Shrinkwrap developments for computational electromagnetics in ICE NITe

Preparing CAD models for electromagnetic analysis remains a complex, time consuming process. Typically, the CAD model will contain excessive detail, contain modelling errors such as overlapping solids, or tiny gaps, and be constructed as a complex assembly of many separate parts. Before such a model can be used for computational electromagnetic analysis (CEM), detail must be removed, errors corrected, and contacting solids joined together. This places a significant burden on the CEM analyst, which is especially frustrating in areas of the model where fidelity isn't critical, and all that is required is an approximated geometry.

Shrinkwrapping is a technique for generating a facetted representation of a CAD model, with a fixed resolution, suppressing details, holes, and gaps below this resolution. It's intended as a "one click" method, where the CEM analyst simply sets some basic parameters, and the remainder of the algorithm is automatic. As such, it provides an alternative to conventional CAD pre-processing operations, offering a faster method for preparing CEM-suitable geometry, but at the expense of reduced control over the process.

1. Shortcomings

Shrinkwrapping is not a new concept, and several tools have demonstrated shrinkwrapping techniques in the past. These are based around a level set/octree approach, and have significant shortcomings, which make them difficult to deploy widely for CEM preprocessing:

1. Poor representation of thin parts – especially thin, sharp features, such as wing trailing edges
2. Lack of associativity with the input geometry; there is typically no relationship between the shrinkwrap and the input geometry, aside from geometric proximity
3. Poor reproducibility – rotating or moving a part will result in a radically different result, due to a different alignment with the background grid

The shrinkwrapping tool described in this report uses a radically different foundation to existing tools, which allows it to avoid these issues, producing a shrinkwrapper which is much better suited to the CAD to CEM process.

2. Shrinkwrapping steps

Shrinkwrapping tools generally construct their facetted representations in three steps:

1. Establish a topology for the model, suppressing gaps and holes below the chosen size. If you imagine shrinkwrap generation as a physical process (i.e., shrinking plastic over an object), this step constructs an initial loose plastic sheeting over the model. Note that any holes which are not suppressed require a “handle” of plastic sheeting which runs through the hole.
2. Construct an initial faceting from this topology, which represents the geometry of the shrinkwrap. This is the “shrink” step in our physical analogy, where heat/vacuum is used to shrink the plastic onto the object. Whilst this faceting represents the geometry of the shrinkwrap to a required tolerance, it will not be of sufficient quality to use for CEM – the triangles will have high aspect ratios, short edges, and the overall triangle count will not be optimal.
3. Re-mesh and smooth this faceting to make the final shrinkwrap. This takes the initial coarse faceting, and converts it into a form suitable for CEM, with an optimum number of good quality triangles.

3. Improved shrinkwrap topology

Typical shrinkwrap processes start by voxelising the input geometry. A regular grid, with a size dictated by the shrinkwrap resolution, is overlaid on the geometry, and the cells which intersect the CAD solids are kept. This collection of cells then determines the topology of the shrinkwrap (how many holes the shrinkwrap should represent) together with a first approximation to the shrinkwrap geometry.

The voxelisation process produces a couple of problems:

- The alignment between the input geometry and the regular grid significantly affects the result; rotating or offsetting the grid will alter the resulting shrinkwrap topology. This is often seen in the output when shrinkwrapping a component with repeated or symmetric features, when the shrinkwrap fails to process identical parts of the model in the same way.
- There is only limited associativity between the input CAD model and the voxels. Each voxel can determine which CAD geometry it contains, but this is not enough to uniquely carry properties across from the CAD model onto the shrinkwrap.

Within ICE NITe, we have evaluated an alternative method for determining shrinkwrap topology, based around the concept of an alpha shape. The alpha shape is, loosely speaking, the shape that you get by rolling a sphere of radius alpha over the outside of a CAD model. All of the places which are touched by the ball are part of the alpha shape, and the remaining areas are deleted. Varying the alpha value alters the level of detail retained by the alpha shape. The alpha shape is a generalisation of a convex hull; in the case where alpha tends to infinity, a convex hull is produced. The concept is illustrated schematically in Figure 1.

