Masters, D., Taylor, N., Rendall, T., & Allen, C. (2017). Three-Dimensional Subdivision Parameterisation for Aerodynamic Shape Optimisation. In *55th AIAA Aerospace Sciences Meeting* [AIAA 2017-0036] American Institute of Aeronautics and Astronautics Inc. (AIAA). <https://doi.org/10.2514/6.2017-0036>

Peer reviewed version

Link to published version (if available):
[10.2514/6.2017-0036](https://doi.org/10.2514/6.2017-0036)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via AIAA at <http://arc.aiaa.org/doi/10.2514/6.2017-0036>. Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Three-Dimensional Subdivision Parameterisation for Aerodynamic Shape Optimisation

D. A. Masters*,

Department of Aerospace Engineering, University of Bristol

N. J. Taylor[†],

MBDA UK Ltd, Filton

T. C. S. Rendall[‡] and C. B. Allen[§]

Department of Aerospace Engineering, University of Bristol

A novel hierarchical wing parameterisation method based on subdivision surfaces is presented and its performance tested on a range of geometric and aerodynamic optimisation test cases. Subdivision surfaces form a limit surface based on the recursive refinement of an initial network of points. This intrinsically creates a hierarchy of control points that can be used to deform the surface at varying degrees of fidelity. This principle is used to create a multi-resolutional surface parameterisation that can make fine and gross surface changes without losing underlying surface detail. This is then extended to allow multi-resolutional control of arbitrary meshes such as computational surface grids. This parameterisation method is then applied to a range of optimisation problems in a ‘multi-level’ procedure that starts with a low fidelity parameterisation and which is then increased sequentially. These cases are compared against a range of ‘single-level’ schemes that use each level in isolation. It was found that by using the multi-level method significant improvements to both convergence rates and robustness were achieved. In some cases this increased robustness lead to improved final results by successfully exploiting high dimensional design spaces that could not be explored using a fixed number of design variables.

I. Introduction and Background

With optimisation becoming more common in aerodynamic design, a significant effort is being made to improve both its effectiveness and its efficiency. Within an optimisation procedure, the choice of shape parameterisation controls the relationship between the optimisation design variables and the aerodynamic surface itself. Consequently the choice of shape parameterisation method can have a significant impact on the effectiveness and efficiency of the overall procedure; this has been subject to much investigation^{1–5}. Many different methods have been used within an aerodynamic optimisation framework though all are subject to the ‘curse of dimensionality’ in some respect. In the context of aerodynamic optimisation this refers to the problems associated with increasing the number of design variables used in the optimisation procedure. For many optimisation schemes the number of objective function evaluations is proportional to the number of design variables used, in conjunction with this a large number of design variables can lead to poor convergence rates and poor design space conditioning. Considering that for aerodynamic optimisation each objective function evaluation can equate to a single, often expensive, aerodynamic computation, the impact of dimensionality can be huge. On the other hand, the fidelity of the parameterisation, and therefore the design space of the problem, is directly linked to the number of design variables. This often leads to a compromise between available resources and desired accuracy of the results.

One approach to reducing this effect is to control the shape with a series of nested, hierarchical parameterisation schemes and increase the fidelity at intervals throughout the optimisation process. This approach is akin to a multi-grid method for parameterisation and was first used in an aerodynamic optimisation setting by Beux and Dervieux⁶. It has since been applied to a range of aerofoil optimisation problems using a variety of different parameterisation frameworks such as Bèzier curves^{7–10}, Bèzier surface FFD^{11–14}, RBFs¹⁵ and B-Splines with a knot insertion algorithm^{16,17}.

*Graduate Student, AIAA Student Member, dominic.masters@bristol.ac.uk, Bristol, BS8 1TR, UK

[†]Capability Leader, Aerodynamic Tools & Methods, AIAA Associate Fellow, nigel.j.taylor@mbda-systems.com, WG3, PO Box 5, Filton, Bristol, BS34 7QW, UK

[‡]Lecturer, AIAA Member, thomas.rendall@bristol.ac.uk, Bristol, BS8 1TR, UK

[§]Professor of Computational Aerodynamics, AIAA Senior Member, c.b.allen@bristol.ac.uk, Bristol, BS8 1TR, UK

In general, work has shown that implementation of these multi-level nested parameterisations can improve the convergence rate, robustness and final result of an optimisation procedure. This paper investigates the application of multi-level shape parameterisation techniques to multi-resolution subdivision surfaces for aerodynamic optimisation procedures.

Subdivision surfaces are a shape parameterisation method used predominantly in computer graphics and animation. They describe a smooth surface based on an initial coarse network of points and a simple subdivision rule of refinement. By successively applying the subdivision rule increasingly fine networks are created which at the refinement limit create a surface. In some cases these limit surfaces are equivalent to B-spline surfaces. For example for a set of regular, rectangular control points Doo-Sabin subdivision²⁰ is equivalent to bi-quadratic B-spline surfaces and Catmull-Clark subdivision²¹ is equivalent to uniform bi-cubic B-spline surfaces. However, subdivision surfaces are generalisable to arbitrary topologies while B-Spline surfaces are not; this is a key benefit of the subdivision framework. In addition, although subdivision curves and surfaces typically create smooth results, creases, corners and cusps can easily be introduced. B-Splines and subdivisions share many characteristics; the method of implementation is, however, one area of difference. B-splines utilise continuous parametric representation whereas subdivisions use a hierarchical process of discrete refinement. It is this innately hierarchical nature of subdivisions that make them easily applicable to multi-resolution analysis.

