Masters, D., Taylor, N. J., Rendall, T., & Allen, C. B. (2017). A Multi-Level Subdivision Parameterisation Scheme for Aerodynamic Shape Optimisation. *AIAA Journal*. <https://doi.org/10.2514/1.J055785>

Peer reviewed version

License (if available):
Unspecified

Link to published version (if available):
[10.2514/1.J055785](https://doi.org/10.2514/1.J055785)

[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 <https://arc.aiaa.org/doi/10.2514/1.J055785>. 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/>

A Multi-Level Subdivision Parameterisation Scheme 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

Subdivision curves are defined as the limit of a recursive application of a subdivision rule to an initial set of control points. This intrinsically provides a hierarchical set of control polygons that can be used to provide surface control at varying levels of fidelity. This work presents a shape parameterisation method based on this principle and investigates its application to aerodynamic optimisation. The subdivision curves are used to construct a multi-level aerofoil parameterisation that allows an optimisation to be initialised with a small number of design variables, and then be periodically increased in resolution throughout. This brings the benefits of a low fidelity optimisation (high convergence rate, increased robustness, low cost finite-difference gradients) while still allowing the final results to be from a high-dimensional design space. In this work the multi-level subdivision parameterisation is tested on a variety of optimisation problems and compared to a control group of single-level subdivision schemes. For all the optimisation cases the multi-level schemes provided robust and reliable results in contrast to the single-level methods that often experienced difficulties with large numbers of design variables. As a result of this the multi-level methods exploited the high-dimensional design spaces better and consequently produced better overall results.

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^{1,2}. 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 procedure³. Many different methods have been used within an aerodynamic optimisation framework, from standard geometric curve definitions such as B-splines⁴ or NURBS⁵ to aerospace-specific methods such as CST^{6,7}, Hicks-Henne bump functions^{8,9} or PARSEC^{9,10} to Free-Form Deformation^{11–14}, proper orthogonal decomposition^{2,15,16} or the discrete method¹⁷. All of these approaches are subject to the ‘curse of dimensionality’; 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 equates to a single, often expensive, aerodynamic solution the impact of dimensionality can be large. 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 was first used in an

^{*}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

aerodynamic optimisation setting by Beux and Dervieux¹⁸ and has since been applied to a range of aerofoil optimisation problems using a variety of different parameterisation frameworks such as Bèzier curves^{19–22}, Bèzier surface FFD^{23–26}, RBFs²⁷ and B-splines with a knot insertion algorithm^{28,29}. In general, it was shown that implementation of multi-level nested parameterisations can improve the convergence rate, robustness and final solution of an optimisation procedure. This paper investigates the application of multi-level shape parameterisation techniques to multi-resolution subdivision curves for aerodynamic optimisation procedures.

Subdivision curves and surfaces are a shape parameterisation method used predominantly in computer graphics and animation. They describe a smooth curve or 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 curve or surface. In some cases these limit curves are equivalent to B-splines, for example Chaikin’s corner cutting scheme³⁰ is equivalent to uniform quadratic B-splines and further extensions to higher order uniform B-splines can also be derived³¹. For subdivision surfaces there are similar equivalences; for a set of regular, rectangular control points Doo-Sabin subdivision³² is equivalent to bi-quadratic B-spline surfaces and Catmull-Clark subdivision³³ is equivalent to uniform bi-cubic B-spline surfaces. However, subdivision surfaces can be generalised to arbitrary topologies while B-spline surfaces can not; this is a key benefit of the subdivision framework. B-splines and subdivisions share many characteristics; the method of implementation is, however, one area of major difference. B-splines utilise continuous parametric representation whereas subdivisions use a hierarchical process of discrete refinement. It is the 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 either coarse geometry changes while maintaining the fine detail, or fine geometry changes while maintaining the overall shape. This approach has been implemented comparably from both B-spline³⁴ and subdivision³⁵ perspectives in both their two-dimensional and three-dimensional forms, though it would seem 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 animation³⁶. It is also slowly being incorporated into some computer aided design (CAD) packages^{37,38}. A comprehensive overview of subdivisions and multi-resolucional analysis can be found in Stollnitz *et al.*³⁹.

The aim of this work is to explore the use of multi-resolution subdivision curves for aerodynamic shape optimisation with particular emphasis on how they can be used to improve both the efficiency, accuracy and robustness of current optimisation procedures.

II. Subdivision Curves

Subdivision curves 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 linear combination of the old points. For this reason the refinements can conveniently be expressed as a simple matrix transformation

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

from old points C_n to new points C_{n+1} . For simple subdivision schemes on closed polygons these matrices are just the two row offset repetition of a subdivision ‘mask’. Two common examples are Chaikin’s rule³⁰ (figures 1a and 1b) and the Cubic B-spline rule (figures 1c and 1d). Figure 2 shows a simple implementation of Chaikin’s rule on a closed polygon. Note how the rows of the transformation matrices all sum to 1 and therefore describe a weighted averaging of the previous points, this is a key feature of all subdivision transformation matrices. Figures 1a and 1c show the matrices for areas of smooth subdivision whereas figures 1b and 1d represent areas with a corner or endpoint, where the non-averaged matrix row represents the corner or endpoint itself.

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

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

for some $n < N$. The limit curve can therefore be described by the relation

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

In practice this calculation must be truncated at some point and the limit curve calculated. A typical solution to this is to calculate an evaluation matrix based on eigenanalysis⁴⁰ that pushes the subdivision points to their limit locations. With this method however the user cannot directly parameterise the curve itself and therefore cannot specify the distribution

$$P_n = \begin{bmatrix} \ddots & \ddots & & & \\ 0.25 & 0.75 & & & \\ & 0.75 & 0.25 & & \\ & 0.25 & 0.75 & & \\ & & 0.75 & 0.25 & \\ & & 0.25 & 0.75 & \\ & & & \ddots & \ddots \end{bmatrix}$$

a) Chaikin's rule for smooth areas

$$P_n = \begin{bmatrix} \ddots & \ddots & & & \\ 0.5 & 0.5 & & & \\ & 1 & & & \\ & 0.5 & 0.5 & & \\ & 0.25 & 0.75 & & \\ & & & \ddots & \ddots \end{bmatrix}$$

b) Chaikin's rule for corners or endpoints

$$P_n = \begin{bmatrix} \ddots & \ddots & & & & \\ 0.125 & 0.75 & 0.125 & & & \\ & 0.5 & 0.5 & & & \\ & 0.125 & 0.75 & 0.125 & & \\ & & 0.5 & 0.5 & & \\ & & 0.125 & 0.75 & 0.125 & \\ & & & \ddots & \ddots & \ddots \end{bmatrix}$$

c) Cubic B-spline rule for smooth areas

$$P_n = \begin{bmatrix} \ddots & \ddots & & & & \\ 0.5 & 0.5 & & & & \\ & 1 & & & & \\ & 0.5 & 0.5 & & & \\ & & 0.75 & 0.25 & & \\ & & 0.1875 & 0.6875 & 0.125 & \\ & & & \ddots & \ddots & \ddots \end{bmatrix}$$

d) Cubic B-spline rule for corners or endpoints

Figure 1. Matrix representations of the Chaikin and Cubic B-spline subdivision schemes.

