

High-Order Curvilinear Mesh Generation From Third-Party Meshes

Kaloyan S Kirilov^{a,b}, Jingtian Zhou^a, Joaquim Peiró^a, David Moxey^b

^a*Department of Aeronautics, Imperial College London, South Kensington Campus, London, SW7 2AZ, U.K.*

^b*Department of Engineering, King's College London, Strand, London, WC2R 2LS, U.K.*

Abstract

Established *a posteriori* mesh generation, high-order mesh curving and some mesh optimization approaches often rely on an accurate CAD parametrization of the boundary of the computational domain. This information, however, is not always available, especially when composite multi-software workflows are employed. To deal with such cases, we propose a method for reconstructing the missing connectivity information between the mesh and the CAD geometry when importing an arbitrarily sourced mesh. The reconstruction is followed by curving methods for order elevation, projections or subsequently optimisations with boundary-conforming node sliding. Lastly, mesh modification techniques are used to achieve the desired mesh resolution and quality for meshes incorporating boundary layers. We illustrate the steps of the proposed end-to-end workflow through two simple geometries coming from different sources and an end-to-end complex automotive mesh generation test case.

Keywords: High-Order, Curvilinear, Mesh Generation, Spectral/hp Element Methods, Industrialization, Complex Geometries, CAD

Email addresses: `kaloyan.kirilov@kcl.ac.uk`,
`kaloyan.kirilov19@imperial.ac.uk` (Kaloyan S Kirilov),
`jingtian.zhou22@imperial.ac.uk` (Jingtian Zhou), `j.peiro@imperial.ac.uk`
(Joaquim Peiró), `david.moxey@kcl.ac.uk` (David Moxey)

1. Introduction

High-fidelity, high-order solvers such as Nektar++ [1], Nek5000 [2] or PyFR [3], require valid high-quality curvilinear meshes that accurately conform to the underlying geometry. Such meshes are commonly generated through an *a posteriori* approach, which begins with a linear straight-sided mesh that is elevated to the required order and then curved to conform to the curves and surfaces defining the CAD geometry. Additional optimisation and untangling techniques are usually adopted during or after the curving process to ensure validity and enhance the quality of the mesh [4, 5, 6, 7]. Different variations of this approach have been the backbone of available open-source tools such as Gmsh [8] and NekMesh [4], and in commercial codes such as GridPro, Pointwise (now Cadence) [9] and more recently ANSA [10]. These developments have led to increased interest from industries, such as aerospace, turbo-machinery and motorsport, to adopt these high-fidelity simulations within their CFD workflows.

In an industrial setting, where geometries can consist of thousands of CAD surfaces and mesh volumes, the development of a methodology capable of automatically, accurately and efficiently resolving the required flow features in fast turn-around time, presents a huge challenge for mesh generation tools. A natural solution is to take advantage of existing fast, robust and flexible mesh generators to produce linear meshes suited to industrial applications. One of the main drawbacks of this approach is that often these mesh generators do not provide the connectivity between the surface mesh and the underlying CAD objects and sometimes do not guarantee CAD compliance. This is of particular importance in the mesh curving process since most CAD related operations are usually performed in the parametric space of the CAD objects. If this information is available, one can elevate the order and curve the mesh to the corresponding CAD. One example is the work by Ruiz-Gironés et al. on the High-Lift Common Research model using Pointwise CAD API and extensive CAD-preparation [11]. Another example is the successful mesh generation, after very extensive de-featuring and BRep simplification, of a full road sport car with NekMesh by the Nektar++ group [12]. However, integration of *a posteriori* approaches in industry is hampered by the need of extensive and time-consuming manual CAD preparation, the absence of CAD-mesh connectivity, and the lack of mesh curving procedures allowing robust mesh optimization on the CAD objects.

To address these issues, we propose a generic methodology for generat-

ing high-order meshes from a very coarse linear mesh through the reconstruction of the missing mesh-CAD connectivity. Moreover, we ensure full CAD-conformity through vertex projections. This allows us to generate valid high-order meshes even from invalid, non-watertight CAD BRep STEP files of complex geometries, assuming a valid linear volume mesh is provided, for example generated from a separate triangulation. As this can be reliably achieved by state-of-the-art mesh generators, the requirements and the need for CAD simplification are significantly relaxed. With this newly acquired CAD-mesh information, we reuse and improve the internal bottom-up approaches for both mesh curving, isoparametric prismatic layer splitting, mesh optimization and untangling as described in [4] and in [13]. This is an extension to our SIAM IMR24 conference publication [14] and adds to it an improved core CAD-reconstruction process with better accuracy and improved robustness, an extension of the method to all common type of structured and unstructured meshes and a new application showcasing the full flexibility of the workflow for mesh curving and optimization with linear and also already high-order meshes generated from third-party software.

The outline of this paper is the following. Section 2 presents a high-level overview of the proposed methodology. Section 3 discusses the implementation details, describing every module of the workflow, while focusing on the challenges and the solutions that allow us to curve and optimise the imported third-party meshes. Section 4 illustrates the methodology with three examples. The first example uses a simple analytical geometry, a semi-sphere, for ease of illustration. The second example is an inverted multi-element wing that demonstrates the flexibility of the proposed approach by employing different mesh sources and mesh types. Finally, we conclude this work with a validation study: the simulation of the flow past the Imperial Front Wing [15]. This is a motorsport benchmark used to demonstrate the suitability for potential industrialization of the method.

2. Overview of the Methodology

To tackle the problem of industrialization of high-order methods for CFD applications, we propose a workflow that combines robust linear mesh generation tools with *a posteriori* mesh curving and optimizations with mesh sliding on the CAD geometry. The starting point of the process is a coarse straight-sided (linear) volume mesh with fully hexahedral or hybrid prism-tetrahedral elements. The main objective is to enable the user to seamlessly

switch between meshing tools from different providers, including both generic (ex: StarCCM+ [16], Gmsh [8]) and application-specific (ex: PADRAM[17], TAU[18]) mesh generators at no additional cost.

When generating the curvilinear mesh, we aim at reusing and, whenever possible, improving current high-order *a posteriori* curving and optimisation techniques. However, these methods often rely on some critical assumptions that are often not satisfied when using a third-party linear meshes:

- *CAD-connectivity*: The mesh curving and optimisation procedures require accurate parametric information (u,v) of the underlying CAD object.
- *CAD-conformity*: The vertices of the linear mesh lie on the CAD curves and surfaces within machine precision (ϵ_M) .
- *CAD-compliance*: The mesh curving processes may assume that every mesh face lies entirely in precisely one BRep surface.

We propose a new intermediate step in the mesh generation process, refer to here as CAD Reconstruction module that is illustrated in Figure 1. This module first associates the vertices of the linear mesh with the closest BRep objects (surface, edge or vertex) and their coordinates in parametric space. Conformity of the linear mesh to the CAD geometry is then ensured through vertex projection. A series of logical steps is subsequently employed to reconstruct the CAD connectivity of all surface entities of the mesh, i.e. edges and faces, by associating each mesh object with its corresponding BRep object. After establishing this link, accurate parametric projections and surface optimisations are performed for mesh curving. To ensure validity, and to increase the mesh quality, we can optionally employ a variational optimiser to untangle any invalid elements. The availability of the newly acquired CAD connectivity now allows the sliding of high-order nodes on the CAD BRep object in mesh optimization.

