Computers and Fluids 148 (2017) 137–156

Contents lists available at ScienceDirect

Computers and Fluids

journal homepage: www.elsevier.com/locate/compfluid

A volumetric geometry and topology parameterisation for fluids-based optimisation

J. Hall, T.C.S. Rendall*, C.B. Allen, D.J. Poole

Department of Aerospace Engineering, University of Bristol, BS8 1TR, UK

ARTICLE INFO

Article history:
Received 21 June 2016
Revised 16 December 2016
Accepted 3 February 2017
Available online 15 February 2017

Keywords:
Parameterization
Optimisation
Topology
Volume of solid
Level set
Phase-field
Material distribution

ABSTRACT

A new geometry and topology parameterisation method is presented which is based on creating a parameterisation grid of cells and reconstructing surfaces from the fraction of the cell volume defined to be solid, with the volume fractions acting as design variables. This method is able to include topological changes alongside fine-level geometric control, and therefore offers a significant increase in flexibility. In this work, the geometric capabilities of the method are confirmed by successfully constructing a variety of surfaces, using both arbitrary object outlines and aerofoil geometries. The method is then used in a range of optimisation problems covering the design of a coastal defence, increasing fluid damping within an oscillating box by the addition of baffles, and design of a multi-body configuration for minimum drag in supersonic flow. These problems demonstrate the benefits of a parameterisation for fluids modelling that is capable of topological changes and which can be used with global search as well as gradient-based methods.

1. Introduction

In optimisation problems a key challenge is how best to represent trial geometries in a way that maximises coverage of the design space and which promotes good convergence of the optimiser. Methods of representing geometries are often referred to as shape parameterisations as the geometry is represented by a vector of parameters; it is these parameters the optimisation method subsequently adjusts to produce new geometries. This process and the links between the methods are shown in Fig. 1.

Having the ability to change topology is a desirable but uncommon feature for a parameterisation scheme. Enabling topological change in an optimisation process is desirable because it allows consideration of designs that would not otherwise be accessible [1]. In the case of high lift configurations for aircraft or racing cars, for example, it is not obvious how many lifting elements may be optimal. Equally, the aerospike design used on Trident II missiles [2,3] would be very difficult to obtain without a parameterisation capable of substantial geometric and/or topological change.

The objective of the work presented is to define a parameterisation scheme that implicitly handles topological change alongside fine-level geometric control. Following a review of existing parameterisation methods in Section 1.1 a new parameterisation method is proposed in Section 2. The method is then applied to a set of geometry reconstruction problems in order to explore the accuracy and flexibility of the parameterisation, before testing on optimisation problems including the design of minimum drag configurations in supersonic flow, the design of a coastal defence and design of baffles to increase damping in an oscillating tank.

1.1. Existing parameterisation methods

Parameterisation raises two fundamental and interlinked problems; first, how to devise a geometry method that designs shapes suitable for the chosen optimisation, and second, how to have confidence that a sufficient but not excessive number of design variables have been introduced.

A parameterisation is a method of representing a geometry by means of a vector of design variables, and these methods can be broadly split between two groups. The first group seeks to *construct* shapes from empty space, while the second aims to *deform* an already existing geometry. Although distinctions can blur, surface point control, level sets and descriptive function approaches are constructive, while control point techniques are deformative (an exception to this is the way in which NURBS can be used to both construct surfaces from control points, and then also to deform them through motion of those control points). It is clear that a constructive approach may usually be linearised to provide a deformative route providing the initial shape may be encoded, but

* Corresponding author.
E-mail addresses: James.Hall@bristol.ac.uk (J. Hall), Thomas.Rendall@bristol.ac.uk, thomas.rendall@bristol.ac.uk (T.C.S. Rendall), C.B.Allen@bristol.ac.uk (C.B. Allen), d.j.poole@bristol.ac.uk (D.J. Poole).

http://dx.doi.org/10.1016/j.compfluid.2017.02.005
0045-7930/

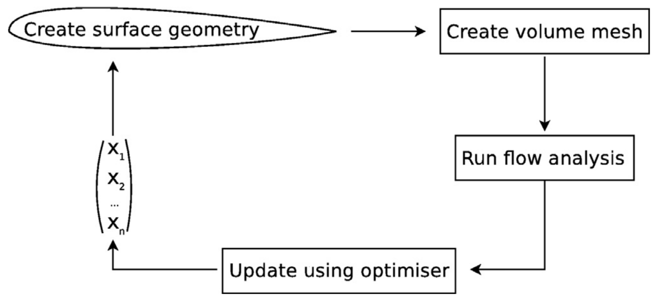

Fig. 1. Diagram of interaction between design variables, **x**, geometry parameterisation scheme, generated shape and an optimisation method.

there is no guarantee a deformative technique will be able to construct a shape from scratch.