Figure 2. Simple closed subdivision using Chaikin's Scheme

of points on it. As the distribution of points around an aerofoil can be very important an alternative method is used in this work. This is done by exploiting the fact B-spline curves and some subdivision formulations create equivalent limit curves for identical control polygons. This means that the discrete operation between the final subdivision level N and the limit curve can be expressed as a continuous B-spline transformation. This B-spline transformation can then be discretised and formulated as matrix P_N^{BS} such that a desired point distribution is achieved. This method is however constrained to subdivisions with B-spline equivalents.

It is then convenient to define

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

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

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

The columns of ϕ_i represent the basis functions of the subdivision scheme. Figure 3 shows the basis functions for a Chaikin subdivision with fixed endpoints, these are identical to the basis functions of an equivalent quadratic B-spline. From this it can be seen that the different subdivision levels can be used to control and deform the limit surface at varying levels of fidelity.

III. Reverse Subdivision Curves

Given a set of points it can often be desirable to obtain the subdivision polygon (or closest possible match) that produced them. Consider a fine set of points C_n ; the coarser set of points C_{n-1} can simply be calculated as the least squares solution of equation 1,

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

Figure 3. The basis functions for the first two levels of a Chaikin subdivision with fixed endpoint conditions with a 4 point initial control polygon

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 of the column space $Q_n = \text{null}(P_n^T)$. It should be noted that Q_n is not uniquely defined and while any choice makes theoretical sense, numerically it is important that it is well conditioned. For this reason an orthonormal basis of the null space is used. 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 . Then by letting $\begin{bmatrix} P_n & Q_n \end{bmatrix}^{-1} = \begin{bmatrix} A_n \\ B_n \end{bmatrix}$, equation 6 can be re-expressed as the reverse subdivision equations

$$C_{n-1} = A_n C_n, \quad (9)$$

$$D_{n-1} = B_n C_n. \quad (10)$$

This importantly creates a one-to-one relationship between the subdivision refinement levels and thus allows information to be propagated uniquely and exactly in either the refinement or coarsening direction (figures 4 and 5).

Figure 4. Process for subdivision refinement.

Figure 5. Process for reverse subdivision.

Equation 5 can then be reformed using equation 8 such that

$$C_\infty = \phi_n C_n + \sum_{i=n}^N \phi_{i+1} Q_i D_i. \quad (11)$$

By storing and including these error terms it means that any discretised shape can now be represented by any subdivision rule as long as the correct error terms D_n are used.

IV. Aerofoil Parameterisation

To parameterise the aerofoils in this work a cubic B-spline subdivision scheme is used with a single, closed initial polygon with ‘corners’ at the leading and trailing edges. This ensures that the position of the leading and trailing edges are equal to the ‘corner’ control points at every subdivision level. This is equivalent to using two distinct subdivision curves for the upper and lower surfaces with shared endpoints at the leading and trailing edges. In physical terms the trailing edge is a corner, however the leading edge is not, for this reason the control points closest to the leading edge are constrained to lie directly above and below it; this enforces the vertical surface required.

In this work the first subdivision level has been defined by six control points with two points defining the leading and trailing edges at $[x/c, z/c] = [0, 0]$ and $[1, 0]$ respectively, and points at $x/c = 0$ and $x/c = 0.5$ on each surface. The maximum control level is defined by the 8th subdivision level with 260 control points.

It was found that if the number of discrete points defining the aerofoil is reduced to the vicinity of the number of design variables, the influence of the point distributions can significantly skew the optimisation results through an aliasing effect. For this reason, unless otherwise specified, the aerofoils have been defined by 601 points, significantly more than the maximum number of design variables. The points were then distributed along the chord in a half-cosine fashion.

Table 1 shows the total number of control points in each level used. All of the control points are then allowed to move in just the z direction apart from the trailing edge which is held fixed.

Subdivision Level	Number of Control Points
1	6
2	8
3	12
4	20
5	36
6	68
7	132
8	260

Table 1. Number of control points at each subdivision level

Given an initial aerofoil $C^{initial}$, the starting initial subdivision positions can then be calculated by the recursive application of equation 9. The resulting control point positions, $C_n^{initial}$, represent the least-squares approximations of the initial aerofoil for the limit surfaces $\phi_n C_n^{initial}$. At each of these subdivision levels a set of fixed error terms D_n can then be calculated by applying equation 10. Then for the set of control points $C_n^{initial}$ calculated, equation 11 implies that

$$C^{initial} \equiv \phi_n C_n^{initial} + \sum_{i=n}^N \phi_{i+1} Q_i D_i, \quad \forall n \leq N. \quad (12)$$

This approach has been used for all the optimisations in this paper, therefore all cases start from the exact prescribed (discrete) initial aerofoil rather than a best approximation. Figure 6 shows the initial positions for the first four subdivision levels.

Each of these available subdivision levels can then be used to parametrise the aerofoil independently, in what have been described hereafter as ‘single-level’ schemes, or in an ascending series, which have been described as ‘multi-level’. For the single-level schemes the aerofoils are defined by

$$C^{aero} = \phi_n C_n + \sum_{i=n}^N \phi_{i+1} Q_i D_i \quad (13)$$

for the desired subdivision level n with control points C_n . For the multi-level schemes the same definition applies but when a ‘refinement trigger’ is activated $n \rightarrow n + 1$. This increases the number of design variables, and therefore the available fidelity, while maintaining the aerofoil shape exactly. The aim of this is to use fewer design variables in the early stages of optimisation to increase the rate of convergence, robustness and, for finite-difference gradients, reduce gradient calculation times. Then, as the lower dimensional design space looks to be fully exploited, the refinement process is applied to allow the larger design space to be explored.