Multi-resolution analysis utilises hierarchical nested data sets to increase efficiency by allowing operations to be performed at varying levels of detail. For geometry and shape parameterisation applications this typically means the ability to implement coarse geometry changes while maintaining the fine detail, and fine geometry changes while maintaining the overall shape. This approach has been implemented comparably from both B-spline²² and subdivision²³ perspectives in both their two-dimensional and three-dimensional forms. It would seem, however, that in three dimensions the advantage of being able to represent arbitrary topologies with subdivision surfaces has lead to it becoming the industry standard choice in multi-resolution computer animation²⁴. It also appears to be slowly being incorporated into some mechanical computer aided design (MCAD) packages²⁵⁻²⁷.

Previous work from the authors has investigated the benefits of using a hierarchical subdivision parameterisation scheme for two-dimensional aerodynamic optimisation^{4,5}. This used a low fidelity parameterisation for the early optimisation stages then periodically refined it to increase the number of design variables. This was compared to equivalent, non-refining, single-level subdivision curves equivalent to using cubic B-Splines. It showed that significant benefits to robustness and convergence rates are achievable. In particular the multi-level subdivision schemes consistently attained the best results which the high-fidelity single-level methods often failed to reach.

The aim of this work is to expand on this previous research and extend the theory to three-dimensional subdivision surfaces. The method is then applied to a range of geometric and aerodynamic test cases.

II. Multi-resolution Subdivision Surfaces

A. Subdivision Refinement

Subdivision surfaces are defined as the limit of a process of repeated subdivision refinement of an initial control polygon. Each subdivision refinement defines a new set of smoother, denser points as a weighted average of the old points. For this reason the refinements can always be expressed as a matrix transformation

$$C_{n+1} = P_n C_n \quad (1)$$

from old points C_n to new points C_{n+1} . A variety of different schemes exist that work primarily with triangles^{28,29} or quads^{20,21} and either approximate^{20,21,28} or interpolate²⁹ the initial control mesh. This work concentrates on the variation proposed by Catmull and Clark²¹ which is an approximating quad subdivision scheme. This method works by inserting vertices at the centre of each edge and face then applying an averaging rule to smooth their positions. Figure 1 shows an example of this process. It should also be noted that sharp edges and corners can be easily integrated into these subdivision refinement schemes by applying a different rule in specified areas²⁴, for example, the red edges in figure 1 are labelled as sharp. Figure 2 then shows the subdivision of a cube.

Given a set of subdivision matrices P the N^{th} subdivision level C_N can then be expressed as

$$C_N = P_{N-1} \dots P_{n+1} P_n C_n \quad (2)$$

for some $n < N$. The limit surface can therefore be described as

$$C_\infty = \lim_{N \rightarrow \infty} C_N = \dots P_{n+1} P_n C_n. \quad (3)$$

Figure 1. Example of subdivision refinement process for the Catmull-Clark scheme

Figure 2. Subdivision refinement of a cube using the Catmull-Clark scheme

In practice this calculation must be truncated at some point and the limit surface calculated. A typical solution to this is to calculate an evaluation matrix (denoted P_N^{Eval}) based on eigenanalysis³⁰ that pushes the subdivision points to their limit locations. It should be noted that this is always a square invertible matrix. The notation C_∞ will, from here on, be used to denote the limit surface itself and the truncation level will be assumed to be N , the maximum subdivision level used. It is then convenient to define

$$\phi_n = P_N^{Eval} P_{N-1} \dots P_n \quad (4)$$

and therefore the relationship between any level of control polygon and the limit surface can be expressed as

$$C_\infty = \phi_n C_n. \quad (5)$$

This relationship then allows each control point polygon to manipulate the limit surface at different levels of fidelity. Figure 3 shows an example of this using the limit surface derived from a cube where each control polygon is equivalent to the coarse surfaces shown in figure 2. Comparable deformations are then made using the different subdivision levels and it can be seen that the fidelity of the control increases with the density of the control polygons. This hierarchy of control, along with the invariance of the limit surface to the refinement of the polygons, creates a system of unidirectional multi-resolutional control of the surface. This means that changes can be made at a coarse level, then refined to increase the fidelity, while keeping the limit surface constant. As a consequence of this, changes can be made at a fine level while keeping the gross shape of the surface. This does not, however, imply that coarse changes can be made while maintaining the fine detail; this is addressed in the next section.

Figure 3. Example deformation of subdivision limit surface.

B. Reverse Subdivision

To achieve surface invariance under the coarsening of control polygons, and consequently enable bidirectional multi-resolutional control, further modifications are required. To illustrate this, consider a fine set of points C_n then the coarser set of points C_{n-1} can simply be calculated as the least squares solution of equation 1,

$$C_{n-1} = P_{n-1}^+ C_n, \quad (6)$$

where $+$ denotes the Moore-Penrose pseudo-inverse. However as P_n is a non-square, non-invertible matrix this leads to some loss of information with the result that, for almost all cases

$$P_{n-1} C_{n-1} \neq C_n. \quad (7)$$

For this reason it is important to retain any errors created through the least squares process and include them in any subsequent refinement. This can be done very conveniently and efficiently if P_n has full column rank by extending the refinement matrices P_n by any orthogonal compliment $Q_n = \text{null}(P_n^T)$. The subdivision refinement (equation 1) can then be reformed such that