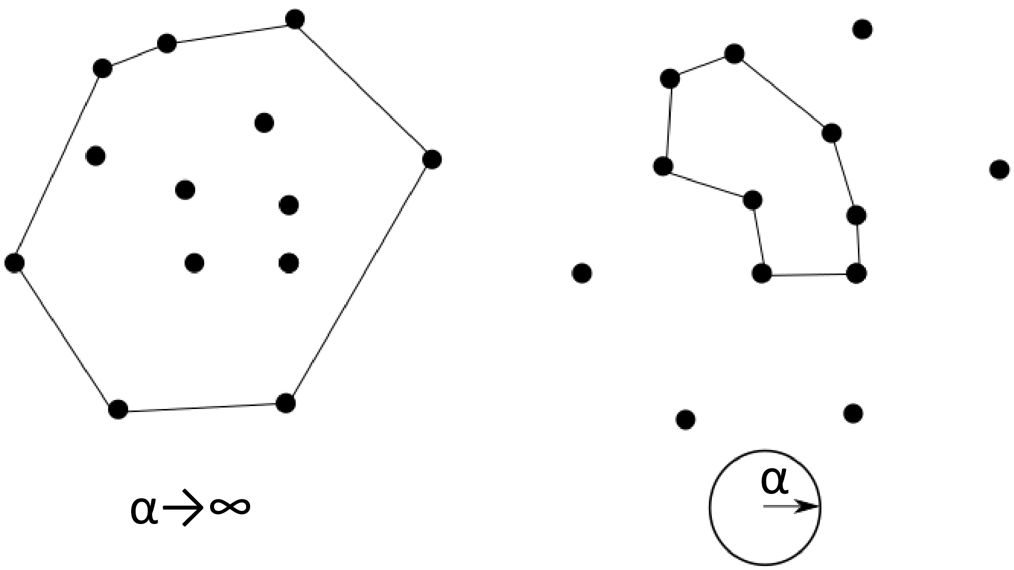

Figure 1: A two-dimensional alpha shape (black lines) of a set of points with alpha → ∞ (left), and a finite value of alpha (right).

The alpha shape approach provides some significant benefits:

- The alpha shape of a CAD geometry is uniquely defined, regardless of orientation, leading to a more stable shrinkwrap process
- The alpha shape calculation re-uses sample points directly taken from the CAD model, allowing CAD properties to be carried directly through onto the shrinkwrap

Published methods for alpha shape calculation use a dense point sampling of the input object. This does not scale to the large, complex assemblies used by the project partners, and so within ICE NITe, ITI have developed an adaptive alpha shape method, using sparse sampling over much of the object, and adding sample density where needed.

4. Better shrinkwrap geometry

Once the shrinkwrap topology has been determined, the next step is to generate appropriate geometry. When following a voxelisation approach, this is typically performed using a contouring procedure, such as marching cubes, or dual contouring. Reproduction of thin, sharp features such as trailing edges of wings can be very difficult to get right using these methods. As when computing the shrinkwrap topology, geometry computed from a voxelisation is subject to aliasing artefacts, such as ripples in the geometry caused by the particular alignment between the geometry and the voxels.

When following the alpha shape method, the output from the alpha shape calculation is a closed, connected set of triangles, whose vertices are samples directly from the input geometry. This makes the alpha shape suitable for use as shrinkwrap geometry directly.

In concave areas, especially 90 degree concave corners, the alpha shape contains unwanted "webs" which smooth off the sharp corners. Various techniques have been explored to remove these artefacts, and the most successful is the "massage splitting" technique, which occurs after the associativity phase, and is described below.

5. Associativity

A significant criticism of conventional shrinkwrap methods is that the resulting faceting only has a loose connection back to the original CAD model. Often, structures such as scratches on the CAD model are lost completely, and data held on the CAD model, such as boundary conditions and material properties, cannot be propagated to the shrinkwrap.