Figure 6. Initial control point positions for the first four subdivision levels for a NACA0012.

For each optimisation case investigated a variety of subdivision schemes have been applied. Each available subdivision level has been applied as a single-level parameterisation as well as four multi-level schemes starting from the first, second, third and fourth levels respectively and refining up to the final level. The single-level schemes are equivalent to using normal cubic B-splines and act as the control group for which the benefits of the multi-level tests have been compared against.

V. Optimisation Methodology

In this work subdivision curves have been applied to a range of aerodynamic optimisation problems. For all of these tests the multi-purpose large-scale optimiser SNOPT⁴¹ was 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. A feasibility tolerance of 10^{-6} was also used and optimiser convergence was determined based on the activation of one of three criteria:

1. The Karush-Kuhn-Tucker (KKT) first-order optimality condition⁴¹ satisfying a tolerance level of 10^{-6}
2. The optimiser unable to improve the objective function
3. For the multi-level schemes *not* on the final level, satisfying the refinement condition (equation 14)

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 local 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 Anderson²⁷ 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 is implemented 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.

$$\left| \frac{1}{w} \sum_{j=0}^{w-1} G_{k-j} \right| < \left| \max_{m \leq l \leq k} \frac{t}{m} \sum_{j=0}^{m-1} G_{l-j} \right|, \quad \forall G \in \{G^{obj}, G^{con}\} \quad (14)$$

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

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 be 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.

VI. Optimisation Results

To test benefits of using the multi-level subdivision parameterisation described above, a series of optimisations of varying difficulty have been performed. A set of geometry matching problems present the simplest test, minimising the RMS geometry error between the aerofoil surface and a target aerofoil. A further set of inverse design problems present a more challenging optimisation with a final set of three drag optimisation cases representing the most difficult tests.

A. Geometry Matching

Three geometry matching problems have been considered in this work. Each one starts from an initial NACA0012 then targets RAE2822, NACA4410 and ONERA M6 aerofoils. For all of these optimisations the objective function is defined as the root-mean-squared difference between the vertical components of the current and target aerofoils, i.e.,

$$J = \sqrt{\frac{1}{n} \sum_{i=1}^n (z_i - z_i^{target})^2}. \quad (15)$$

The gradients are then calculated analytically and the refinement parameters used are $w = m = 1$ and $t = 0.3$.

Figures 7a, 7b, 7c show the results of these optimisations. For all of the cases tested, final convergence was triggered through reduction of the KKT optimality condition below the prescribed tolerance and all of the cases achieve consistent minima. For the single-level parameterisations this means that as more design variables are used the results improve and for the multi-level methods this means that optimums equivalent to the single-level 8 cases are always achieved; these are all of order 10^{-7} . It can also be seen that for the single-level cases the rate of convergence typically reduces when large numbers of design variables are used. For the multi-level cases however, convergence seems to be fairly consistent regardless of the numbers of design variables used. As a result they generally show better overall convergence rates than the equivalent single-level cases; in particular the multi-level $4 \rightarrow 8$ method shows the best convergence rates for all three test cases.

Other work by the authors⁴² has shown that for aerofoil approximations the errors in drag coefficient prediction for transonic Euler flow relative to a target aerofoil are roughly equivalent to the maximum geometric error. These results therefore suggest that if the eight subdivision level is used for an inviscid drag reduction case, the design space should include aerofoils within the order of a thousandth of a drag count of any optimum.

B. Inverse Design

Three inverse design problems have also been considered to compare the subdivision aerofoil parametrisations. Each one starts from an initial NACA0012 then targets RAE2822, NACA4410 aerofoils with $\alpha = 0$ and an ONERA M6 with $\alpha = 3$. For all of these the objective function is defined as the root-mean-squared difference between the current and target pressure distributions, i.e.,

$$J = \sqrt{\frac{1}{n} \sum_{i=1}^n (Cp_i - Cp_i^{target})^2}. \quad (16)$$

The pressure distributions are calculated with a potential flow panel code and the gradients are calculated through finite-forward-differencing with a step size of 10^{-8} . The refinement parameters used are $w = m = 1$ and $t = 0.1$.

Figures 8a, 8c, 8e show the inverse design test case results. For all of the cases the final convergence was triggered due to an inability to improve the solution, the optimality had typically reduced to around 10^{-4} at this point. For the single-level methods it can again be seen that, in general, the final results improve with increased fidelity. There are however some instances where this is not true, for example, for the NACA4410 case (figure 8c) the 6th, 7th and 8th level methods fail to improve on the final result obtained using just the 5th level. This shows that the optimiser has

a) NACA0012 → RAE2822

b) NACA0012 → NACA4410

c) NACA0012 → ONERAM6

Figure 7. Results for the geometry matching optimisations. Numbers represent the number of design variables at the end of each optimisation phase.

failed to reach an improved solution. It can also be seen that the rate of convergence tends to increase with fidelity for these cases, in particular it can be seen for the RAE2822 case (figure 8a) where the ‘single-level 8’ case takes a significantly longer to converge than any of the other methods. The multi-level methods converged to the optimum result consistently for all cases investigated and have better convergence rates than the high fidelity single-level cases. This represents an improvement to both the robustness and efficiency of the optimisation by using the multi-level methods.

C. Inviscid Drag Reduction

A series of three inviscid drag optimisation cases have also been tested. A symmetric NACA0012 drag reduction case, a lifting NACA0012 case with a constraint on lift and a lifting RAE2822 case with lift and moment constraints. A full description of the test specifications is presented in table 2.

The unstructured CFD solver SU²⁴³ has been used for all of these optimisations with the design variable gradients calculated using the continuous adjoint method⁴³ and convergence acceleration using three multi-grid levels.

The computational meshes were generated using a structured conformal mapping approach where all surface cells have aspect ratio one and the farfield distance is 50 chord lengths. The mesh was then deformed at every iteration using an RBF method⁴⁴ with a support radius of 10 chord lengths and Wendland’s C4 RBF⁴⁵. This ensures the surface deformations are dissipated smoothly across the mesh to maintain mesh quality.