$$\begin{aligned} C_{n+1} &= \begin{bmatrix} P_n & Q_n \end{bmatrix} \begin{bmatrix} C_n \\ D_n \end{bmatrix} \\ &= P_n C_n + Q_n D_n, \end{aligned} \quad (8)$$

for some set of error coefficients D_n . By letting

$$\begin{bmatrix} P_n & Q_n \end{bmatrix}^{-1} = \begin{bmatrix} A_n \\ B_n \end{bmatrix} \quad (9)$$

equation 6 can be re-expressed as the reverse subdivision equations

$$C_{n-1} = A_n C_n, \quad (10)$$

$$D_{n-1} = B_n C_n. \quad (11)$$

Given equation 8, a two level refinement can be expressed as

$$C_{n+2} = P_{n+1} (P_n C_n + Q_n D_n) + Q_{n+1} D_{n+1}, \quad (12)$$

which can then be applied recursively such that

$$C_\infty = P_N^{Eval} (P_{N-1} (\cdots (P_n C_n + Q_n D_n) + \cdots) + Q_{N-1} D_{N-1}). \quad (13)$$

This can then be simplified to the final subdivision refinement equation

$$C_\infty = \phi_n C_n + \sum_{i=n}^{N-1} \phi_{i+1} Q_i D_i. \quad (14)$$

Importantly this creates a one-to-one relationship between the subdivision refinement levels, thus creating a bidirectional multi-resolutional system. This allows information to be propagated uniquely and exactly in either the refinement (figure 4) or coarsening (figure 5) direction.

Figure 4. Process for subdivision refinement.

Figure 5. Process for reverse subdivision.

There are two key consequences of this:

- Coarse changes can be made while maintaining fine detail

- The subdivision system can be initialised for any compatible limit surface

Figure 6 shows the deformation of a limit surface with some arbitrary fine detail. Here the control polygons and error coefficients are calculated from the limit surface using the reverse subdivision equations (10 and 11). Interestingly it can be seen that the control polygons C_1 and C_2 are almost identical to those shown in figure 3 while C_3 is significantly different. This is because in this case the surface roughness is on a scale that cannot be captured by the first and second level control polygons. As a result the roughness information is constrained to the error coefficients and is invariant to any movement of the control polygon. This can be seen as the surface roughness is maintained under the deformation of polygons C_1 and C_2 .

Figure 6. Example deformation of subdivision limit surface with fine detail.

III. Multi-resolutional Control of Arbitrary Topology Meshes

The method described in the previous section is limited by the fact that the surface mesh is restricted to match the subdivision surface topology. In this section a method for extending the multi-resolution theory to arbitrary surface meshes is described. This allows the variable fidelity control of subdivision surfaces to be applied to arbitrary surfaces, such as triangulations or computational meshes.

A. Continuous Parameterisation of Subdivision Limit Surface

A key building block of this method is the ability to continuously parameterise the limit of a subdivision surface. This work follows the method outlined in Stam³¹. To do this it is important to first differentiate between regular faces, that have exactly four regular vertices (valence four), and extraordinary faces, that do not. This is key because regular faces can be reduced to a bi-cubic B-spline patch which makes their parametrisation simple; extraordinary patches on the other hand do not have a direct B-spline representation.

Figure 7 shows an example of a B-spline representation of a regular face both with and without sharp edges. For faces with no sharp edges a set of 16 B-spline control points are constructed from the four face vertices plus the surrounding 12 vertices as shown in figure 7a. When an edge face is parameterised a bzier end condition must also be applied. This is done by introducing a set of ghost control points as shown in figure 7b. These surfaces are each parameterised by $[u, v] \in [0, 1] \times [0, 1]$ allowing a full and continuous representation of any regular face.

Figure 7. Continuous parameterisation of regular patches. The grey surface with black lines represent a subdivision surface and control polygon; the magenta surface and red points represent the continuously parameterised B-spline patch and associated control points.

For extraordinary faces a single continuous B-spline patch cannot be constructed due to difficulties around extraordinary vertices. This stems from the fact that B-spline surfaces are only defined for regular rectangular patches thus extraordinary vertices are incongruent with this approach. This influence only affects surface regions directly adjacent to an extraordinary vertex; consequently the regions that cannot be parameterised with standard bi-cubic patches decreases with every refinement. Figures 8a-8c show an example of this around a valence five vertex. A naïve approach to this would be to subdivide the geometry until the extraordinary patches are sufficiently small, this however would result in an extremely fine discretisation of the whole surface at a significant computation cost. A more efficient approach would be to apply local refinement scheme that only refines the area around an extraordinary valence as shown in figure 8d, this method still requires significant computational cost relative to the parameterisation of a regular face. This is however solved by the approach outlined by Stam³¹. He recognised that the operation of locally refining around an extraordinary vertex (figure 8d) can be described by the repeated application of a single local subdivision matrix. This means that it is therefore open to using eigenanalysis to understand its limit behaviour. The eigenvalues of this local subdivision matrix are then used to scale a set of B-spline basis functions to continuously parameterise the extraordinary faces by continuous variables $[u, v] \in [0, 1] \times [0, 1]$ as required.

Figure 8. Parameterisation of subdivision surface around an extraordinary vertex. Extraordinary patches (green) cannot be parameterised using a single B-spline patch.

B. Matching Subdivision Limit Surfaces to Arbitrary Meshes

Now that the full subdivision surface can be continuously parameterised a matrix transformation can be formed from the final subdivision surface to an arbitrary set of points on the surface. This however requires each point to be mapped to a subdivision face with accompanying parameterisation $[u, v] \in [0, 1] \times [0, 1]$ defining its position on the face. This section describes a surface matching technique used to calculate this information for every point on an arbitrary surface mesh.