The alpha shape approach brings significant advantages for associativity. Because the alpha shape is constructed using sample points from the original CAD model, these points can be used to carry associativity information directly into the shrinkwrap. It's also much easier to ensure that lines and scratches in the CAD model carry through to the resulting shrinkwrap.

Additional work has been carried out in ICE NITe to improve the associativity between a shrinkwrap and the original CAD model, using a re-parenting technique. For many models, much of the shrinkwrap will closely reproduce the original CAD model. In these areas, it's possible to find a correspondence between the shrinkwrap facets and the input geometry. We have developed "re-parenting" techniques to do this, and tag the shrinkwrap facets with parent face information where a direct correspondence is possible. We can also identify the common situations where a shrinkwrap facet spans two parent faces (for example, across an edge in the CAD model which has no geometric significance), and split the shrinkwrap facet, allowing the two halves to be successfully reparented.

6. Splitting and smoothing

In "un-parented" regions of the shrinkwrap, where a direct correspondence to CAD is not immediately possible, we attempt to improve it, to bring it closer to the input geometry. For example, a common issue with the alpha shape is that it will generate facets which round off sharp, concave, 90 degree corners. An approach we call "massage splitting" is deployed here to improve these situations. The incorrect facets are identified, split along their centerline, and projected onto the correct, sharp, intersection line. A round of smoothing, based on a mesh massaging technique[1] is then deployed across all un-parented facets, to improve the shape and fit to the input geometry.

---

[1] T. Winkler, K. Hormann, C. Gotsman. Mesh massage: A versatile mesh optimization framework. *The Visual Computer, 24(7-9):*775-785. July 2008

We then run the re-parenting step again, as the modified facets will often be very close to the input CAD geometry after massage splitting. An example of this is shown in Figure 2.

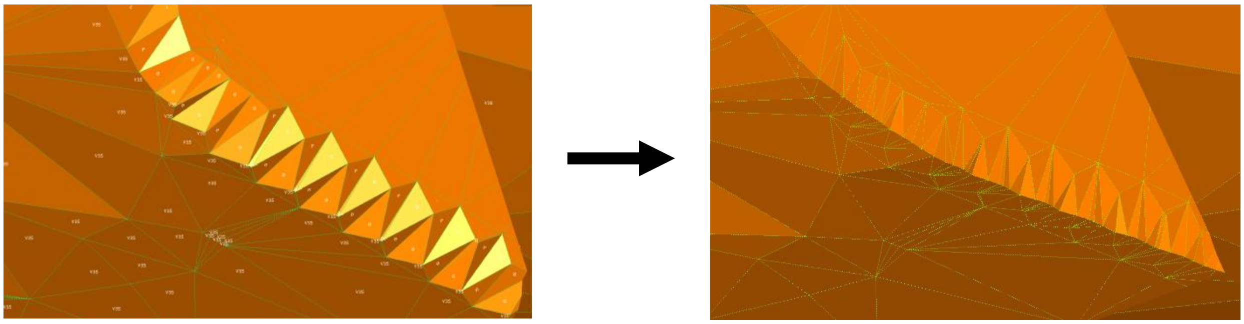

Figure 2: An example of "massage splitting". On the left, a wing/fuselage junction is spanned by a "web" of facets. After the smoothing process, the facets are split and projected onto the correct intersection line.

7. Re-meshing

The facets that are generated by shrinkwrap procedures – both by voxelisation-based approaches, and by our new alpha-shape one – are typically not suitable for direct use in simulation. The facets may have undesirable angles and aspect ratios, and have no control over facet size, being too dense in some areas, and too sparse in others. Re-meshing and smoothing techniques can be employed to improve the facet sizes and shapes, and make them suitable for simulation.