For each test case a mesh resolution study has been performed on four different resolution meshes: 257×129 , 513×257 , 1025×513 , 2049×513 (*surface points $\times$ points to far field*). For each of these calculations the flow was converged to approximately machine zero and it should be noted that for the symmetric NACA0012 case half meshes were used so the number of surface points was reduced by a factor of two. For all three cases the 513×257 mesh was chosen for optimisation to provide a compromise between accuracy and computation time. Figure 10 shows the initial mesh used for the lifting NACA0012 optimisation; an identical halved mesh was used for the symmetric case and an equivalent topology mesh was used for the RAE2822 case.

A gradient accuracy test was also performed for each case. This calculated the gradients for all of the level 8 control points for the initial aerofoils using both the continuous adjoint method that will be used for the optimisation as well as finite-forward-differencing. Again, all of the flow and adjoint solutions were converged to machine zero. A finite-difference step-size of 10^{-5} was used, this was calculated as the root of the precision, which for SU² is the output precision of 10^{-10} . Figure 9 shows the comparison of these gradients for both the objective function and the constraints. It can be seen that for all three cases the two gradient types match closely though with some discrepancy, particularly in the shocked regions. A further gradient comparison was performed at the increased mesh resolution of 1025×513 , however errors of a similar magnitude were found. The errors seem to be associated with the continuous adjoint method itself, which is in agreement with results found by other researchers^{46–50}.

It was found that, much like for the inverse design cases, final convergence drag for these cases was always triggered due to an inability to improve the result rather than the convergence of KKT optimality condition. This was typically reduced to an order of 10^{-2} .

For the final optimisations the flow and adjoint residuals were reduced by 6 orders. This produced lift and drag force adjoints converged to an order of 10^{-7} . A preliminary study found that, for a single multi-level optimisation, reducing the residuals by 12 orders compared to 6 only improved the final solution by 7 hundredths of a drag count while increasing the runtime by a factor of 2.2. A 6 order reduction was therefore considered satisfactory.

	Symmetric NACA0012	Lifting NACA0012	RAE2822
Initial Aerofoil:	NACA0012	NACA0012	RAE2822
Flow Conditions:	$\alpha = 0$ $M = 0.85$	$\alpha = 1.25$ $M = 0.82$	$\alpha = 2.79$ $M = 0.734$
Minimise:	C_D	C_D	C_D
Subject to:	$C_L = 0$ $z \geq z_{NACA0012}$	$C_L \geq C_{Lorig}$ $A \geq A_{orig}$	$C_L \geq C_{Lorig}$ $C_M \geq C_{Morig}$ $A \geq A_{orig}$
Refinement Parameters:	$t = 0.1, w = 3, m = 3$	$t = 0.1, w = 3, m = 3$	$t = 0.1, w = 3, m = 3$

Table 2. Test specifications for inviscid drag reduction cases.

a) NACA0012 → RAE2822

b) NACA0012 → RAE2822 (Close-up)

c) NACA0012 → NACA4410

d) NACA0012 → NACA4410 (Close-up)

e) NACA0012 → ONERA M6

f) NACA0012 → ONERA M6 (Close-up)

Figure 8. Results for the inverse design optimisations. Numbers represent the number of design variables at the end of each optimisation phase.

a) Symmetric NACA0012		b) Lifting NACA0012			
Mesh Res.	C_D	Mesh Res.	C_D	C_L	C_M
129×129	469.41	257×129	359.84	0.394	0.0781
257×257	469.35	513×257	360.49	0.396	0.0787
513×513	469.41	1025×513	361.13	0.397	0.0792
1025×513	469.41	2049×513	360.03	0.396	0.787

c) RAE2822			
Mesh Res.	C_D	C_L	C_M
257×129	233.38	1.016	1.464
513×257	234.08	1.021	1.477
1025×513	234.13	1.022	1.479
2049×513	234.02	1.022	1.479

Table 3. Mesh convergence results for all three inviscid drag reduction cases.

1. Symmetric NACA0012

This test case was originally outlined by the Aerodynamic Design Optimisation Discussion Group⁵¹ and describes the inviscid drag reduction of a NACA0012 at $M = 0.85$ and $\alpha = 0$ under the constraint that the local thickness can only increase. A considerable number of papers have been published on this case^{1,2,52–60} and it has been shown to be a particularly difficult problem due to a range of characteristics such as multiple local minima⁵⁸, non-symmetric solutions⁵⁹ and hysteresis⁶⁰. In this previous work, results ranging from 100 to 25 drag counts have been reported with the best cases produced by Lee and Zingg *et al.*⁵⁹, 42 counts, Bisson and Nadarajah⁵⁵, 25 counts, and Masters *et al.*⁵⁸, 25 counts.

Due to the symmetry of the problem a symmetry plane was aligned with the aerofoil chord and just half the mesh was solved. This meant that only half of the aerofoil surface, and therefore only half of the subdivision curve, needed to be modelled. For this reason each subdivision level has half the number of control points and design variables used in the other optimisation cases.

The local thickness constraint for this case can be expressed as a linear transformation of the design variables so is enforced as linear constraint in SNOPT (at every chord-wise percentile). This enables SNOPT to efficiently satisfy it at every iteration major and minor iteration. For this optimisation each CFD run was initialised from the previous best solution. This was found to decrease the computation time for solutions close to the best geometry and have a small negative impact on those far away. This resulted in a significant reduction in overall computation time. It was also found to increase the chances of a similar solution being found when multiple solutions exist^{59,60}. This is believed to be the main contributor to the improvement in results compared to previous work from the authors^{58,61}.

For this case no improvement was made at subdivision level 1 (with 2 design variables) therefore only levels two and above are included in the results (displayed in figures 11–13). The drag convergence histories in figure 11 show that the single-level methods improve with fidelity up to level 5 after which they fail to successfully exploit the available design space. This can further be seen in figures 12 and 13 where the single-level results show significantly different aerofoil shapes and C_p distributions to the optimum found. In particular the aerofoil shapes for levels 6, 7 and 8 are not very smooth. The formation of these non-smooth shapes is likely to be the reason for their premature stagnation which is in agreement with the conclusions of Masters *et al.*⁵⁸. It can then be seen that the multi-level methods all converge to good results and all surpass those from the single-level methods. In this case the ‘3 → 8’ method produces the best result of 15.7 drag counts. Figures 12 and 13 show that all three of the multi-level methods converge to similar aerofoil shapes though there are some small discrepancies in the C_p distributions. Interestingly, even though the ‘2 → 8’ and ‘4 → 8’ methods converge to very similar drag values they have slightly different C_p distributions at the leading and trailing edges. This indicates that there may be some degree of multi-modality around the optimum solution.