The premise is to ‘shrink wrap’ the subdivision surface onto the surface mesh in the normal direction, this does however require that the subdivision limit surface and the surface mesh are reasonably closely matched from the start; this is assumed throughout this process. Given a subdivision limit surface C_∞ and a surface mesh M , a normal projection of the subdivision surface on to the surface mesh C^{Mesh} can be calculated. Figure 9 shows a graphical representation of this process. The nearest subdivision face on C^{Mesh} for each mesh point M can then be found as well as the continuous parameters $[u, v] \in [0, 1] \times [0, 1]$ to evaluate the subdivision surface at its closest point. This can then be formulated into a matrix transformation P_N^{Mesh} from subdivision level $N - 1$ to the mesh M . Though this matrix can be used to directly replace P_N^{Eval} , P_N^{Mesh} is not guaranteed to have full column rank inhibiting the calculation of A_N^{Mesh} and B_N^{Mesh} (equation 9). For this reason P_N^{Mesh} must be appended to P_N^{Eval} such that

$$P_N^{EM} = \begin{bmatrix} P_N^{Eval} \\ P_N^{Mesh} \end{bmatrix}. \quad (15)$$

This allow a final surface topology congruent to the surface mesh to be calculated and a full multi-resolutional system to be defined.

IV. Wing Parameterisation

The following section outlines how the multi-resolution subdivision surfaces described above are applied to a range of wing geometries. For each case a full set of subdivision matrices (P_n , Q_n , A_n and B_n) are created to match the desired topology of the surface, though two approaches are used depending on the desired structure of the limit surface.

If the objective function can be calculated on the standard subdivision limit surface then the wing is constructed to match the required surface topology. This was done by splining the target wing based on aerofoil sectional data, then

Figure 9. Procedure for matching subdivision surface (blue) to arbitrary surface mesh (red)

evaluating it with the chord-wise and span-wise distributions of the subdivision surface. This creates the limit surface C_∞ which can then be used to calculate the control point positions C_n and error coefficients D_n using the reverse subdivision equations (10 and 11).

If the objective function is calculated on an arbitrary surface mesh then the subdivision system itself must be altered to match the surface. For this an initial surface approximation is made, then the final refinement matrix P^{Eval} is extended to P^{EM} by the method described in section III. Matrices Q^{EM} , A^{EM} and B^{EM} can then be formed and control point positions C_n and error coefficients D_n calculated. This allows all of the benefits of multi-resolutional control but with added access to an arbitrary surface meshes.

In this work two different surface topologies are considered: a tube-like wing with an open tip (figure 10) and a wing with a closed tip (figure 11). The hierarchical control is then applied in two ways throughout an optimisation procedure. As a ‘multi-level’ method, this starts an optimisation using a coarse level of control and then increases the fidelity sequentially. The benefit of this is that in the early stages large coarse changes can be made very easily aiding the speed and convergence of the optimisation. In the latter stages fine changes can easily be made to ensure that the best result is achieved.

It is also applied in a ‘single-level’ configuration where no refinement takes place. This is akin to a more traditional parameterisation method and for rectangular topology is equivalent to using B-Spline surfaces. This is used by the authors as a control group to compare the multi-level techniques against.

Figure 10. Subdivision parameterisation of un-swept NACA4410 wing with open tip.

V. Optimisation Methodology

In this work subdivision surface parameterisation have been applied to a range of geometric and aerodynamic optimisation problems. For all of these tests the multi-purpose large-scale optimiser SNOPT³² has been used. This is a gradient-based sequential-quadratic programming (SQP) method that employs a reduced-Hessian BFGS search-direction and, in this work, a non-derivative line-search technique.

Optimiser convergence was determined based on the activation of one of three criteria:

1. The Karush-Kuhn-Tucker (KKT) first-order optimality condition³² satisfying a tolerance level
2. The optimiser unable to improve the objective function

Figure 11. Subdivision parameterisation of ONERA M6 wing with closed tip and arbitrary surface mesh.

3. For the multi-level schemes *not* on the final level, satisfying the refinement condition (equation 16)

The refinement condition aims to trigger the refinement of the subdivision scheme when the optimisation has exploited most of the available gains from the current design space and is approaching the minimum. This moment can be difficult to identify as it is very hard to differentiate between the optimiser converging to a local optimum and the optimiser traversing a difficult area of the design space. If refinement is triggered too early the under-exploitation of design space can result in slower convergence and possibly a poorer final result, and if it is triggered too late, over-exploitation of the design space can waste resources. A method for approximating the optimum refinement time was proposed by Anderson¹⁵ where refinement was triggered when the convergence of the objective function with respect to the iterations dropped below some proportion $t < 1$ of the maximum attained. A slightly modified version of this trigger was implemented in Masters *et al.*⁴ and is used in this work. This triggers refinement if the rolling average of the slope of the objective function (and constraints) is less than a proportion t of their max rolling average; i.e.

$$\left| \frac{1}{w} \sum_{j=0}^{w-1} G_{k-j} \right| < \left| \max_{m \leq l \leq k} \left(\frac{t}{m} \sum_{j=0}^{m-1} G_{l-j} \right) \right|, \quad \forall G \in \{G^{obj}, G^{con}\} \quad (16)$$

where $G_i^{obj} = \log_{10}(J_{i-1}) - \log_{10}(J_i)$
and $G_i^{con} = \xi_{i-1} - \xi_i$

for objective function J with constraints ξ , iteration k (at current refinement level) and parameters t , w and m . It should be noted that the slope of the objective is calculated on a log scale. This was found to work better than a linear scale as the convergence of the objective functions considered in this work are typically logarithmic.