Special consideration is taken when a linear element (boundary face or edge) crosses two BRep surfaces. In these cases, no BRep object is assigned to the mesh entity and a simplified projection process follows with no further optimisations on the surface. This projection may result in a slight deterioration of the surface quality in cases where the multi-patch BRep has only C_0 continuity between patches. However, it also allows for significantly increased flexibility on the linear mesh generation side, decreasing the need




Figure 1: The proposed mesh generation workflow starts from third-party linear meshes (blue) which are processed by in house *a posteriori* high-order modules (orange).

for CAD defeaturing and also facilitating the use of two separate geometrical representations for the linear mesh generation (e.g. a triangulation) and for the high-order part of the process (invalid STEP/IGES). Whether the right balance between surface quality and volume mesh generation flexibility is achieved will depend on the application and complexity of the case. Lastly, the process is extended to high-order meshes, that can be produced either by in-house or external high-order mesh generators, see Figure 2.




Figure 2: The proposed mesh generation workflow working directly on external or internal high-order meshes.

This offers two main advantages. Firstly, it provides the means for implementing a mesh quality suite where the geometrical accuracy of any high-order mesh can be evaluated if a CAD file is present. Secondly, it permits the user to start the workflow from any internal or external high-order curved meshes, assess their quality and then use *a posteriori* in-house optimization techniques with node sliding on the CAD geometry. This is also crucial for the mesh modifications required in *a posteriori* mesh adaptation.

3. Proposed Workflow Implementation

This section presents the proposed workflow, covering all stages - from CAD preparation and linear mesh generation, mesh curving, optimization and finally ensuring the required "y+" spatial resolution in the wall normal direction, see Figure 3. This work is within the high-order mesh generator NekMesh [13], part of Nektar++ [1].

3.1. Linear Mesh Generation & Common Issues




Figure 3: Overview of the end-to-end mesh generation pipeline, showing all steps/modules in generating the valid high-order mesh. Processes marked in blue are third-party software, whereas all orange ones are internal to NekMesh.

The starting point of the process is the creation of the underlying CAD model and the generation of the linear straight-sided mesh. The way these are generated is left completely to the desires and needs of the user. However, there are best practices that have been discovered during the development and the applications of this work. In terms of the CAD BRep, there are no specific requirements on the user for the linear generation. The only necessity on the high-order side of the process is the availability of a BRep in the form of a STEP/IGES file.

The workflow allows the generation of a linear mesh from various CAD file formats or even from a triangulation. This is advantageous when handling complex geometries because it does allow some imperfections in the BRep. Their presence might lead to a slight deterioration of the surface quality of the mesh, but undertaking CAD repair could lead to unacceptable time scales in an industrial environment. As a best practice, it is recommended to use a single thick O-type prismatic or hexahedral that is with sufficient thickness to account for the mesh curving. While not strictly necessary, this streamlines the presented workflow drastically.

Currently, NekMesh supports the input from the following mesh standard formats: CGNS (.cgns), Gmsh (.msh), Nektar++ (.xml) and StarCCM+ through the CCM to OpenFoam API [19]

libraries. Supporting other input formats is a relatively straight forward task to implement due to the modularity of NekMesh [1].

Ensuring CAD Conformity

In regions of high curvature, the mesh vertices often do not lie precisely on the CAD, which we will denote the minimum distance to the target CAD object as δ . Such deviations may arise naturally from discrepancies between the linear mesh CAD source and the CAD BRep, or from surface meshing algorithms tailored for finite-volume methods, which ensure the centroids being on the CAD.

These geometrical errors significantly impact the geometrical accuracy of the mesh and, potentially, of the simulation. For instance, they could introduce waviness on the boundary of the mesh and trigger transition of the flow. They can worsen C_1 discontinuities at the interfaces between elements. Figure 4 shows an example where a tessellation generates a high deviation of $\delta \approx 2 \times 10^{-4}$ on the low curvature regions of a wing. A naive projection of only the high-order nodes onto the BRep results in an unacceptable spike-like pattern on the curved surface mesh, which subsequently decreases the geometrical accuracy but also the volume mesh quality.




Figure 4: An illustration of the effect of vertex CAD conformity. A “very fine” surface triangulation generates an oscillatory pattern of a wing, once the mesh is curved, due to $\delta = 2 \times 10^{-4}$ between CAD sources of truth and lack of vertex projection on the BRep.

CAD Compliance and Multi-patch Elements

Watertight and valid complex CAD BRep often exhibit small and skewed surfaces due to certain CAD operations such as fillets, see Figure 5a). If the linear mesh is assumed to be CAD compliant in such instances, the generated mesh could have numerous highly skewed surfaces which will deteriorate volume mesh quality, see Figure 5b).




Figure 5: Area with 7 highly skewed CAD surfaces leading to sliver elements if CAD compliance is strictly enforced. Note the deterioration in mesh quality in elements if projections are inaccurate (blue)

A remedy employed by many linear meshers is to combine these troublesome CAD surfaces into multi-patch or virtual geometries. This creates multi-patch surface elements that span between two CAD surfaces and could lead to inaccurate projections and self-intersections during the mesh curving stage. This is especially relevant in cases where the surfaces are $C1$ -discontinuous, hence special treatment should be considered for these elements. Such an example, where a simple projection has been performed even in a relatively fine mesh, can be seen in Figure 6.

Observation 1: The CAD conformity of the linear mesh is of crucial importance for the surface and volume mesh quality of the high-order mesh. Therefore, one must ensure appropriate treatment for the linear vertices on the high-order side of the process.

Observation 2: Any high-order process benefits from the CAD compliance and minimizing multi-patch element faces but also requires keeping reasonable mesh quality by employing multi-patch elements. These conflicting requirements must be carefully accounted for the linear meshing process




Figure 6: Mesh distortions and self-intersections due to naive projection on multi-patch elements (BRep line in white) at polynomial order $p = 4$.

through the use of mesh modification techniques *a posteriori* or approximative mesh curving inside NekMesh.

3.2. CAD Reconstruction

This section describes a novel method for reconstructing the lost CAD mesh connectivity which combines bottom-up approaches for vertex BRep association through the CAD API with a set of top-down logical operations derived from the BRep topology. This module is designed to be applied on complex geometries and to be compatible with any linear mesh generator. To achieve this we have eliminated a large number of restrictions and assumptions of the current processes. For example, it allows for lack of CAD conformity of the vertices and for the presence of multi-patch elements. Moreover, the user can provide usually invalid CAD. Further, the new modules of NekMesh permit a full integration with third-party curved meshes. The final outcome of the process is the connectivity between the BRep and the mesh objects, as illustrated in Table 1.

Table 1: The final association between BRep and mesh objects.