A limited study was also done on the effects of increasing the mesh resolution. For this the ‘3 → 8’ multi-level optimisation was run with increased mesh resolutions of ‘513 × 513’ and ‘1025 × 513’. The refinement parameters were also relaxed to $t = 0.01$, $w = 7$ and $m = 7$. This was done to ensure the best final results were achieved, though at the cost of some efficiency. The convergence histories for these optimisations are displayed in figure 14 and show that the all three mesh resolutions produce similar results up to approximately 16 drag counts where the increased mesh resolution allows a further reduction in drag over the initial result. The final drag results are then summarised in table

Figure 9. Comparison of finite-difference and adjoint gradients calculated for the level 8 control points on 513×257 mesh.

5 which show that the finest mesh produces the best result of 4.2 drag counts. Figures 15a and 15b then show the aerofoil shapes and pressure distributions for these cases. They show that the optimised trailing edge geometries are very similar though show some notable differences in their pressure distributions.

It appears that these increased resolution results are converging to zero as the number of surface points is increased. As the solutions appear to be shock free this would represent a reduction in the numerical drag as would be expected. An attempt was made to run the final shapes with an increased mesh resolution in an attempt further reduce the final drag. Difficulties were however encountered with the hysteresis present at the design point for these solutions. It was found that the ‘lower branch’ drag solution was unattainable from a uniform initialisation and could not even be achieved through a sweep in Mach number (as was used by Meheut *et al.*). It was therefore deemed that the full optimisation history of solutions was crucial in attaining these solutions.

2. Lifting NACA0012

This second case describes the optimisation of a NACA0012 under non-symmetric flow conditions. The optimisation is constrained such that the lift coefficient and total area must not decrease below their initial values. Due to the strong influence of angle of attack on lift and drag it is included as a design variable for this case and the leading edge control point is held fixed. Any deviation from the initial value of 1.25 is then implemented as a rotation of the aerofoil surface mesh rather than an alteration of the the solver settings. This is done to preserve the boundary conditions

Figure 10. Mesh used for Lifting NACA0012 optimisations.

	C_D (Counts)
<i>Baseline</i>	469.4
Single-Level 2 (3 DV)	332.2
Single-Level 3 (5 DV)	148.5
Single-Level 4 (9 DV)	65.9
Single-Level 5 (17 DV)	25.1
Single-Level 6 (33 DV)	246.5
Single-Level 7 (65 DV)	247.0
Single-Level 8 (129 DV)	144.7
Multi-Level 2→8	18.8
Multi-Level 3→8	15.7
Multi-Level 4→8	18.6

Table 4. Table of final drag results (in counts) for the inviscid NACA0012 optimisation at $\alpha = 0$ and $M = 0.85$ with 257×257 mesh.

exactly throughout the optimisation. The lift and area constraints are implemented as non-linear in SNOPT for this case. This enforces them by transforming the constrained optimisation of the objective function into the unconstrained optimisation of an augmented Lagrangian merit function⁴¹. A consequence of this is that objective function is not guaranteed to decrease for every iteration like in the symmetric NACA0012 problem.

The CFD solutions were again initialised from previous solutions, though as the objective function was not guaranteed to monotonically decrease it could not be taken as the solution with the lowest drag. The optimisations are however guaranteed to improve (based on the constraints as well as the objective function) at every major iteration, for this reason the previous major iteration was used for initialisation.

Figures 16-19 and table 6 show the results of this optimisation. From table 6 and the drag convergence history (figure 16) it can be seen that all four multi-level methods and single-levels 3 to 5 achieve good results. They each obtain a final drag value of between 3 and 4 counts and converge to approximately this solutions within 20-40 iterations. The Level 6 case does also produce a comparable result though is slightly infeasible and takes 290 iterations to converge.

Considering the pressure distributions for these low drag solutions (figure 19) it can be seen that all of them produce shock free solutions. It is however clear from these distributions as well as the final aerofoil shapes (figure 18) that there are large differences in the solutions. This shows that there is a range of shapes that eliminate the shock for this test case, which by d'Alemberts paradox should exert zero drag, therefore a definitive optimum point may not be well defined. It should also be noted that the remaining drag calculated in these cases is a result of numerical error and does

	Mesh Resolution	C_D (Counts)
Multi-Level 3→8	257×257	15.7
Multi-Level 3→8	513×513	8.2
Multi-Level 3→8	1025×513	4.2

Table 5. Final drag results (in counts) for the symmetric NACA0012 optimisation with increased mesh resolution.

Figure 11. Drag convergence for the symmetric NACA0012 optimisation. Numbers represent the number of design variables at the end of each optimisation phase.

Figure 12. Final aerofoil shapes for the symmetric NACA0012 optimisation.

Figure 13. Final pressure distributions for the symmetric NACA0012 optimisation.

Figure 14. Drag convergence for the symmetric NACA0012 case with increased mesh resolution. Numbers represent the number of design variables at the end of each optimisation phase.

Figure 15. Trailing edge aerofoil shapes (a) pressure distributions (b) for the symmetric NACA0012 case with increased mesh resolution.

not represent a physical quantity.

The single-level 7 and 8 results display similar trends to previous test cases, showing a very slow rate of convergence and stopping a long way from their best possible result. In this instance the level 8 case does particularly poorly, only achieving a final result of 80 drag counts. Looking at the final aerofoil and pressure distribution produced in figure 19 it can be seen that a sharp step is formed on the surface around the shock and high frequency oscillations are present in the pressure. Looking at this shape it is unsurprising that the optimiser could not find a way to navigate away from it towards the smooth solutions found by the other cases.

It can be seen that there are some large variations in both the final objective function and constraint values. This is surprising as it would be expected that all of the successful cases would converge to similar shapes and for those solutions to lie on the constraint boundary. The most likely conclusion for these difficulties derive from sources of error or noise in the objective and gradient calculations. Firstly, as many of the cases find shock-free solutions the drag is dominated by numerical noise which can not be accounted for in the continuous gradient calculation. This will mean that even with a physically correct gradient it may not be possible to find an improved solution along the search direction. Secondly the continuous gradient itself has some degree of error, identified in figures 9c and 9d, which may be more influential in this area where low gradients are expected.

Nevertheless these results again show good evidence for the robustness of using a multi-level method and the difficulties that can arise by using high-fidelity design variables throughout a full optimisation.

