by: Brewer, M. L. and Marcum D., Springer, Berlin, Heidelberg, Germany, 3–23, 2008.
- Lai, Z., Chen, C., Cowles, G. W., and Beardsley, R. C.: A nonhydrostatic version of FVCOM: I. Validation experiments, *J. Geophys. Res.-Oceans*, 115, C11010, <https://doi.org/10.1029/2009JC005525>, 2010.
- Lambrechts, J., Comblen, R., Legat, V., Geuzaine, C., and Remacle, J.-F.: Multiscale mesh generation on the sphere, *Ocean Dynam.*, 58, 461–473, 2008.
- Lawson, C. L.: Software for C^1 Surface Interpolation, in: *Mathematical Software III*, edited by: Rice, J. R., 161–194, Academic Press, New York, USA, 1977.
- Lipnikov, K., Manzini, G., and Shashkov, M.: Mimetic finite difference method, *J. Comput. Phys.*, 257, 1163–1227, 2014.
- Löhner, R.: Regridding surface triangulations, *J. Comput. Phys.*, 126, 1–10, 1996.
- Madec, G.: NEMO ocean engine (version 3.3), Tech. Rep. 27, Note du Pole de moélisation, Institut Pierre-Simon Laplace (IPSL), Paris, France, 2011.
- Marshall, J., Adcroft, A., Hill, C., Perelman, L., and Heisey, C.: A finite-volume, incompressible Navier-Stokes model for studies of the ocean on parallel computers, *J. Geophys. Res.*, 102, 5753–5766, 1997.
- Murray, R. J.: Explicit generation of orthogonal grids for ocean models, *J. Comput. Phys.*, 126, 251–273, 1996.
- Pain, C., Piggott, M., Goddard, A., Fang, F., Gorman, G., Marshall, D., Eaton, M., Power, P., and De Oliveira, C.: Three-dimensional unstructured mesh ocean modelling, *Ocean Model.*, 10, 5–33, 2005.
- Persson, P.-O.: Mesh Size Functions for Implicit Geometries and PDE-based Gradient Limiting, *Eng. Comput.*, 22, 95–109, <https://doi.org/10.1007/s00366-006-0014-1>, 2006.
- Persson, P.-O. and Strang, G.: A simple mesh generator in MATLAB, *SIAM Rev.*, 46, 329–345, 2004.
- Piggott, M., Gorman, G., Pain, C., Allison, P., Candy, A., Martin, B., and Wells, M.: A new computational framework for multi-scale ocean modelling based on adapting unstructured meshes, *Int. J. Numer. Meth. Fl.*, 56, 1003–1015, 2008.
- Putman, W. M. and Lin, S.-J.: Finite-volume transport on various cubed-sphere grids, *J. Comput. Phys.*, 227, 55–78, 2007.
- Randall, D. A., Ringler, T. D., Heikes, R. P., Jones, P., and Baumgardner, J.: Climate modeling with spherical geodesic grids, *Comput. Sci. Eng.*, 4, 32–41, 2002.
- Rebay, S.: Efficient Unstructured Mesh Generation by Means of Delaunay Triangulation and Bowyer-Watson Algorithm, *J. Comput. Phys.*, 106, 125–138, <https://doi.org/10.1006/jcph.1993.1097>, 1993.
- Ringler, T., Ju, L., and Gunzburger, M.: A multiresolution method for climate system modeling: application of spherical centroidal Voronoi tessellations, *Ocean Dynam.*, 58, 475–498, 2008.
- Ringler, T., Thuburn, J., Klemp, J. B., and Skamarock, W. C.: A unified approach to energy conservation and potential vorticity dynamics for arbitrarily-structured C-grids, *J. Comput. Phys.*, 229, 3065–3090, 2010.
- Ringler, T., Petersen, M., Higdon, R. L., Jacobsen, D., Jones, P. W., and Maltrud, M.: A multi-resolution approach to global ocean modeling, *Ocean Model.*, 69, 211–232, 2013.
- Rypl, D. and Krysl, P.: Triangulation of 3D surfaces, *Eng. Comput.*, 13, 87–98, 1997.
- Schreiner, J., Scheidegger, C. E., Fleishman, S., and Silva, C. T.: Direct (Re)Meshing For Efficient Surface Processing, *Comput. Graph. Forum*, 25, 527–536, <https://doi.org/10.1111/j.1467-8659.2006.00972.x>, 2006.
- Sein, D. V., Danilov, S., Biastoch, A., Durgadoo, J. V., Sidorenko, D., Harig, S., and Wang, Q.: Designing variable ocean model resolution based on the observed ocean variability, *J. Adv. Model. Earth Syst.*, 8, 904–916, <https://doi.org/10.1002/2016MS000650>, 2016.
- Skamarock, W. C., Klemp, J. B., Duda, M. G., Fowler, L. D., Park, S.-H., and Ringler, T. D.: A multiscale nonhydrostatic atmospheric model using centroidal Voronoi tessellations and C-grid staggering, *Mon. Weather Rev.*, 140, 3090–3105, 2012.
- VanderZee, E., Hirani, A. N., Guoy, D., and Ramos, E.: Well-centered Planar Triangulation – An Iterative Approach, in: *Proceedings of the 16th International Meshing Roundtable*, edited by: Brewer, M. L. and Marcum, D., Springer, Berlin, Heidelberg, 121–138, https://doi.org/10.1007/978-3-540-75103-8_7, 2008.
- VanderZee, E., Hirani, A. N., Guoy, D., and Ramos, E. A.: Well-centered triangulation, *SIAM J. Sci. Comput.*, 31, 4497–4523, <https://doi.org/10.1137/090748214>, 2010.
- Vitousek, S. and Fringer, O. B.: A nonhydrostatic, isopycnal-coordinate ocean model for internal waves, *Ocean Model.*, 83, 118–144, 2014.
- Wang, Q., Danilov, S., Sidorenko, D., Timmermann, R., Wekerle, C., Wang, X., Jung, T., and Schröter, J.: The Finite Element Sea Ice-Ocean Model (FESOM) v.1.4: formulation of an ocean general circulation model, *Geosci. Model Dev.*, 7, 663–693, <https://doi.org/10.5194/gmd-7-663-2014>, 2014.