Mesh/CAD	Vertex	Edge	Face	HO-EdgeNode	HO-FaceNode
CAD-Vertex	✓	✗	✗	✗	✗
CAD-Curve	✓	✓	✗	✓	✗
CAD-Surface	✓	✓	✓	✓	✓

Algorithm 1 CAD-Reconstruction

In: Mesh $\mathcal{M}$, CAD-BRep (STP/IGES)**Out:** Bnd Conforming $\mathcal{M} + \text{CAD-mesh link}$

- 1: **Load** CAD
 - 2: **Create Bounding Boxes** - *CADSurf*, *CADCurve*, and *CADVert*
 - 3: **Link** bnd vertices to CAD objects, (ϵ_{loose})
 - 4: **Project** bnd vertices to closest CAD object
 - 5: **Link** bnd vertices to CAD objects, (ϵ_{tight})
 - 6: **Link** mesh faces to a CADSurf
 - 7: **Link** mesh edges to a CADCurv/Surf
 - 8: **if** High-Order (HO) nodes exist **then**
 - 9: **Link** HO-nodes to parent edge CAD
 - 10: **Link** HO-nodes to parent face CAD
 - 11: **end if**
 - 12: **if** edge !CADobject **then**
 - 13: **Project** multi-patch edges to closest CADSurf (node-by-node)
 - 14: **end if**
-

BRep and CAD API

The CAD APIs handle the geometric information provided in Boundary Representation (BRep) format, typically via STEP or IGES files. As the linear volume mesh is already available before the curving stage, the existing mesh connectivity data can be leveraged to relax traditionally strict geometric requirements imposed on STEP. This flexibility significantly improves robustness against common STEP-related artifacts, such as free edges, pierced surfaces, non-manifold edges and other topological errors that typically pose challenges in the initial volume meshing. Moreover, a user can provide a non-watertight geometry covering only certain parts of interest of the body for mesh curving or can reuse the geometry multiple times by excluding the outer CFD domain from the BRep. The only requirement for the BRep surfaces is a valid underlying parametric definitions.

All geometrical operations and queries performed within NekMesh rely on dedicated CAD APIs, which interface directly with established CAD kernels. Currently, NekMesh supports two distinct CAD engines: OpenCascade [20] and ITI CADfix [21]. In this study, the OpenCascade one was used. While some extensions were part of this work, we refer to the one by Turner [4] and Marcon [22] for the details of these CAD APIs.

Bottom-up BRep Pre-Processing

The vertices of the initial linear mesh are the only mesh entity sufficiently close to the BRep. Hence, the first step of the reconstruction process is to associate every boundary vertex of the linear mesh with the CAD surfaces, CAD Curves and CAD vertices. A naive approach that would query the distance to every BRep object is unfeasible for industrial geometries that may contain more than 10^4 surfaces. To reduce the cost of such queries, a preprocessing of the BRep is first performed. All CAD surfaces are sampled with nodes from which their bounding boxes are calculated and are enlarged. Then all these bounding boxes are stored in a k-d tree, as in the original NekMesh workflow [23]. Here we extend this procedure to for every CAD curve and every CAD vertex, all stored in three separate k-d trees.

Vertex Associations

This way every boundary vertex can be efficiently associated with only a few candidate surfaces when it is inside their bounding box, see Figure 7. These can vary from 3 to 20 per vertex roughly depending on the CAD complexity. The vertex is then virtually projected only onto the candidate surfaces, the surfaces are ordered according to the distance δ , and discard those where $\delta > \delta_{tol}$. A default value used for that is $\delta_{tol} = \max(0.5L, 10^{-5})$ where L is the length of the edge, but these can be adjusted depending on the application. If more than one CAD Surfaces is shortlisted, we proceed with a look-up for CAD Curves and CAD Vertex. Finally, we project the vertex to the closest CAD Curve or CAD surface. To ensure even further reduction of false positive associations, we run the same cycle of associations without the final projection with a 10 times stricter tolerance and reevaluate the list with the parametric location for all sufficiently close candidate CAD objects. This way we ensure the CAD conformity of the linear mesh and a minimal amount of false positives, see Figure 8

Face-Associations

Now that the mesh vertices have been linked to their corresponding CAD objects, we move forward to boundary faces. From a CAD connectivity standpoint, boundary faces are classified into three distinct types: internal elements, corner elements (also referred to as Case 2), and multi-patch elements spanning over two CAD surfaces (Case 3). These cases are illustrated in Figure 8 with a triangular face, but the implementation is applied for




Figure 7: Diagram showing the process of vertex association to BRep objects before projection. Note that none of the vertices is exactly on the BRep. Vertices V_1 and V_3 will have only one associated CAD curve, whereas V_2 will have two.

both triangular and quadrilateral faces. Thus achieving a full coverage for all common mesh elements types.




Figure 8: Diagram showing the 3 different element cases: (a) internal face, (b) corner (Case 2), and (c) multi-patch (Case 3) with initial relaxed and consequently strict association tolerances.

The CAD reconstruction for “internal” elements is the easiest as all its vertices share a single common surface. In addition, at least one vertex is associated only with the common surface. These two conditions are usually met by more than 95% of the boundary elements, enabling for better parametric mesh curving, surface and volume optimisations.

Case 2 elements, with vertices having more than one common CAD surface, are more difficult to handle. These often appear when one of the surfaces is very small or with a very high aspect ratio like, for example in blunt trailing edges or sliver surfaces as illustrated in Figure 9. In this case, the algorithm samples the mid-point on all edges and the centroid of the face, calculates the distance to the candidate CAD surfaces with the CAD API

and then picks the one with the minimum overall deviation. The surface ID is obtained as follows

$$\text{ID} = \arg \min_k \sum_{j=1}^3 \|\mathbf{x}_j - \mathbf{P}_k(\mathbf{x}_j)\| \quad (1)$$

where $\mathbf{P}_k(\mathbf{x}_j)$ denotes the coordinates of the projection of the point $\mathbf{x}_j$ on the surface of ID k . This is effective for both straight-sided and curved surfaces, given a with a reasonably deviation from their linear approximation, which we allow to be $\delta_{\max} = 0.5 \times L_{\text{edge}}$.




Figure 9: Specific treatment of the objects with multiple CAD surface candidates in a region of trailing edge or corner elements (Case 2).

Finally, if boundary face vertices do not share any common CAD surfaces or the element fails any of the previously described filters, no BRep object is associated with the element, Case 3 face. Hence, a simpler mesh curving algorithm is applied to this element without local surface optimisation and these faces do not slide on the CAD during the global volume optimisation.

Edge-Associations:

The edge association to CAD surfaces is relatively straightforward. By revisiting Figure 7, it can be observed that when the vertices have a single associated common CAD Curve or CAD Surface, this edge is internal to the BRep object and can be directly linked to it. In the case of CAD Surface association, a cross-check is performed with the parent element CAD surface. A priority is given to the CAD Curve link.