	$\alpha(^{\circ})$	C_D	C_L
<i>Baseline</i>	1.25	360.49	0.396
Single-Level 1 (5 DV)	2.20	205.53	0.349
Single-Level 2 (7 DV)	1.60	33.69	0.396
Single-Level 3 (11 DV)	0.15	3.82	0.396
Single-Level 4 (19 DV)	0.14	3.75	0.400
Single-Level 5 (35 DV)	-0.15	3.27	0.397
Single-Level 6 (68 DV)	-0.47	3.27	0.393
Single-Level 7 (131 DV)	-0.46	11.29	0.395
Single-Level 8 (259 DV)	0.08	80.79	0.400
Multi-Level 1→8	1.14	4.03	0.407
Multi-Level 2→8	1.12	3.82	0.405
Multi-Level 3→8	0.08	3.30	0.399
Multi-Level 4→8	0.23	3.30	0.399

Table 6. Final results for the lifting NACA0012 optimisation. Red cells represent constraint violations.

Figure 16. Drag convergence for the lifting NACA0012 optimisation. Numbers represent the number of design variables at the end of each optimisation phase.

Figure 17. Lift convergence for the lifting NACA0012 optimisation. Numbers represent the number of design variables at the end of each optimisation phase.

Figure 18. Final aerofoils for the lifting NACA0012 optimisation.

Figure 19. Final pressure distributions for the lifting NACA0012 optimisation.

3. RAE2822

The final optimisation presented in this work is based on ADODG test case 2⁵¹, it has however been modified to exclude viscosity. As with the lifting NACA0012 optimisation the angle of attack is included as a design variable and the CFD solutions are again initialised from the previous major iteration. The lift and pitching moment coefficients and total area were also constrained such that they must not decrease. Some difficulties were however experienced with some preliminary optimisations where a cusped leading edge was created. This led to the CFD solutions failing due to an overlapping mesh in this area. To avoid this behaviour a further constraint was added to enforce positive curvature around the leading edge. This was calculated using a three-point Menger scheme⁶² and was applied to all the points in the first 1% of the chord.

The results for this test case are shown in figures 20 to 24, and table 7. The drag results show that all of the multi-level methods, as well as single-level methods for levels 3 to 5, converge to results of between 16 and 18 drag counts. The pressure distributions in figure 24 suggest that all of these solutions are shock free. These cases also show good convergence rates, reducing the drag to below 20 counts in 10 to 40 iterations. Figures 23 and 24 do however display some quite large differences in the aerofoil shapes and pressure distributions. This suggests, in a similar fashion to the lifting NACA0012 case, there are a range of solutions that eliminate the shocks.

The single-level 7 case also achieved a comparable result of 19 drag counts though interestingly figure 24 shows that a significant shock is still present. It appears that for this case the optimiser has found a solution with a portion of back-pressure that cancels some of the wave drag. The single-level 8 case again displays difficulties exploiting its available design space, obtaining a final drag of just 157 counts. Similarly the single-level 6 case stagnates after only 5 iterations with 131 counts.

Much like for the lifting NACA0012 case some difficulties were experienced converging to consistent, feasible minima; this can most likely be attributed to similar difficulties obtaining accurate gradients. These difficulties are however compounded by the additional pitching moment constraint. As a result a significant number of the constraints have been violated by small amounts (table 7). Despite these difficulties, this case again shows the benefits achieved by using the multi-level methods relative to the high-dimensional single-level parametrisations.

	$\alpha(^{\circ})$	C_D	C_L	C_M
<i>Baseline</i>	2.79	234.07	1.021	0.1477
Single-Level 1 (5 DV)	2.88	96.13	1.007	0.1411
Single-Level 2 (7 DV)	2.34	35.96	1.021	0.1478
Single-Level 3 (11 DV)	2.12	16.93	1.020	0.1475
Single-Level 4 (19 DV)	2.23	17.38	1.020	0.1480
Single-Level 5 (35 DV)	2.17	18.02	1.019	0.1424
Single-Level 6 (68 DV)	2.64	130.91	0.961	0.1136
Single-Level 7 (131 DV)	2.06	18.72	1.019	0.1465
Single-Level 8 (259 DV)	2.44	156.89	1.037	0.1471
Multi-Level 1→8	2.15	17.50	1.016	0.1466
Multi-Level 2→8	1.92	16.51	1.020	0.1475
Multi-Level 3→8	2.18	16.82	1.018	0.1472
Multi-Level 4→8	2.23	17.76	1.021	0.1476

Table 7. Final results for the RAE2822 optimisation. Red cells represent constraint violations.

VII. Conclusions

In this work a multi-level subdivision parameterisation scheme has been tested on a series of optimisation problems and have been compared to a range of single-level subdivision parameterisations that are equivalent to cubic B-splines. In total nine test cases have been investigated, three geometric shape matching cases, three inverse design cases and three inviscid drag reduction cases. These provide a range of tests that vary significantly in their computational requirements and complexity.

The geometric shape matching cases represent the simplest of the optimisations. For all three of the tests all of the individual cases reach a consistent local minimum. It was however found that by using the multi-level methods compared to the single-level methods the optima could be reached with 20-60% fewer iterations.

The three inverse design cases then represent a significant increase in complexity compared to the geometry matching optimisations. For these cases it was found that some of the high-fidelity single-level methods stopped prematurely

Figure 20. Drag convergence for the RAE2822 optimisation. Numbers represent the number of design variables at the end of each optimisation phase.

Figure 21. Lift convergence for the RAE2822 optimisation. Numbers represent the number of design variables at the end of each optimisation phase.

Figure 22. Pitching moment convergence for the RAE2822 optimisation. Numbers represent the number of design variables at the end of each optimisation phase.

Figure 23. Final aerofoils for the RAE2822 optimisation.

Figure 24. Final pressure distributions for the RAE2822 optimisation.

and failed to fully exploit the available design space. The multi-level methods, however, did not have this problem and all reached a consistent minimum. In some instances this meant that the multi-level methods found solutions orders of magnitude better than the single-level cases.

The three drag reduction cases then represent the most complex optimisations investigated. For all three of these optimisations the multi-level cases, and the first five single-level cases all appeared to converge to good, consistent results. This however was not the case for the higher level single-level methods, which typically struggled to even achieve results equivalent to their lower dimensional alternatives.