The parameter $0 < t < 1$ controls the change in gradient required to trigger the scheme and the positive integers w and m control the size of the rolling average windows for the maximum and current slope. If small values are used for w and m this defines a very aggressive triggering system, for well behaved, consistently converging optimisations this ensures that iterations are not wasted converging areas close to a local minimum. For more complex optimisation procedures this can however cause premature triggering when the optimiser only makes a small improvement through a highly non-linear area. For this reason these parameters can be increased to average the gradients and only trigger refinement when improvements are consistently small. For this work the parameters $m = w = 3$ and $t = 0.1$ are used throughout.

VI. Results

To test the effects of using this multi-level subdivision parameterisation scheme a range of optimisations of varying complexity have been carried out. A geometric shape matching optimisation represents the least complex of these where there is a linear relationship between the design variables and the objective and the gradients can be calculated analytically. A series of inverse design studies were then performed that aim to match the surface pressure of the parameterised wing to a target distribution. Finally a fully constrained drag optimisation case is investigated.

A. Geometric Shape Matching

For this optimisation a subdivision wing is created to match a prescribed initial geometry and is then matched to a ‘target’ wing geometry (X^{target}) through an optimisation. Importantly the surface discretization of the target geometry

is also created to match that of the initial subdivision wing. The objective function can then be defined as

$$J(\gamma) = \sqrt{\frac{1}{n} \sum_{i=1}^n \|C_{\infty}(\gamma)_i - X_i^{target}\|^2}. \quad (17)$$

for $C_{\infty}(\gamma), X^{target} \in \mathbb{R}^{n \times 3}$ and design variables γ . The design variable gradients for this function are calculated analytically.

For both of the cases investigated a simple, open-tipped subdivision topology is used (as shown in figure 10) with 12 control points on the first level and 1632 control points on the fifth level, the finest level used for parameterisation. The final surface then represents one further refinement (and evaluation at limit) producing a limit surface with 6336 points.

Each of the subdivision control points is then free to move in all three coordinate directions meaning the number of design variables is triple the number of control points at any stage; this therefore varies between 36 and 4896.

A NACA0012 wing with unit square planform is used as the initial geometry for both cases with two target geometries; a NACA4410 wing with unit chord, span of three, a taper ratio of 0.6 and leading edge sweep angle of 18.4° and an ONERA M6 as described in Schmitt³³.

The convergence of the objective function for these two cases is presented in figure 12. They show that for the single level cases the optima achieved decrease monotonically for increases in the number of design variables, they do however show a decrease in convergence rate. The multi-level cases achieve final results equivalent to the highest fidelity single-level results though with significantly better convergence and approximately a third of the number of iterations.

Figure 12. Objective function convergence for the geometry matching optimisations.

B. Inverse Design

A set of inverse design optimisation have also been performed. These take an initial subdivision wing and target defined with equivalent surface distributions and simple open tipped tube topology, as before, then minimises the pressure difference at each surface point. The surface pressures are calculated using an incompressible, inviscid Hess-Smith panel method and the objective function is defined as

$$J(\gamma) = \sum_{i=1}^{N_{panels}} \frac{a_i}{a_{total}} (C_P(\gamma)_i - (C_P^{target})_i)^2 \quad (18)$$

where a_i represents the area of panel i and $a_{total} = \sum_j a_j$ and γ represents the design variables. The objective function gradients were then calculated with forward-finite-difference with a step-size of 10^{-8} .

Due to the computational requirements associated with calculating finite-difference gradients the limit surface is only defined to the fifth level for this case (1632 surface points) and the deformations are only applied to a maximum

of the fourth level. Additionally, it was found that chordwise and spanwise deformations and sectional translations only had a limited effect on this objective function. The surface was therefore restricted to deformations in just the z (vertical) direction and the trailing edge was held fixed. This was done by only selecting the vertical, non trailing edge design variables. The planforms of the initial and target wings were therefore defined as equal for all cases tested to ensure an achievable objective.

Two cases were investigated: a NACA0012 to NACA4410 case (both un-swept with unit chord and span of three) and a NACA0012 to ONERA M6 case (both with ONERA M6 planform). The results for these two cases are shown in figure 13. They show similar results to the geometry matching cases with the single-level methods improving monotonically for increases in design variables though at some reduction in rate of convergence. The multi-level methods again show improved convergence rates and for the ONERA M6 case equivalent final results. For the NACA4410 case however it can be seen that the multi-level results achieve a slightly improved final result over the single-level 4 case. This indicates that the single-level 4 method was unable to successfully exploit the full design space available.

This is in agreement with results found by the authors for two-dimensional cases^{4,5} where the single-level methods often struggled to fully exploit the available design space with large numbers of design variables. It was however observed that the multi-level cases were much more robust to these problems, always achieving improved or equivalent results to the single-level control cases.

Figure 13. Objective function convergence for inverse design optimisations.

C. Drag Reduction

A final constrained drag reduction case is also investigated. This minimises the drag on an ONERA M6 at $M = 0.84$ and $\alpha = 3.06$ under inviscid flow conditions subject to lift constraint $C_L \geq C_{L0}$ and volume constraint $V \geq V_0$. This test case has been investigated previously by Chauhan³⁴ and Straathof³⁵ who achieved reductions in drag of 25% and 17% respectively.