If the vertices share two or more CAD surfaces, but not a CAD Curve, we apply a redundancy association approach. Topologically, this indicates a corner edge likely generated on a CAD curve. Therefore, we loop over all

curves, forming the EdgeLoops around the candidate CAD surfaces creating a shortlist of CAD curve candidates. If the edge has a single common one, it is connected to this CAD curve. If two curves are present (around junctions), another voting process is performed, choosing the closest one within a tolerance of δ_{tol} . In the case of more than two common CAD curves, we prefer the safer choice of assigning the elemental CAD surface to the edge to facilitate a more accurate projection and surface optimisation. Lastly, if the vertices of the edge do not share a CAD surface or fail the CAD curve tolerance test, this indicates that the edge is crossing multiple patches. Hence no BRep object is associated at this stage, and a simpler curving process is applied as per the Case 3 elements.

Extension to High-Order Meshes

Once the mesh-CAD connectivity is available for the linear mesh, we can extend the method to high-order mesh entities such as edge nodes and face nodes. This is straightforward since the edge nodes will adopt the CAD curve or CAD surface of the parent edge, and the face nodes the CAD surface of the parent face. Interestingly, this brings about the unique workflow capability to restart the high-order meshing from any point of the process. Moreover, it allows to restart from any kind of a third-party mesh, being it linear or high-order. To the best of our knowledge, this is a unique capability and has not been reported previously.

3.3. Mesh Curving and Surface Quality

Once the connectivity between the linear mesh and the CAD BRep is established, we can apply any high-order mesh curving module that uses the CAD BRep. In this workflow, we modify and enhance the interpolative mesh curving algorithm, which is based on the parametric projection of high-order nodes onto the BRep, followed by surface mesh optimization, part of NekMesh [13]. This has already been proven to produce high-quality meshes in the in-house bottom-up mesh generation workflow [4, 24] that generates hybrid tetrahedral and prism-tet linear meshes. Moreover, we extend that capability for also fully hexahedral or hex-dominant meshes.

The mesh curving starts with the parametric projection, which is applied hierarchically for every boundary face in the parametric space of the corresponding CAD object (u, v) . First the edges are curved, followed by the internal face nodes. When we consider the edges, the high-order nodes are created in the reference space according to the user-defined order and

quadrature rule, usually Gauss-Lobatto-Legendre. Then following the edge vertices in (u, v) or (t) the nodes are mapped to the parametric space. Given the (u, v) coordinates of every high-order edge node, we use the CAD API to map back its location back to Cartesian coordinates (x, y, z) . The process is illustrated in Figure 10.

After the parametric projection of the edge nodes, we apply the same type of parametric projection on the corresponding internal face nodes, see Figure 10. Similarly, the desired quadrature nodes are generated according to the user-defined order and quadrature, then mapped to the parametric space of the corresponding BRep surface with the parametric knowledge of the vertices. Last, these are projected on the CAD and the validity of the element is checked. This way, the workflow provides very high geometrical surface accuracy with the user-defined quadrature order, tested up to polynomial order $P = 14$.




Figure 10: The 3D parametric projection of the high-order face nodes in parametric (u, v) to physical space for triangular p4 and quadrilateral p3 faces).

In the rare cases where the boundary face does not possess a CAD surface association, i.e. they are not CAD compliant, we rely on a simplified secondary mesh curving process. In this case, the face has at least one edge crossing two BRep surfaces. As discussed earlier, this is not a desirable situation where the exact parametric projection is not possible. However, when faced with BRep slivers or meshing from triangulations, it is crucial to account for them and devise strategies to avoid having linear or tangled ele-

ments. Therefore, we populate the desired quadrature on the linear element and then project the high-order edge nodes one by one to the closest CAD surface, as depicted in Figure 11. This fail-safe process relies on the curved edges to reconstruct the underlying geometry for robustness purposes. The reason being that projecting high-order face nodes on the closest CAD surface is more likely to result in self-intersections, especially when the two surfaces have high curvature and are only C_0 continuous. No surface optimisation is applied to these mesh faces.




Figure 11: Diagram showing the 3D projection of edge nodes from Cartesian space (x, y, z) onto the closest CAD surface.

Surface Mesh Optimizations

During the mesh curving, the workflow moves the high-order nodes on the BRep object to minimise the distortion induced by the parametric projection according to the optimization method described in references [25, 26]. On an edge, we achieve that by introducing a virtual spring network system between the edge nodes and moving them on the BRep object to minimise the potential energy, where the w_i weights are purely geometrical - with $w_i = \xi_{i+1} - \xi_i$, where ξ_i is the node location in the reference space as defined by the nodal distribution. A minimum of this energy is calculated in Cartesian space using a classical BFGS optimization algorithm [27].

$$F = \sum_{i=1}^P \frac{\|\mathbf{x}(u_{i+1}, v_{i+1}) - \mathbf{x}(u_i, v_i)\|^2}{w_i}. \quad (2)$$

This process becomes important when coarse elements are present in regions with very high two-directional curvature, such as wingtips or rounded trailing edges. In these cases, we observe a significant increase in the surface mesh quality, geometrical accuracy and even untangling of self-intersections,

as shown in Figure 12. Naturally, this translates to an improved overall volume mesh quality.




Figure 12: The curved surface mesh generated with (b) and without (a) the optimisation. Note the presence of invalid elements when the optimisation is not applied.

3.4. Assessing Geometrical Mesh Accuracy through CAD-Reconstruction

The geometrical errors in the mesh discretization of the underlying CAD could have a significant impact on the accuracy of simulations. Often it is difficult to identify these errors. The high-order extension of the CAD-Reconstruction algorithm provides a tool for assessing the geometrical accuracy of the mesh - the deviation from the target CAD object $\delta \mathbf{x}$ and from the target CAD Normal $\delta \mathbf{n}_i$.

The first stage of the process is to elevate (sample) the boundary face to an order of choice with the desired quadrature points. Then the CAD connectivity is reconstructed and all vertices, high-order edges and face nodes receive their BRep parametric coordinates (u_i, v_i) . This backward mapping allows us to then query through the CAD API for the target location $\mathbf{x}_i$ and the unit normal vector $\mathbf{n}_i$ on the associated CAD object. Last, we calculate the corresponding nodal deviations for $\mathbf{x}_i$ and $\mathbf{n}_i$ between the isoparametric and BRep nodes, see Figure 13. Through the Nektar++ high-order utilities, we can visualise the elemental values of these deviations and assess if any further interventions are necessary. This also allows the user to visualise and investigate if, where and why an element does not have CAD associated with it which will indicate the presence of a multi-patch face or a rare failure of the CAD-Reconstruction algorithm.




Figure 13: Diagram showing the mechanism of assessing the geometrical accuracy $\delta \mathbf{x}_i$ of a $P = 2$ mesh evaluated at $P = 6$ with an evenly distributed quadrature rule.

3.5. Variational Optimizer

To improve mesh quality and untangle invalid elements, we (re)use the variational optimiser module in the NekMesh framework [26]. This module minimizes the deformation energy functional W , which is a function of the mapping ϕ from the ideal straight-sided element Ω_I^e to the curvilinear element Ω^e , and this is equivalent to finding

$$\min_{\phi} \mathcal{E}_i(\nabla \phi) = \sum_{e \in \mathcal{E}} \int_{\Omega_I^e} W(\nabla \phi) dy. \quad (3)$$




Figure 14: The mappings ϕ from standard reference element Ω_{st} to ideal straight-sided Ω_I^e and the final curved element Ω^e . Note that the mapping used for optimization is $\phi = \phi_M \circ \phi_I^{-1}$ [23].