The main difference between the three cases was whether or not the optimum result could be achieved with a low number of design variables. In the case of the lifting NACA0012 and the RAE2822 optimisations a shock free solution was achieved with only 11 design variables (single-level 3), therefore the multi-level methods could only match these results. For the symmetric NACA0012 case however, it seems that a high dimensional design space was required to achieve the optimum, this was not possible using single-level methods due to the difficulties of navigating the high dimensional space. As a result, all of the multi-level methods outperformed the best single-level case with the ‘3→8’ method producing the best result of 15.7 drag counts. Further improvement on this result was then achieved by increasing the mesh resolution by a factor of 8. This yielded a final drag count of just 4.2; this is a significant improvement on the previously published results for this test case.

In conclusion, it was found that significant robustness and convergence difficulties were experienced if a large, fixed number of design variables were used for the full duration of an optimisation process. As a consequence this

meant that potentially beneficial high-dimensional design spaces could not be utilised reliably. This was however overcome by using a multi-level subdivision parameterisation that started the optimisation with a low number of design variables then increased them throughout the process. It was found that by using this approach robust and efficient performance could be achieved across all nine of the optimisations, with the high-dimensional design space available consistently exploited where possible. In some instances this enabled the multi-level methods to achieved results orders of magnitude better than the single-level control group.

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/> and the data necessary to support the conclusions are included in the paper. 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] Carrier, G., Destarac, D., Dumont, A., Meheut, M., Salah El Din, I., Peter, J., Ben Khelil, S., Brezillon, J., and Pestana, M., "Gradient-based aerodynamic optimization with the elsA software," *52nd Aerospace Sciences Meeting*, , No. January, Jan. 2014, pp. 1–31.
- [2] Poole, D. J., Allen, C. B., and Rendall, T. C. S., "Application of control point-based aerodynamic shape optimization to two-dimensional drag minimization," *52nd Aerospace Sciences Meeting*, No. January, 2014, pp. 1–13.
- [3] 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.
- [4] Braibant, V. and Fleury, C., "Shape optimal design using B-splines," *Computer Methods in Applied Mechanics and Engineering*, Vol. 44, No. 3, Aug. 1984, pp. 247–267.
- [5] Lépine, J., Trépanier, J.-Y., and Pépin, F., "Wing aerodynamic design using an optimized NURBS geometrical representation," *The 38th AIAA Aerospace Sciences Meeting and Exhibit*, No. AIAA–2000–699, Reno, Nevada, 2000.
- [6] Kulfan, B. M. and Bussoletti, J. E., "Fundamental parametric geometry representations for aircraft component shapes," *AIAA-2006-6948*, Sept. 2006.
- [7] Kulfan, B. M., "A universal parametric geometry representation method-CST," *AIAA-2007-0062*, Jan. 2007, pp. 1–36.
- [8] Hicks, R. M. and Henne, P. A., "Wing design by numerical optimization," *Journal of Aircraft*, Vol. 15, No. 7, 1978, pp. 407–412.
- [9] Wu, H. H.-Y., Yang, S., Liu, F., and Tsai, H.-M. H., "Comparison of three geometric representations of airfoils for aerodynamic optimization," *16th AIAA Computational Fluid Dynamics Conference, Orlando, Florida*, 2003.
- [10] Sobieczky, H., "Parametric airfoils and wings," *Recent Development of Aerodynamic Design Methodologies*, Vol. 68, Springer, 1999, pp. 71–87.
- [11] Morris, A. M., Allen, C. B., and Rendall, T. C. S., "CFD-based optimization of aerofoils using radial basis functions for domain element parameterization and mesh deformation," *International Journal for Numerical Methods in Fluids*, Vol. 58, No. 8, Nov. 2008, pp. 827–860.
- [12] Morris, A. M., Allen, C. B., and Rendall, T. C. S., "Aerodynamic shape optimization of a modern transport wing using only planform variations," *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, Vol. 223, No. 6, 2009, pp. 843–851.
- [13] Allen, C. B. and Rendall, T. C. S., "CFD-based optimization of hovering rotors using radial basis functions for shape parameterization and mesh deformation," *Optimization and Engineering*, Vol. 14, No. 1, March 2013, pp. 97–118.
- [14] 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, pp. 1–8.
- [15] Poole, D. J., Allen, C. B., and Rendall, T. C. S., "Metric-based mathematical derivation of efficient airfoil design variables," *AIAA Journal*, Vol. 53, No. 5, 2015, pp. 1349–1361.
- [16] J. Toal, D. J., Bressloff, N. W., Keane, A. J., and E. Holden, C. M., "Geometric filtration using proper orthogonal decomposition for aerodynamic design optimization," *AIAA Journal*, Vol. 48, No. 5, 2010, pp. 916–928.

- [17] Jameson, A., "Aerodynamic design via control theory," *Journal of Scientific Computing, also ICASE Report No.88-64*, Vol. 3, 1988, pp. 233–260.
- [18] Beux, F. and Dervieux, A., "A Hierarchical Approach For Shape Optimization," *Engineering Computations*, Vol. 11, No. 1, 1993, pp. 25–48.
- [19] Desideri, J. and Zolesio, J., "Inverse shape optimization problems and application to airfoils," *Control and Cybernetics*, Vol. 34, No. 1, 2005, pp. 165.
- [20] 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.
- [21] 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.
- [22] 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.
- [23] Andreoli, M., Ales, J., and Désidéri, J.-A., "Free-form-deformation parameterization for multilevel 3D shape optimization in aerodynamics," 2003.
- [24] Duvigneau, R., Chaigne, B., and Désidéri, J.-A., "Multi-Level Parameterization for Shape Optimization in Aerodynamics and Electromagnetics using a Particle Swarm Optimization Algorithm," *Sophia*, 2006.
- [25] Duvigneau, R., "Adaptive parameterization using free-form deformation for aerodynamic shape optimization," 2006.
- [26] Abou, B., Majd, E. L., Duvigneau, R., and Désidéri, J.-a., "Aerodynamic Shape Optimization using a Full and Adaptive Multilevel Algorithm," *ERCOTAC 2006 Design Optimization Methods Application*, 2006, pp. 1–11.
- [27] 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.
- [28] 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.
- [29] Sherar, P. A., Thompson, C. P., 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.
- [30] Chaikin, G., "An algorithm for high-speed curve generation," *Graphical Models /graphical Models and Image Processing /computer Vision, Graphics, and Image Processing*, Vol. 3, No. 4, 1974, pp. 346–349.
- [31] Cashman, T. J., Hormann, K., and Reif, U., "Generalized Lane–Riesenfeld algorithms," *Computer Aided Geometric Design*, Vol. 30, No. 4, 2013, pp. 398–409.
- [32] Doo, D. and Sabin, M., "Behaviour of recursive division surfaces near extraordinary points," *Computer-Aided Design*, Vol. 10, No. 6, 1978, pp. 356–360.
- [33] 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.
- [34] Forsey, D. R. and Bartels, R. H., "Hierarchical B-spline refinement," *ACM SIGGRAPH Computer Graphics*, Vol. 22, ACM, 1988, pp. 205–212.
- [35] 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.
- [36] DeRose, T., Kass, M., and Truong, T., "Subdivision surfaces in character animation," *Proceedings of the 25th annual conference on Computer graphics and interactive techniques SIGGRAPH 98*, Vol. pages, No. Annual Conference Series, 1998, pp. 85–94.
- [37] Ma, W., "Subdivision surfaces for CAD - An overview," *Computer-Aided Design and Applications*, Vol. 1, No. 1-4, 2004, pp. 223–232.
- [38] Cashman, T. J., *NURBS-compatible subdivision surfaces*, Ph.D. thesis, Cashman, Thomas J., 2010.
- [39] Stollnitz, E. J., DeRose, T. D., and Salesin, D. H., *Wavelets for computer graphics : theory and applications*, Morgan Kaufmann Publishers, 1996.