In this work the unstructured CFD solver SU2³⁶ has been used to calculate the flow solutions with the design variable gradients calculated using the continuous adjoint method³⁶. Each flow (or adjoint) solution is converged down to a residual reduction of 6 orders and convergence is accelerated using multi-grid.

The initial computational mesh was generated using the method outlined in Allen³⁷ with 300,000 mesh points (113×33 on wing surface). To create the initial subdivision wing the method outlined in section III was followed. The resulting surface mesh and the initial two control polygons are shown in figure 11.

Design variables were set up such that only the z coordinates of the control points can change, this implicitly fixes the planform. Furthermore the trailing edge points were held stationary to stop sectional translations. No explicit twist or pitch design variables are included; pitch and twist changes are however achievable through perturbations of the subdivision control points.

Figure 14 shows the surface Mach number for the initial wing which displays a typical lambda shock formation. This initial solution produces a drag of 128 counts with a lift coefficient of 0.311. The drag coefficient produced here is the sum of the wave drag from the shocks and the induced drag as a result of the lift. The wave drag can

theoretically be reduced to zero, however, due to the lift constraint, the induced drag can only be reduced to a minimum of $C_{Dmin} = C_L^2 / \pi AR$ for aspect ratio AR . For this reason a quantity $C_D^* = C_D - C_{Dmin}$ is defined that represents the total reducible drag.

Figure 14. Initial surface pressure distribution for ONERA M6 drag reduction case.

Figure 15 shows the adjoint sensitivities for the initial solution where the yellow represents areas where a positive z deformation will increase the lift or drag and the dark blue represents areas where an equivalent deformation will decrease the lift or drag. Therefore any area where they differ significantly signifies a region with the large potential to reduce the drag while increasing the lift. The main regions where this is true are along the shocks themselves.

Figure 15. Initial adjoint sensitivities for ONERA M6 drag reduction case in z direction

The results of the drag minimisation optimisation are then shown in figure 16 and table 1. They show that the multi-level cases produce the best results reducing the drag by 2-3 drag counts over the best single-level result. The best result (multi-level 1 $\rightarrow$ 4) achieves a total drag count of 85 which represents a 36% reduction in the overall drag. Of the single-level results, case 2 produces the best final drag count with the third and fourth level optimisations stopping at poorer results despite being mathematically able to represent the single-level 2 geometry. This again highlights the difficulties found by the authors in previous work^{4,5} with high fidelity single level optimisations.

Due to the difference in final results the benefits to rate of convergence from using the multi-level methods are not as clear as for the other other test case. It does however seem that the high-fidelity single-level cases do experience some reduction in convergence rate which is not evident in the multi-level convergence histories.

In SNOPT non-linear constraints, such as those used in this problem, are handled by modifying the initial constrained problem such that it can be expressed as the unconstrained optimisation of an augmented Lagrangian merit function. A consequence of this is that the constraints are not guaranteed to be strictly satisfied and in some instances the optimiser may make a trade-off between some small infeasibility and the improve of the objective function. An example of this can be seen in figure 16b which shows the convergence for the lift constraint. It shows that all of the cases run show some small violation of the lift constraint however none of them exceed 0.04%. This is however within a 0.1% tolerance deemed acceptable by the authors.

The final Mach number distributions are shown in figure 17 and it can be seen in all cases that the strength of the shocks relative to the initial solution (figure 14) has been reduced significantly. The single-level 3 and 4 case do

	C_D	C_D^*	C_L	$C_M^\dagger$	V
<i>Baseline</i>	128.00	54.80	0.3111	-0.1028	0.0554
Single-Level 1 (15 DV)	90.61	17.43	0.3111	-0.1005	0.0554
Single-Level 2 (60 DV)	87.92	14.81	0.3109	-0.1013	0.0555
Single-Level 3 (236 DV)	90.44	17.60	0.3103	-0.1014	0.0554
Single-Level 4 (936 DV)	94.88	22.06	0.3103	-0.1010	0.0554
Multi-Level 1→4	85.07	12.47	0.3098	-0.0961	0.0555
Multi-Level 2→4	86.11	13.19	0.3105	-0.1021	0.0554

[†]Pitching moment not explicitly constrained

Table 1. Final results for the ONERA M6 drag reduction optimisation.

Figure 16. Convergence of reducible drag and lift coefficients for the ONERA M6 drag reduction case.

however show some remnant of the front and rear shocks that are not present in the multi-level solutions. The multi-level 2 → 4 case does also appear to show a small shock around the tip that is not present in the multi-level 1 → 4 solution (figure 18), this may be the source of the difference in drag.

The surface changes between the initial and final solutions are presented in figure 19. Interestingly this shows that the single-level 1 and multi-level 1 → 4 cases have large leading edge deformations not present in the other solutions. This represents an increase in the angle of attack that the other methods failed to achieve and may be the reason for the improved result achieved by the 1 → 4 case. It is not surprising that achieving this kind of global deformation is important to the optimisation as the angle of attack is very strongly linked to both the lift and the drag and will have a trimming effect on the result. The other cases show increasingly local surface changes centred around the high sensitivity regions identified earlier. This is most evident in the single-level 4 surface changes where a very thin surface deformation has been applied to the rear shock region. This suggests that this kind of deformation, despite being an accurate representation of the adjoint sensitivity information, is not effective for achieving global minimisation. Previous work from the authors suggests that the creation of high frequency surface undulations may have a negative impact on the overall performance of aerodynamic optimisation procedures³.