The various mappings involved in this setting are depicted in Figure 14. The optimization calculates the positions of the nodes that lead to a minimum of the energy in Equation 3. We direct the reader to reference [4] for a detailed description of the method. The energy functional adopted here is based on hyperelastic materials, see Equation 4.

$$W = \frac{\mu}{2}(\text{Tr}(\mathbf{C}) - 3) - \mu \ln J + \frac{\lambda}{2}(\ln J)^2 \quad (4)$$

where $\mathbf{J} = \nabla\phi(\mathbf{y})$ is the gradient tensor of the mapping, its determinant $J = \det(\nabla\phi)$ is referred to as the Jacobian, λ and μ are the Lamé elastic constants, and $\mathbf{C}$ is the Cauchy-Green stress tensor.

The untangling capability of the method is achieved through the regularisation of the Jacobian for values $J \leq 0$, and its performance is significantly improved by allowing the movement of the nodes on the BRep object, except for the mesh vertices associated with the CAD vertices. A minor limitation of our process is that we guarantee a perfect CAD surface, but not always a perfect CAD curve association. To avoid this problem and increase robustness, we fix the nodes on the boundary linked to two or more surfaces or curves, and the boundary nodes without a CAD object. While certain enhancements, such as enabling nodes to slide between CAD surfaces will improve mesh quality further in problematic regions like CAD slivers, we use the module in its current form and leave these for future work.

3.6. An Improved Isoparametric Boundary-layer Splitting

Viscous flow simulations require stretched elements with very high aspect ratio near the wall boundaries. The main challenge in generating these boundary-layer meshes is to propagate the curvature inside the domain and to ensure the quality and validity of the final mesh. For this purpose, we employ the NekMesh isoparametric approach for generating and curving a thick single prismatic layer around the body, the macro-prism layer [28]. Every macro-prism is mapped to reference space and split according to a user-defined progression ratio (r) and a number of boundary layers. This work introduces an improved implementation in NekMesh by propagating not only the edge curvature through the edge nodes but also all face nodes, which results in a more accurate approximation of the CAD surface by the elemental boundary faces.




Figure 15: Isoparametric boundary-layer splitting in reference space (ξ) and the mappings from reference to physical space [28].

4. Verification and Validation

To demonstrate the quality and effectiveness of the proposed meshing technique, we present three examples of increasing geometrical complexity: one academic and two automotive. These include a simple domain around a hemisphere, an inverted multi-element extruded wing to show the flexibility of the approach and the full version of the Imperial Front Wing as a complex motorsport geometry [29, 15] representative of an industrial application.

4.1. A Semi-spherical Geometry Test Case - Consistency

This elementary geometry is a domain around a hemisphere. The process works as follows. A very coarse linear mesh is generated from STEP using a third-party software, consisting of only 4 prism elements representing the hemisphere, see Figure 16a. The CAD-Reconstruction process successfully associated the necessary connectivity between all the mesh elements and their associated BRep objects, and performed parametric projections with surface optimisation, as described in Section 3, to polynomial order $P = 9$ producing the mesh shown in Figure 16b.




Figure 16: The high-order mesh generation of a domain around a hemisphere: (a) Initial linear mesh , and (b) the curved surface mesh optimised at polynomial order $P = 9$.

Given the analytical nature of the geometry, we can evaluate the surface mesh geometrical accuracy against the CAD model but also evaluate the error of the model itself against the analytical hemisphere, see Figure 17a. The CAD BRep we used had a maximum error of $\delta_{CAD} = 3.2 \times 10^{-5}$ located around the base and the apex. Our mesh shows a similar error distribution to the CAD reference with slightly smaller δ_{max} against the analytical hemisphere, but also marginally increased δ around edge, see Figure 17b.

To quantify geometric accuracy against the CAD model, we use the proposed quality suite. Each P=9 isoparametric element is sampled with evenly spaced interpolation points (N=16 and N=21). This high sample number was chosen due to the non-uniform distribution of the δ inside the element and the sensitivity to their location. Then the deviations δ are computed against the parent CAD object, see 17c. The analysis revealed negligible deviations for the element edge nodes, while maximum ones concentrated primarily around face nodes close to the element edges, likely due to the P=9 and the subsequent Runge Phenomena. Overall, the maximum deviation is $\delta_{max} = 9 \times 10^{-6}$, roughly three times lower than the approximation error of the original B-Rep model. Furthermore, 74% of the sampled high-order nodes exhibited deviations of less than $\delta = 3 \times 10^{-6}$, demonstrating the workflow’s capability to accurately represent the CAD BRep to values comparable to or even lower than the original CAD model with very coarse meshes utilizing high polynomial orders.




Figure 17: Geometrical error of the CAD model a) and the P=9 HO-mesh b) against the analytical hemisphere showing comparable max errors at the base and the apex with slightly increased δ around the mesh edges. And the geometrical accuracy of the same P=9 mesh against the CAD c), evaluated at evenly spaced points HO-mesh, displaying δ_{max} against the CAD being three times lower than the CAD model approximation error.

Using the same geometry but with a finer all-tetrahedral linear mesh created separately in NekMesh, we verify the consistency between the existing mesh generation workflow, as implemented by Turner [4], and the new proposed workflow. The linear mesh is depicted in Figure 18a. The process successfully reconstructed all the CAD objects: 1650 vertices, 2474 edges, 1650 faces, and 6908 elements. This allowed accurate parametric mesh curving, surface optimisation, and the variational optimiser with sliding of the

internal to the CAD surfaces mesh nodes and vertices. As the histogram of Figure 19a shows, the initial curved mesh is valid but certain elements have low quality in terms of scaled Jacobian, with the minimum value $J_{\min}=0.36$. After the application of the variational optimiser, see Figure 19b, the overall quality of the mesh increases drastically, with a minimum scaled Jacobian $J_{\min} = 0.75$. Both the geometrical accuracy and the final mesh quality are similar to the legacy internal bottom-up mesh generation workflow with $J_{\min} = 0.74$, which is a difference of 1%. This shows the consistency of the proposed methodology with the high-quality internal bottom-up one, which is to be expected due to the reuse of the high-order modules and the successful CAD-Reconstruction.




Figure 18: The high-order mesh generation of a domain around a semi-sphere starting from a linear mesh (a) and the curved surface mesh (b).




Figure 19: The mesh quality of the high-order mesh generation of a domain around a semi-sphere: (a) before and (b) after variational optimisation.

4.2. Imperial Front Wing

This workflow is aimed at the robust and automatic generation of high-order meshes for high-fidelity simulations of flow about complex geometries. We demonstrate the suitability of the workflow from a design perspective using an industrial example of such geometry: the front wing of a motorsport car. The Imperial Front Wing, is an open geometry from a legacy F1 car provided by McLaren Racing [15]. The front wing, shown in Figure 20, consists of a main plane, two wing flaps, a nose cone, a Gurney flap, a canard, a side endplate with a connected footplate.