- [40] Sabin, M., "Eigenanalysis and artifacts of subdivision curves and surfaces," *Tutorials on multiresolution in geometric modelling*, Springer, 2002, pp. 69–92.
- [41] 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.
- [42] Masters, D. A., Taylor, N. J., Rendall, T. C. S., Allen, C. B., and Poole, D. J., "Geometric Comparison of Aerofoil Shape Parameterization Methods," *AIAA Journal*, Jan. 2017, pp. 1–15.
- [43] 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.
- [44] Rendall, T. C. S. and Allen, C. B., "Unified fluid–structure interpolation and mesh motion using radial basis functions," *International Journal for Numerical Methods in Engineering*, Vol. 74, No. 10, 2008, pp. 1519–1559.
- [45] Wendland, H., *Scattered data approximation*, Cambridge University Press Cambridge, 2005.
- [46] Reuther, J., *Aerodynamic shape optimization using control theory*, Ph.D. thesis, University of California, Davis, 1996.
- [47] Kim, S., Alonso, J., and Jameson, A., "A gradient accuracy study for the adjoint-based Navier-Stokes design method," *37th Aerospace Sciences Meeting and Exhibit*, No. c, American Institute of Aeronautics and Astronautics, Reston, Virginia, Jan. 1999.
- [48] Kim, S., Alonso, J. J., Jameson, A., Usaf, A., and Issmo, N., "Aerodynamic Optimization Using the Continuous Adjoint Method Two-dimensional High-Lift Aerodynamic Optimization Using the Continuous Adjoint Method," *AIAA Aerospace Sciences Meeting*, , No. AIAA 2000-4741, 2000.
- [49] Papadimitriou, D. I. and Giannakoglou, K. C., "A Continuous Adjoint Method for the Minimization of Losses in Cascade Viscous Flows," *Current*, , No. January, 2006, pp. 1–11.
- [50] Ceze, M., Hayashi, M., and Volpe, E., "A Study of the CST Parameterization Characteristics," *27th AIAA Applied Aerodynamics Conference*, American Institute of Aeronautics and Astronautics (AIAA), June 2009.
- [51] Nadarajah, S. K., "Aerodynamic Design Optimization: Drag Minimization of the NACA 0012 in Transonic Inviscid Flow," .
- [52] Zhang, M., Rizzi, A. W., and Nangia, R. K., "Transonic Airfoils and Wings Design Using Inverse and Direct Methods," *53rd AIAA Aerospace Sciences Meeting*, , No. January, Jan. 2015, pp. 1–17.
- [53] Vassberg, J., Harrison, N., Roman, D., and Jameson, A., "A Systematic Study on the Impact of Dimensionality for a Two-Dimensional Aerodynamic Optimization Model Problem," *29th AIAA Applied Aerodynamics Conference*, June 2011.
- [54] Bisson, F. and Nadarajah, S., "Adjoint-Based Aerodynamic Optimization of Benchmark Problems," *52nd Aerospace Sciences Meeting*, , No. January, Jan. 2014, pp. 1–20.
- [55] Nadarajah, S., "Adjoint-Based Aerodynamic Optimization of Benchmark Problems," *53rd AIAA Aerospace Sciences Meeting*, Jan. 2015.
- [56] Poole, D. J., Allen, C. B., and Rendall, T. C. S., "Control Point-Based Aerodynamic Shape Optimization Applied to AIAA ADODG Test Cases," *53rd AIAA Aerospace Sciences Meeting*, Jan. 2015.
- [57] Anderson, G. R., Nemec, M., and Aftosmis, M. J., "Aerodynamic Shape Optimization Benchmarks with Error Control and Automatic Parameterization," *53rd AIAA Aerospace Sciences Meeting*, Jan. 2015.
- [58] Masters, D. A., Taylor, N. J., Rendall, T. C. S., and Allen, C. B., "Impact of Shape Parameterisation on Aerodynamic Optimisation of Benchmark Problem," *54th AIAA Aerospace Sciences Meeting*, Jan. 2016.
- [59] Lee, C., Koo, D., Telidetzki, K., Buckley, H., Gagnon, H., and Zingg, D. W., "Aerodynamic Shape Optimization of Benchmark Problems Using Jetstream," *53rd AIAA Aerospace Sciences Meeting*, 2015.
- [60] Meheut, M., Destarac, D., Ben Khelil, S., Carrier, G., Dumont, A., Peter, J., Anderson, G., Nadarajah, S., Poole, D., Vassberg, J., and Zingg, D. W., "Gradient-Based Single and Multi-points Aerodynamic Optimizations with the elsA Software," *53rd AIAA Aerospace Sciences Meeting*, , No. January, Jan. 2015, pp. 1–34.
- [61] 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.
- [62] Leger, J. C., "Menger curvature and rectifiability," *Annals of Mathematics*, Vol. 149, 1999, pp. 831–869.