VII. Conclusion

In this work a multi-level subdivision surface parameterisation is outlined and tested on a range of optimisation test cases. This method defines a multi-resolutional surface parameterisation using the Catmull-Clark subdivision scheme. The subdivision parameterisation ‘levels’ can be used to control the surface at varying degrees of fidelity. The ‘multi-level’ parameterisation methods describe the use of these sequentially in an optimisation. They therefore start with low fidelity control over the surface which then increases over the course of the optimisation. For each optimisation results for a set of single-level parameterisations are also presented which use each subdivision level in isolation.

Figure 17. Final optimised Mach distributions for the ONERA M6 drag reduction case.

Figure 18. Mach number distribution on optimised ONERA M6 wings at 88% span

Figure 19. Final optimised surface changes for the ONERA M6 drag reduction case.

This represents a control group for the tests.

The first optimisations investigated describe a geometry matching test case that minimises the distance between a target wing and the parameterised subdivision wing. This found that the multi-level parameterisation significantly improved convergence rates over the highest fidelity single-level parameterisation and achieved equivalent results.

Further tests were performed on two inverse design optimisation cases. These minimised the difference in surface pressure between the parameterised subdivision wing and a target. The multi-level methods again showed improvements to the rate of convergence over the high-fidelity single level cases. For the NACA4410 case the multi-level method also achieved an improvement in the final result. This shows that the single-level 4 case (the highest fidelity case used) cannot have reached a true optimum. This highlights the improvements to robustness achieved with the multi-level cases.

A final constrained drag reduction optimisation case was also presented. This minimised the drag for an ONERA M6 at transonic flow conditions subject to the constraints that the lift and volume must not reduce. The best result was achieved by the multi-level 1 $\rightarrow$ 4 case, achieving a final drag result marginally better than the 2 $\rightarrow$ 4 case. It appears that this may be due to a difference in global pitch and suggests that the addition coarse level used in 1 $\rightarrow$ 4 may have made this sort of global deformation more achievable.

Of the single-level methods it was found that the level 2 case produced the best final drag result. This implies that these higher fidelity level 3 and 4 cases failed to achieve their optimum as the level 2 optimum geometry is within their available design space. This emphasises the difficulties associated with the high fidelity single-level cases and consequently highlights the benefits to robustness achieved with the multi-level parameterisation.

In conclusion this work has found that using the multi-level subdivision parameterisation has shown significant improvements in rate of convergence and robustness. The benefits to robustness are particularly important as this lead to exploitation of high-dimensional design spaces not always achievable by the single-level control group. Ultimately this lead to improved final results in two of the test cases investigated.

VIII. Acknowledgements

This work was carried out using the computational facilities of the Advanced Computing Research Centre, University of Bristol - <http://www.bris.ac.uk/acrc/>. The authors also wish to acknowledge the financial support provided by Innovate UK: the work reported herein has been undertaken in GHandI (TSB 101372), a UK Centre for Aerodynamics project.

References

- [1] Masters, D. A., Taylor, N. J., Rendall, T. C. S., Allen, C. B., and Poole, D. J., “Review of Aerofoil Parameterisation Methods for Aerodynamic Shape Optimisation,” *53rd AIAA Aerospace Sciences Meeting*, Jan 2015, doi: 10.2514/6.2015-0761.
- [2] Masters, D. A., Taylor, N. J., Rendall, T. C. S., Allen, C. B., and Poole, D. J., “A Geometric Comparison of Aerofoil Shape Parameterisation Methods,” *54th AIAA Aerospace Sciences Meeting*, Jan 2016.
- [3] Masters, D. A., Taylor, N. J., Rendall, T. C. S., and Allen, C. B., “Impact of Shape Parameterisation on Aerodynamic Optimisation of Benchmark Problems,” *54th AIAA Aerospace Sciences Meeting*, Jan 2016.
- [4] Masters, D. A., Taylor, N. J., Rendall, T. C. S., and Allen, C. B., “Progressive Subdivision Curves for Aerodynamic Shape Optimisation,” *54th AIAA Aerospace Sciences Meeting*, Jan 2016.
- [5] Masters, D. A., Taylor, N. J., Rendall, T. C. S., and Allen, C. B., “A Locally Adaptive Subdivision Parameterisation Scheme for Aerodynamic Shape Optimisation,” *34th AIAA Applied Aerodynamics Conference*, 2016.
- [6] Beux, F. and Dervieux, A., “A hierarchical approach for shape optimization,” *Engineering Computations*, Vol. 11, No. 1, 1994, pp. 25–48.
- [7] Desideri, J. and Zolesio, J., “Inverse shape optimization problems and application to airfoils,” *Control and Cybernetics*, Vol. 34, No. 1, 2005, pp. 165.
- [8] Desideri, J. and Dervieux, A., “Hierarchical methods for shape optimization in aerodynamics I: Multilevel algorithms for parametric shape optimization,” *LECTURE SERIES-VON KARMAN INSTITUTE FOR FLUID DYNAMICS*, Vol. 3, 2006, pp. 10.
- [9] Désidéri, J.-A., El Majd, B. A., and Janka, A., “Nested and self-adaptive Bézier parameterizations for shape optimization,” *Journal of Computational Physics*, Vol. 224, No. 1, 2007, pp. 117–131.
- [10] Martinelli, M. and Beux, F., “Multi-level gradient-based methods and parametrisation in aerodynamic shape design,” *European Journal of Computational Mechanics/Revue Européenne de Mécanique Numérique*, Vol. 17, No. 1-2, 2008, pp. 169–197.