Figure 20: The CAD BRep geometry of the Imperial Front Wing (IFW).

The principal characteristics of the STEP BRep geometry are given in Table 2. Based on this information, the geometry can be thought of having medium complexity, but it is sufficient complex to pinpoint many of the previously described problems arising from industrial CAD such as CAD slivers, parametric mesh curving distortions and the presence of BRep curves in close proximity requiring the use of multi-patch elements, see Figures 5, 6, and 12.

Table 2: Statistics on the BRep of the Imperial Front Wing.

CAD Object	Count	Parameter	Value [m]
Vertices	351	$L_{\min \text{ edge}}$	0.4×10^{-3}
Curves	559	$L_{\text{main plane}}$	0.2
Surfaces	209	$R_{\min \text{ curvature}}$	0.2×10^{-3}

To allow for a quicker testing of numerical models and benchmark performance, a simplified 3D geometry of the IFW was created by extruding the cross-sectional slice at the mid-span location of the IFW. The 3D extruded

Imperial Front Wing with $L = 0.05c$ thickness, referred to as “eIFW” in the following, will also be used in this study for verification and testing purposes.

3D Extruded Imperial Front Wing - V&V Flexibility

This section uses the simplified extruded wing eIFW to demonstrate the flexibility of the proposed workflow in handling different types of elements: tetrahedral, hexahedral, and hybrid prismatic-tetrahedral meshes. Moreover, the straight-sided volume meshes are generated in three different software: NekMesh for the linear tetrahedral meshes, a combination of Gmsh and NekMesh for the fully hexahedral one and StarCCM+ for the hybrid mesh. The linear vertices of the NekMesh tetrahedral mesh is on the CAD, whereas the StarCCM+ and Gmsh ones have a maximum deviation of $\delta_{\max} = 2 \times 10^{-6}$ and $\delta_{\max} = 4 \times 10^{-7}$ respectively. The CAD-mesh connectivity is neither retained nor provided *a priori*. Despite the lack of CAD-mesh connectivity, the CAD reconstruction is successfully applied to all three meshes. Mesh curving and surface optimization are subsequently performed to polynomial order $P = 4$ with Gauss-Lobatto-Legendre quadrature. A comparison between the initial straight-sided and the resulting curvilinear surface meshes is shown in Figure 22. The mesh quality of the final high-order meshes can be seen in Figure 21 with a summary given in Table 3. From these, we can observe a good final mesh quality for all meshes $\min J_{sc} > 0.4$, with the hybrid prism-tetrahedral one showing the highest $\min J_{sc} = 0.68$ and not a single highly deformed element ($J_{sc} < 0.5$). This is despite the highly non-uniform curvature of the two flap leading edges being resolved by only two to four elements, respectively.

Table 3: Comparison of high-order meshes of different types for the eIFW

Mesh Type	Elements	$\min J_{sc}$	$J_{sc} < 0.5$	Comment
Tetrahedral	50750	0.42	1	
Hexahedral	25140	0.44	20	No Volume Optimisation
Prism-Tet	21885	0.68	0	
HO Tet	88054	0.32	4	58 $J_{sc} < 0$ on input (GMSH)

Last, we demonstrate the ability to import and optimize high-order meshes directly from external mesh generators. For this purpose, we use a third-order high-order curved mesh created with Gmsh without optimizing it there. Initially, the mesh contains 58 elements where $J_{sc} < 0$, making the mesh invalid. We then apply the CAD-reconstruction recreating the parametrization also of the high-order nodes. This allows our variational optimizer to effectively untangle the mesh helped by sliding the surface nodes on CAD surfaces. The final mesh is valid with $\min J_{sc}=0.32$ and only 4 highly deformed elements. Additionally, we observe a minor improvement of the geometrical accuracy on the second flap of the wing.







Figure 21: High-order mesh quality distribution for the final optimized four variants of the eIFW, evaluated through the scaled Jacobian.




(a) Tetrahedral mesh. Linear mesh generator: NekMesh.



(b) Hexahedral mesh. Linear mesh generator: 2D Gmsh + 3D extrusion: NekMesh



(c) Prism-tetrahedral mesh. Linear mesh generator: StarCCM+.



(d) HO-Tetrahedral. HO-mesh generator: Gmsh

Figure 22: The 3D linear/initial (left) and high-order $P = 4$ surface mesh (right) of a fully tetrahedral, hexahedral and hybrid prism-tetrahedral meshes, respectively. Last a $P = 3$ invalid ($J_{sc} < 0$) high-order Gmsh mesh improved with the proposed NekMesh pipeline.

Finally, the proposed workflow has been used to generate multiple hybrid prism-tet high-order meshes of the extruded IFW that have been successfully utilised to perform incompressible large-eddy simulations by Liosi, Khurana and Wüstenberg [30, 31, 32] with Nektar++. Their simulations are in good agreement with the large eddy simulations by Slaughter and also with simulations by other flow solvers such as CharLES and Horses3D [33, 34]. Several positive effects were observed with our new meshes during the simulation campaign. First and foremost, the improved geometrical accuracy enabled the correct capture of the transition mechanism through a short laminar separation bubble. Smooth shear stresses were observed without the previously present non-physical fishnet-like pattern of increased τ_{xz} that occurred due to C^1 discontinuities between the curved faces of legacy meshes. This was of great importance for the project’s success, as the legacy fishnet effect resulted in an unphysical transition. Second, the improved robustness of the pipeline generated a reduced number of invalid elements, even without the use of the variational optimiser, for all tested geometries.

In conclusion, the eIFW case demonstrated that our method handles all common unstructured meshes, manages to reconstruct the CAD-mesh link, curves the linear meshes and produces a final mesh suitable for simulation meshes from both commercial and open-source third-party sources. Moreover, the process can be restarted, can improve and even untangle third-party high-order curved meshes.

Imperial Front Wing - Industrialization

The final test case is the full Imperial Front Wing. First, a linear straight-sided mesh is generated directly from the STEP BRep in the third-party finite volume mesh generator, StarCCM+, see Figure 23. No CAD de-featuring or simplifications are applied, and CAD compliance is imposed on the whole geometry to demonstrate the reconstruction algorithm. The only region where this is not enforced is the leading edge of the footplate to avoid sliver elements, see Figure 5. A single thick prismatic layer is generated near the body, which is sufficiently thick to accommodate the mesh curving process, and tetrahedral elements are used in the domain outside the boundary layer. Additional information about the mesh entities can be found in Table 4.