Once a parameterisation has been selected, it is possible to consider design space (dimension) reduction (DSR) in an effort to ensure the number of design variables is minimised. The first approach, proposed by Diez et al. [4,5], is to randomly sample a wide range of geometries constructed via the chosen parameterisation and use a Karhunen-Loève expansion (although other matrix approaches may also be used [6]) to find a series of eigenvectors and values ordered in terms of decreasing importance. From this, a subset of the 'early' eigenvectors may then be used as design variables, and this approach has been applied for marine hull design [4,5] using free-form deformation. A second approach is to take a representative geometric library and parameterise this with any method through a fitting procedure, before interpolating the calculated parameter values against a chosen, smaller set of variables of interest. For example, an aerofoil library might be parameterised using CST [7], and then the CST parameters interpolated in terms of lift coefficient and thickness to chord ratio, as demonstrated by Sóbester and Powell [8]. It is an important result of these approaches that even if a parameterisation initially introduces too many design variables, these can later be reduced to an acceptable level through DSR. It is therefore not compulsory to address accuracy (ability to reconstruct any shape) and efficiency (ability to reconstruct using a small number of design variables) simultaneously, although clearly a link will always exist.

The goal in this section is to illustrate where an approach handling topological change as part of the parameterisation can fit in to the existing tool box of methods, so the relative merits of existing methods and their most important results shall now be considered.

1.1.1. Surface points as design variables

A natural approach is to consider moving every surface point, so that the design variables are simply the surface point locations in x,y,z. Doing so ensures that the complete design space is retained, as any surface can be represented (in a discrete sense) if all surface points are free to move.

This has been applied extensively in aerodynamic optimisation [9–13]. The gradient of the objective function must however be smoothed in order to ensure smooth surfaces are generated, as with such a high number of design variables is it easy to produce noisy surface shapes. It can also be demonstrated on a test example, such as the brachistochrone problem, that in the absence of a smoothing operation surface quality will progressively deteriorate [10].

The benefit of being able to represent a large design space creates difficulties as high dimensional spaces are time consuming to explore, with slow optimiser convergence being common, and it is imperative to use an adjoint route for gradient computation. Using this parameterisation it is also not clear how to create topological changes moving from the starting surface in the absence of any other descriptive mechanism; furthermore, including topological variation in an adjoint frame work would be difficult. Topological derivatives have been defined for general optimisation [14], but not applied with a surface point parameterisation.

1.1.2. Level sets

Level sets represent the boundary of a geometry as a level set of a function and their use in geometry optimisation is reviewed by van Dijk et al. [15]. For reasons of convenience the zero level set is usually chosen as this can be detected by a sign change. The use of level sets as a method of representing fronts was first suggested by Osher and Sethian [16] with applications to tracking the behaviour of propagating flame fronts, and this led naturally into Sethian and Wiegmann's [17] work using level sets to parameterise a geometry for optimisation. The level set method operates by evolving the level set function in a time like manner using the Hamilton–Jacobi equation such that at each time step a new geometry is produced as the position of the zero level set changes. The evolution of the level set function is controlled by a velocity term which acts normal to the surface and typically the level set function is initialised as a signed distance function from the initial surface geometry.

Level set based optimisation has been largely applied to steady state structural problems and often requires an adjoint solution. This is especially restrictive when considering the use of an optimisation scheme with a general solver as code specific modifications must be made to solve both the adjoint problem along with the original objective function evaluation.

Sethian and Wiegmann [17] calculated velocity directly from the local stress such that the boundary moves to remove the maximum amount of material subject to constraints; this approach was applied to the optimisation of a cantilever beam. Wang et al. [18] found velocities by using the adjoint method to calculate the sensitivity of the objective function to the geometry defined by the level set; this sensitivity was then used to define the level set velocity. Allaire et al. [19] present similar results but, as with methods described below in Section 1.1.5, used low density material rather than a true absence in regions defined as outside the part. Allaire et al. [20] show that level sets evolved by the velocity method are unlikely to add new holes and so create a different topology; this can be improved by the inclusion of the topological derivative [14]. Rather than move a boundary defined by a level set, Wei and Wang [21] define a piecewise constant level set function which has the advantage of easier topology changes through creation of holes.

In addition to the evolution based level set parameterisation, a less common explicit level set approach has been used in optimisation such as that suggested by Kreissl et al. [22] where weighted radial basis functions (RBFs) were used to represent the level set function with the weights acting as design variables. Kreissl et al. combined this parameterisation with a lattice Boltzmann flow solver to optimise pipe geometry for minimum pressure loss. This method allows standard gradient based solvers to be used with a level set approach but a large number of design variables are needed (on the order of 200) to produce a simple final geometry and as presented the method cannot be used to produce solid regions in areas that were previously entirely void. Further work by Kreissl and Maute [23] developed a framework based on the extended finite element method (XFEM) which permitted changing topologies to be included in an optimisation based on a Navier–Stokes solver, as XFEM permits discontinuities in shape functions.

1.1.3. Control point methods

An intuitive parameterisation is to have a surface constructed by joining a set of control points on that surface [24]. However,