We have deployed a re-meshing technique which relies on the re-parenting work above. Regions of the shrinkwrap which are associated with a single CAD face can be simply re-meshed by inscribing the boundary of the region onto the embedding geometry for the CAD face, and using a standard CAD meshing algorithm to generate a suitable mesh over this region. The re-meshed regions can then be stitched back together to create a high quality re-meshed faceting over the whole model. A selection of Delaunay triangulation algorithms has been made available for the re-meshing step: minimal curvature-sensitive triangulation, quality triangulation with interior refinement, and quality triangulation with curvature-sensitive refinement.

8. Results

The CADfix alpha shrinkwrapper was run on a selection of models. Here, we present three examples (Figures 3, 4 and 5). Because the original CAD geometry of each model was watertight, an alpha value of zero was chosen without risk of the wrap “leaking” inside the model through any holes. The re-meshing algorithm used was a quality triangulation with curvature-sensitive refinement.

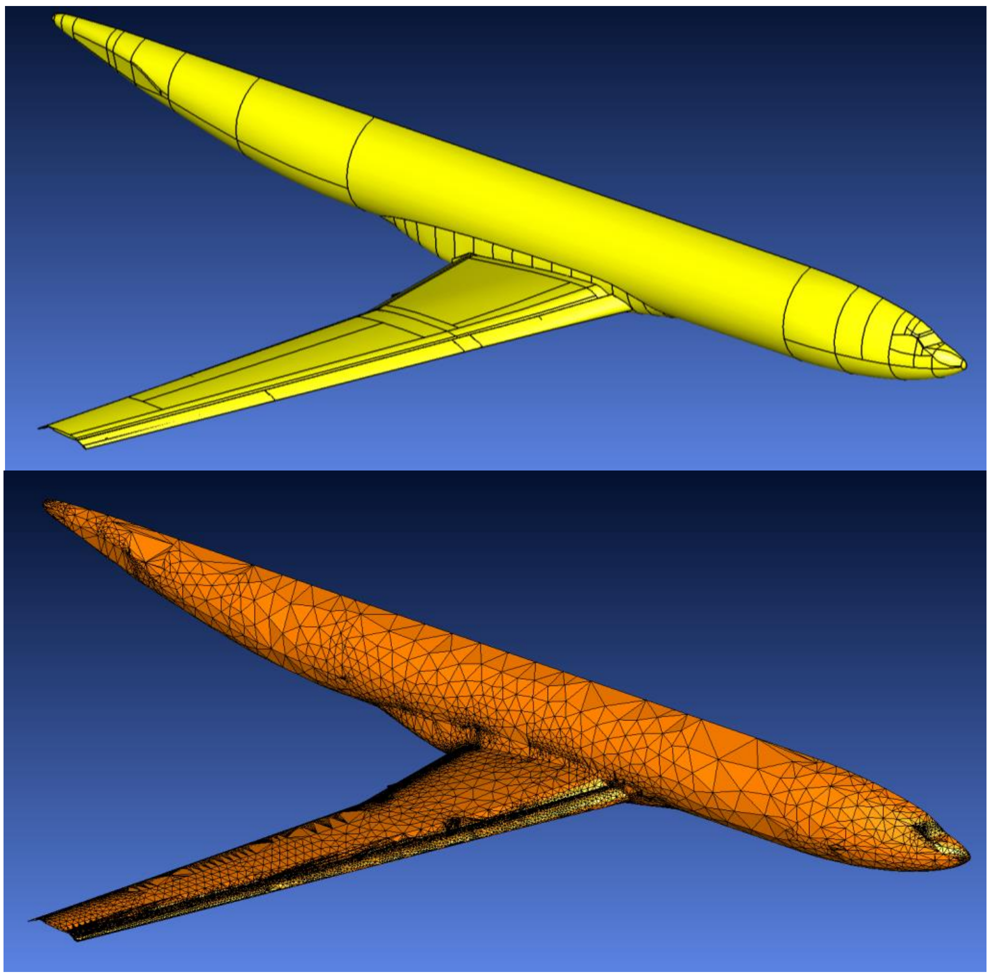

Figure 3: A CAD model (top) and its shrinkwrap (bottom).