Once converted to the NekMesh format, the CAD-Reconstruction algorithm is successfully applied as shown in Table 4. The application of bounding boxes extending the surfaces facilitated the association of all vertices to at least one BRep object, even if vertex deviation was initially with a




Figure 23: A linear mesh of the Imperial Front Wing. The surface mesh is coloured in grey, the prism layers in blue and the tetrahedra in yellow.

maximum deviation $\delta = 3 \times 10^{-6}$ from the BRep. The only exceptions are the 5 elements located at the end-plate region where the CAD conformity was deliberately not enforced, leading to multi-patch linear elements. These 5 elements account for less than 0.01% of the prismatic elements on the boundary. The module did not associate CAD objects to 34 edges, representing 0.04% of the surface edges. Note that the discrepancy between elements without CAD object and edges without CAD objects is expected due to the retrofitting nature of the process and the independent, but only approximate, queries of edges and elements.

Table 4: Statistics on CAD-Reconstruction module. Vertices, Edges and Faces are only counted on the IFW.

Mesh Bnd Obj.	Overall	No-CAD/ $J < 0$	Fail[%]
Vertices	27 839	0	0%
Edges	83 339	34	0.04%
Faces	55 499	5	0.01%
Elements Tets	398 548	0	0%
Elements Prism	56 592	0	0%

The most critical step for these geometries is the polynomial elevation during mesh curving, in this case to $P = 5$ with Gauss-Lobatto-Legendre quadrature. As there are also elements and edges without CAD object,

the curving module first uses the fail-safe simplified projection on these 34 edges and isolates them from the optimisation processes. For the rest of the boundary edges and elements, we can use the more accurate parametric projection and the surface optimisation presented in Section 3.3. This leads to a highly smooth isoparametric surface, see Figure 24, even for coarse mesh resolution in regions with high non-uniform curvatures.




Figure 24: Surface quality shown for an enlargement of the isoparametric surface mesh.

As expected, the smooth isoparametric surface from Figure 24 results in high geometrical accuracy once it is estimated. The vertices are ensured to be CAD-conforming and lie within a distance $\delta \approx 10^{-10}$, hence within the CAD-API tolerance. The deviation from the CAD is $\delta = 6.4 \times 10^{-7}$ when tested against evenly spaced points at $P = 6$. The maximum values appear in the regions with the highest non-uniform curvatures and coarsest resolution, with leading edges reconstructed by just two or three faces, and around the simplified projections, see Figure 25. The module indicates some higher values around the symmetry plane, but after observation, these were found to be artifacts of the secondary run of face-nodes associations and not real surface discrepancies. Depending on the applications and the Reynolds number, the user might use this information and assess if mesh resolution or polynomial order should be increased.

It is also possible to apply the variational optimiser to improve mesh quality. Initially, the mesh is invalid and has five elements with negative Jacobians. The quality after applying the optimiser and linearising the leftover tangled element is displayed in Figure 26. The histogram does show improvement in prisms with overall good quality ($J_{sc} > 0.8$), but not as good as the quality obtained for the hemisphere. This can be explained by two factors. Firstly, the initial linear mesh is highly distorted because we enforced CAD compliance despite the presence of very thin CAD sliver surfaces. Secondly, these slivers only allow for very limited CAD sliding in the example. These




Figure 25: The maximum elemental geometrical deviation, δ , from the BRep of the isoparametric mesh as evaluated at evenly spaced intervals with $P = 6$. Only surface elements with $\delta > 10^{-9}$ are visualised.

issues indicate the need for improvements to CAD sliding between two BRep objects and to high-order mesh modifications. However, such improvements are left for future work.




Figure 26: Histogram of the distribution of scaled Jacobians in a high-order IFW mesh with $P = 5$ after mesh curving and the end of the meshing workflow.

Finally, following the generation of the isoparametric boundary layer that ensures a suitable mesh resolution for the required local distance y^+ , this workflow provides high-order meshes that are currently employed for incompressible flow simulations. Meshes generated using this methodology

have been used extensively within the Nektar++ group for studying the flow around different Imperial Front Wing configurations. Multiple simulations with different h/p resolutions were performed and the studies resulted in the work by Liosi and Khurana [30], [35], [31], etc. The simulated meshes are finer and rely extensively on multi-patch elements for the sliver CAD regions. An illustrative example is the work by Liosi et al. [30], where a final mesh with over 1.3M elements produced the time-averaged vortical structures shown in Figure 27.




Figure 27: Time-averaged $\lambda - 2$ vortical structures on the solution at $Re = 220 \times 10^3$ with the Nektar++ incompressible solver using a mixed formulation with $P = 3$ for pressure and $P = 4$ for velocity. [30]

5. Conclusions

We have presented and discussed a new improved workflow for high-order mesh generation from third-party meshes focusing on the CAD-Reconstruction process, its advantages and the potential risks when using third-party meshes. A new method for reconstructing BRep parametrisation of mesh entities allows us to fully merge third-party linear mesh generators, or even high-order meshes, with modern high-order methods. Its implementation also provides a suitable metric for assessing the geometrical accuracy of high-order meshes with reference to a given CAD BRep. We demonstrated some of the capabilities of the proposed workflow for both simple academic and complex

industrial automotive geometries with different mesh elements. Finally, the improved geometrical accuracy and robustness of the proposed workflow resulted in a successful Nektar++ campaign with better physical correlation compared to legacy meshes.

Acknowledgements

This project has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 955923. The authors greatly appreciate the contribution of Alexandra Liosi who produced the solution snapshot for the simulation of the Imperial Front Wing shown here.

References

- [1] D. Moxey, C. Cantwell, Y. Bao, A. Cassinelli, G. Castiglioni, S. Chun, E. Juda, E. Kazemi, K. Lackhove, J. Marcon, G. Mengaldo, D. Serson, M. Turner, H. Xu, J. Peiró, R. Kirby, S. Sherwin, Nektar++: Enhancing the capability and application of high-fidelity spectral/ hp element methods, *Computer Physics Communications* (2019) 107110.
- [2] P. Fischer, J. Lottes, S. Kerkemeier, nek5000 Web page, <http://nek5000.mcs.anl.gov> (2008).
- [3] F. Witherden, A. Farrington, P. Vincent, PyFR: An open source framework for solving advection–diffusion type problems on streaming architectures using the flux reconstruction approach, *Computer Physics Communications* 185 (11) (2014) 3028–3040.
- [4] M. Turner, J. Peiró, D. Moxey, Curvilinear mesh generation using a variational framework, *Computer-Aided Design* 103 (2018) 73–91.
- [5] T. Toulorge, C. Geuzaine, J.-F. Remacle, J. Lambrechts, Robust untangling of curvilinear meshes, *Journal of Computational Physics* 254 (2013) 8–26.
- [6] E. Ruiz-Gironés, X. Roca, J. Sarrate, High-order mesh curving by distortion minimization with boundary nodes free to slide on a 3D CAD representation, *Computer Aided Design* 72 (2016) 52–64.