- [11] Andreoli, M., Ales, J., and Désidéri, J.-A., “Free-form-deformation parameterization for multilevel 3D shape optimization in aerodynamics,” 2003.
- [12] Duvigneau, R., Chaigne, B., and Désidéri, J.-A., “Multi-level parameterization for shape optimization in aerodynamics and electromagnetics using a particle swarm optimization algorithm,” 2006.
- [13] Duvigneau, R., “Adaptive parameterization using free-form deformation for aerodynamic shape optimization,” 2006.
- [14] El Majd, B. A., Duvigneau, R., and Désidéri, J., “Aerodynamic shape optimization using a full and adaptive multilevel algorithm,” *ERCRAFT Conference Design Optimization: Methods and Applications, Canary Island, Spain*, 2006.
- [15] Anderson, G. R. and Aftosmis, M. J., “Adaptive Shape Control for Aerodynamic Design,” *56th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Jan 2015.
- [16] Han, X. and Zingg, D. W., “An adaptive geometry parametrization for aerodynamic shape optimization,” *Optimization and Engineering*, Vol. 15, No. 1, 2014, pp. 69–91.
- [17] Sherar, P., Thompson, C., Xu, B., and Zhong, B., “An Optimization Method Based On B-spline Shape Functions & the Knot Insertion Algorithm,” *World congress on engineering*, Citeseer, 2007, pp. 862–866.
- [18] Chaikin, G., “An algorithm for high-speed curve generation,” *Graphical Models /graphical Models and Image Processing /computer Vision, Graphics, and Image Processing*, Vol. 3, 1974, pp. 346–349.
- [19] Cashman, T. J., Hormann, K., and Reif, U., “Generalized Lane–Riesenfeld algorithms,” *Computer Aided Geometric Design*, Vol. 30, No. 4, 2013, pp. 398–409.
- [20] Doo, D. and Sabin, M., “Behaviour of recursive division surfaces near extraordinary points,” *Computer-Aided Design*, Vol. 10, No. 6, 1978, pp. 356–360.
- [21] Catmull, E. and Clark, J., “Recursively generated B-spline surfaces on arbitrary topological meshes,” *Computer-aided design*, Vol. 10, No. 6, 1978, pp. 350–355.
- [22] Forsey, D. R. and Bartels, R. H., “Hierarchical B-spline refinement,” *ACM SIGGRAPH Computer Graphics*, Vol. 22, ACM, 1988, pp. 205–212.
- [23] Finkelstein, A. and Salesin, D. H., “Multiresolution curves,” *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, ACM, 1994, pp. 261–268.
- [24] DeRose, T., Kass, M., and Truong, T., “Subdivision surfaces in character animation,” *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, ACM, 1998, pp. 85–94.
- [25] Ma, W., “Subdivision surfaces for CAD,” *Computer-Aided Design and Applications*, Vol. 1, No. 1-4, 2004, pp. 223–232.
- [26] Cashman, T. J., *NURBS-compatible subdivision surfaces*, Ph.D. thesis, Cashman, Thomas J., 2010.
- [27] Quinlan, G., “Using Freestyle in PTC Creo Parametric to Create Award-Winning Designs,” <http://creo.ptc.com/2013/12/06/using-ptc-creo-parametric-freestyle-to-create-award-winning-designs/>, December 2013, Accessed: 01-Jun-2015.
- [28] Loop, C., *Smooth Subdivision Surfaces Based on Triangles*, Master’s thesis, University of Utah, 1987.
- [29] Dyn, N., Levin, D., and Gregory, J. A., “A Butterfly Subdivision Scheme for Surface Interpolation with Tension Control,” 1990.
- [30] Sabin, M., “Eigenanalysis and artifacts of subdivision curves and surfaces,” *Tutorials on multiresolution in geometric modelling*, Springer, 2002, pp. 69–92.
- [31] Stam, J., “Exact evaluation of Catmull-Clark subdivision surfaces at arbitrary parameter values,” *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, ACM, 1998, pp. 395–404.
- [32] Gill, P. E., Murray, W., and Saunders, M. A., “SNOPT: An SQP algorithm for large-scale constrained optimization,” *SIAM journal on optimization*, Vol. 12, No. 4, 2002, pp. 979–1006.
- [33] Schmitt, V. and Charpin, F., “Pressure Distributions on the ONERA-M6-Wing at Transonic Mach Numbers,” *Experimental Data Base for Computer Program Assessment, Report of the Fluid Dynamics Panel Working Group 04*, AGARD AR 138, 1979.

- [34] Chauhan, D., Chandrashekarappa, P., and Duvigneau, R., “Wing shape optimization using FFD and twist parameterization,” *12th Aerospace Society of India CFD Symposium*, Bangalore, India, Aug. 2010.
- [35] Straathof, M. and van Tooren, M., “Adjoint Optimization of a Wing Using the CSRT Method,” *29th AIAA Applied Aerodynamics Conference*, Jun 2011.
- [36] Economou, T. D., Palacios, F., and Alonso, J. J., “A viscous continuous adjoint approach for the design of rotating engineering applications,” *AIAA Paper*, Vol. 2580, 2013, pp. 24–27.
- [37] Allen, C., “Towards automatic structured multiblock mesh generation using improved transfinite interpolation,” *International Journal for Numerical Methods in Engineering*, Vol. 74 (5), 2008, pp. 697 – 733, Publisher: John Wiley & Sons.