The alpha shrinkwrapper’s ability to preserve thin, sharp edges (such as the trailing edges on the wings and tails) is clearly demonstrated in each case. As well as this, the “massage splitting” procedure ensures that the junctions between wings and the fuselage (for example) are well preserved. A large proportion of the wrap has successfully re-meshed; this is thanks to the one-to-one correspondence between the shrinkwrap facets and the input geometry (re-parenting).

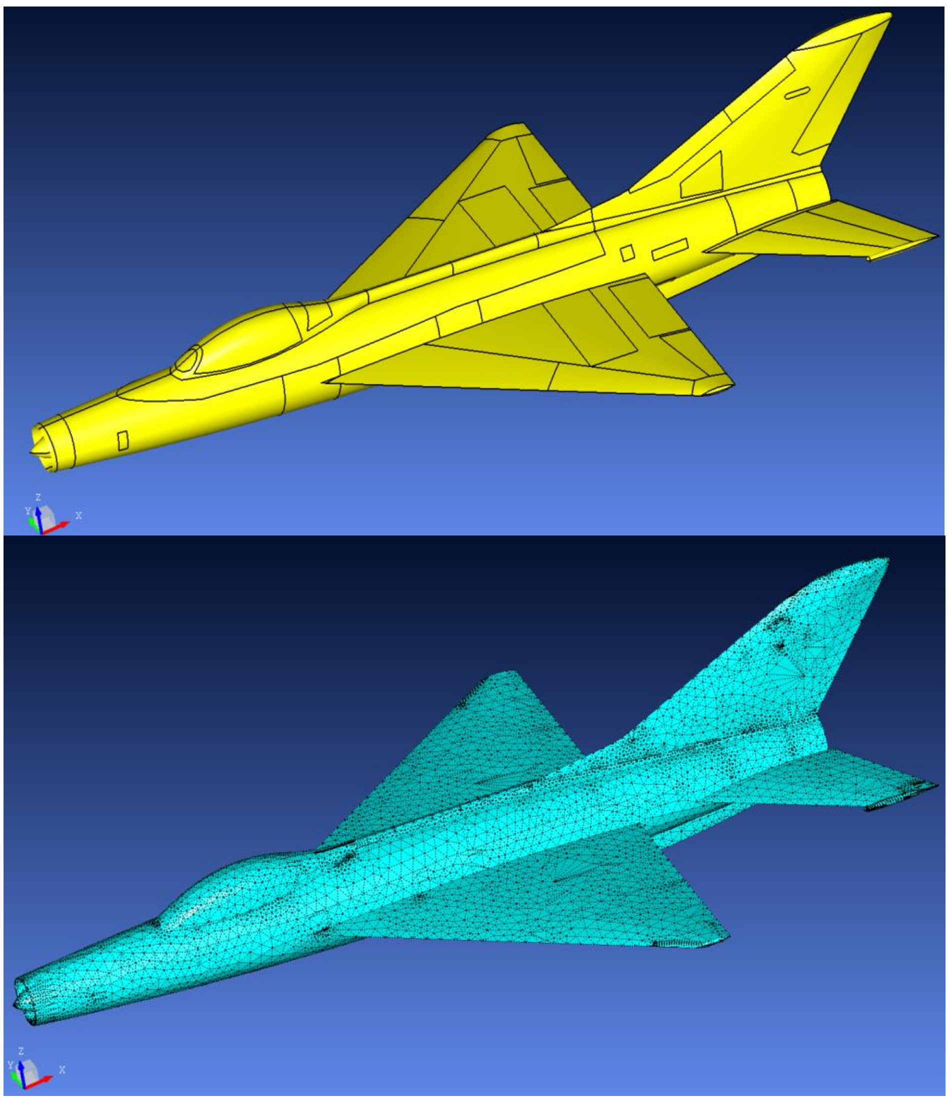

Figure 4: A CAD model of a MiG-21 (top) and its shrinkwrap (bottom).