- [7] V. Dobrev, P. Knupp, T. Kolev, K. Mittal, V. Tomov, The target-matrix optimization paradigm for high-order meshes, *SIAM Journal on Scientific Computing* 41 (1) (2019) B50–B68.
- [8] C. Geuzaine, J.-F. Remacle, Gmsh: A 3-D finite element mesh generator with built-in pre- and post-processing facilities, *International Journal for Numerical Methods in Engineering* 79 (11) (2009) 1309–1331.
- [9] S. Karman, K. Karman-Shoemake, C. Woeber, Mixed order mesh curving, *SEMA SIMAI Springer Series* 30 (2022) 1–21. doi:10.1007/978-3-030-92540-6_1.
- [10] Beta CAE Systems, ANSA v23.1.0 User Guide and Version Release Webinar, <https://www.beta-cae.com/ansa.htm>, high-order meshing capabilities (2025).
- [11] E. Ruiz-Gironés, X. Roca, Generation of curved meshes for the high-lift common research model, in: *AIAA Aviation Forum 2022*, 2022.
- [12] G. Mengaldo, D. Moxey, M. Turner, R. Moura, A. Jassim, M. Taylor, J. Peiró, S. Sherwin, Industry-relevant implicit large-eddy simulation of a high-performance road car via spectral/*hp* element methods, *SIAM Review* 63 (4) (2021) 723–755.
- [13] M. Green, K. Kirilov, M. Turner, J. Marcon, J. Eichstäedt, E. Laughton, C. Cantwell, S. Sherwin, J. Peiro, D. Moxey, Nekmesh: An open-source high-order mesh generation framework, *Computer Physics Communications* (2024).
- [14] K. S. Kirilov, J. Peiró, J. Zhou, M. D. Green, D. Moxey, High-Order Curvilinear Mesh Generation From Third-Party Meshes, *SIAM*, 2024, pp. 93–105. arXiv:<https://epubs.siam.org/doi/pdf/10.1137/1.9781611978001.8>, doi:10.1137/1.9781611978001.8.
URL <https://epubs.siam.org/doi/abs/10.1137/1.9781611978001.8>
- [15] F. Buscariolo, J. Hoessler, D. Moxey, A. Jassim, K. Gouder, J. Basley, Y. Murai, G. Assi, S. Sherwin, Spectral/*hp* element simulation of flow past a formula one front wing: Validation against experiments, *Journal of Wind Engineering and Industrial Aerodynamics* 221 (2022) 104832.

- [16] Siemens Digital Industries Software, Simcenter STAR-CCM+, version 2022.2, <https://www.plm.automation.siemens.com/global/en/products/simcenter/STAR-CCM.html> (2025).
- [17] PADRAM: Parametric Design and Rapid Meshing System for Complex Turbomachinery Configurations, Vol. Volume 8: Turbomachinery, Parts A, B, and C of Turbo Expo: Power for Land, Sea, and Air. doi: 10.1115/GT2012-69030.
- [18] T. Gerhold, J. Neumann, The parallel mesh deformation of the DLR TAU-code, in: C. T. et al. (Ed.), New Results in Numerical and Experimental Fluid Mechanics VI, Springer Berlin Heidelberg, 2008, pp. 162–169.
- [19] OpenFOAM, API guide: ccmToFoam.C file reference, https://www.openfoam.com/documentation/guides/latest/api/ccmToFoam_8C.html, last accessed October 2023.
- [20] Capgemini Engineering, Opencascade technology, <https://www.opencascade.com>, software platform for 3D CAD, CAM, CAE. Last accessed: January 2023 (2023).
- [21] CADfix: CAD translation, healing, repair, and transformation, www.iti-global.com/cadfix (2025).
- [22] J. Marcon, M. Turner, J. Peiró, D. Moxey, C. R. Pollard, H. Bucklow, M. Gammon, High-order curvilinear hybrid mesh generation for CFD simulations, in: 2018 AIAA Aerospace Sciences Meeting, 2018.
- [23] M. Turner, D. Moxey, J. Peiró, Automatic mesh sizing specification of complex three dimensional domains using an octree structure, in: 24th International Meshing Roundtable, 2015.
- [24] S. Chun, J. Marcon, J. Peiró, S. Sherwin, Reducing errors caused by geometrical inaccuracy to solve partial differential equations with moving frames on curvilinear domain, Computer Methods in Applied Mechanics and Engineering 398 (2022) 115261. doi:<https://doi.org/10.1016/j.cma.2022.115261>.

- [25] S. Sherwin, J. Peiró, Mesh generation in curvilinear domains using high-order elements, *International Journal for Numerical Methods in Engineering* 53 (1) (2002) 207–223.
- [26] M. Turner, D. Moxey, J. Peiró, M. Gammon, C. Pollard, H. Bucklow, A framework for the generation of high-order curvilinear hybrid meshes for CFD simulations, in: *Procedia Engineering*, Vol. 203, 2017, pp. 206–218.
- [27] R. Byrd, P. Lu, J. Nocedal, C. Zhu, A limited memory algorithm for bound constrained optimisation, *SIAM Journal on Scientific Computing* 16 (5) (1995) 1190–1208.
- [28] D. Moxey, M. Green, S. Sherwin, J. Peiró, An isoparametric approach to high-order curvilinear boundary-layer meshing, *Computer Methods in Applied Mechanics and Engineering* 283 (2015) 636–650.
- [29] J. Pegrum, Experimental study of the vortex system generated by a Formula 1 front wing, Ph.D. thesis, Imperial College London (2007).
- [30] A. I. Liosi, P. Khurana, A. Swift, A. Chatzopoulos, F. Bottone, J. Hoessler, S. J. Sherwin, Requirements & current capabilities for implicit large eddy simulation of the imperial front wing using spectral h/p element methods, in: *Direct and Large Eddy Simulation IV Proceedings*, 2025.
- [31] A. I. Liosi, H. Wüstenberg, P. Khurana, S. Sherwin, J. Hoessler, A. Swift, A. Chatzopoulos, Comparing efficient implicit time-stepping techniques for highly resolved simulations using industrial geometries, in: *AIAA SCITECH 2025 Forum*, 2025. doi:10.2514/6.2025-1486.
- [32] P. Khurana, S. Sherwin, J. Hoessler, D. Moxey, A. Chatzopoulos, Lower-order refined preconditioning for spectral/hp element methods for complex, 3D geometries, in: *Proceedings of the 9th European Congress on Computational Methods in Applied Sciences and Engineering (ECCOMAS 2024)*.
- [33] A. S. O’Sullivan, K. Puri, E. Ferrer, Best practices and mesh convergence on the extruded imperial front wing with CharLES solver, in: *AIAA Aviation Forum and Ascend 2024*, 2024. doi:10.2514/6.2024-4259.

- [34] G. Ntoukas, G. Rubio, O. Marino, A. Liosi, F. Bottone, J. Hoessler, E. Ferrer, A comparative study of explicit and implicit large eddy simulations using a high-order discontinuous Galerkin solver: Application to a Formula 1 front wing, *Results in Engineering* 25 (2025) 104425. doi:<https://doi.org/10.1016/j.rineng.2025.104425>.
- [35] P. Khurana, A. Liosi, S. Sherwin, J. Hoessler, A. Chatzopoulos, A. Swift, F. Bottone, Industrialization of spectral/hp element method for incompressible, transitional flow around Formula 1 geometries, in: *AIAA SCITECH 2025 Forum*, 2025. doi:[10.2514/6.2025-0880](https://doi.org/10.2514/6.2025-0880).