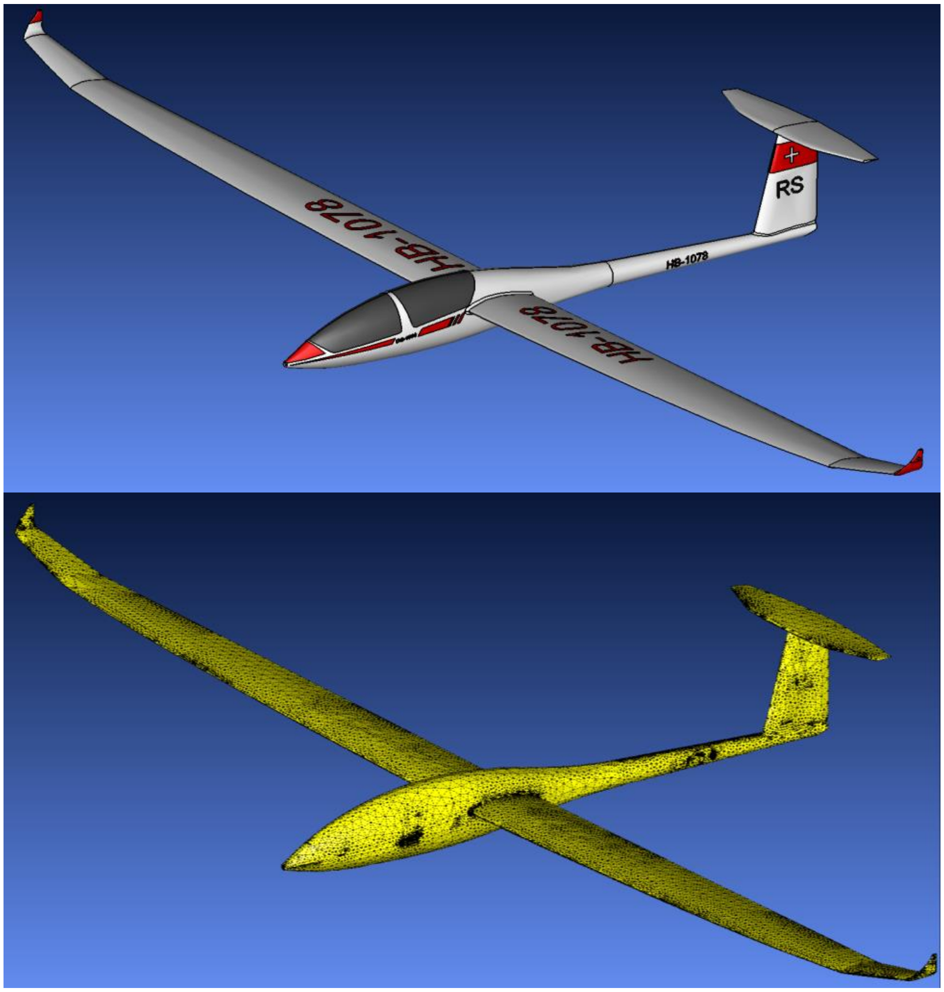

Figure 5: A CAD model of a glider (top), and its shrinkwrap (bottom).

9. Conclusions

This report has detailed the advances in shrinkwrap generation which have been researched in the ICE NITe project. A completely new shrinkwrap method has been researched, developed, and implemented inside the CADfix platform, based around the use of alpha shapes, CAD associativity, and re-meshing using the CAD surfaces. This tool will prepare a faceting of CAD model which ignores small details, holes, and gaps, using a single click, and can be controlled by two straightforward parameters. Compared to previous approaches, the new shrinkwrap algorithm is invariant under rotation and translation of the input CAD model. It can better approximate sharp edges and thin-walled objects, has strong associativity to the CAD model, and provides a result with high-quality facets.