Optimisation using smoothed particle hydrodynamics with volume-based geometry control

J. Hall¹ · T. C. S. Rendall¹ · C. B. Allen¹

Received: 13 September 2016 / Revised: 22 May 2017 / Accepted: 24 May 2017 / Published online: 7 July 2017
© The Author(s) 2017. This article is an open access publication

Abstract This work explores how smoothed particle hydrodynamics (SPH) may be applied for fluids optimisation problems. To achieve this, a newly developed volume of solid geometric parameterisation is applied that implicitly allows large geometric changes as well as topological changes. The meshless nature of SPH has long been an advantage, but when combined with the parameterisation presented here, optimisation calculations are able to make unlimited changes in the geometry without user intervention. To demonstrate the benefits this pairing of techniques affords for free-surface problems, three model optimisations and objective functions are considered: improvement of discharge coefficient through a nozzle, maximisation of damping in a pivoting tank and minimisation of wave overtopping for a simplified coastal defence. For the wave problem, varying constraints are explored and a time-recorded particle boundary condition is applied to accelerate the optimisation process. In each of the cases the optimisation finds a significant improvement in the objective function and shows how constraint selection influences the performance of the final design.

Keywords Optimisation · Parameterisation · Smoothed particle hydrodynamics · Volume of solid

1 Introduction

The aim of optimisation is to determine the topology and shape of a final design which maximises/minimises a chosen objective, for example minimising drag of an aerofoil, as well as to gain deeper understanding of the physics-based compromises involved during design. With continued improvements in computational capability, the use of automatic methods to find better performing designs is becoming more common. Applications within aerodynamics (Jameson et al. 1998), structural mechanics (Wang et al. 2003) and multidisciplinary design (Samareh 2000) are now well established. A variety of numerical methods are used in these processes, with finite volume and finite element methods particularly common. Meshless methods such as SPH have yet to be fully exploited, but offer access to a repertoire of physical behaviour for which mesh based methods can often be unsuitable. In particular, optimisation using mesh based Eulerian methods requires continuous remeshing, which is labour intensive and costly, or mesh deformation, which although fast and reliable will strongly limit the geometries that may be explored. In comparison SPH presents no such restriction and is straightforward to apply no matter how significantly the surface changes. The goal of this work is to demonstrate a framework in which SPH may operate for optimisation and the benefits that this affords, especially for free-surface problems, but doing so requires a great degree of flexibility in how the optimised shapes are defined or ‘parameterised’.

Numerical optimisation in fluids is a powerful technique but a critical aspect to the process shown in Fig. 1 is how the method converts between design variables, which are

✉ T. C. S. Rendall
thomas.rendall@bristol.ac.uk

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

Fig. 1 Interaction between vector of design variables, $\mathbf{x}$, geometry parameterisation scheme, generated shape and an optimisation method

usually real or integer numbers on which the optimiser operates, and the physical shape. For many general optimisations this physical parameterisation is simple or may not even be required (an example might be optimisation of resistors, capacitors or other components of an electrical circuit), but for fluids optimisation it is critical because the desired result is usually a shape. For example, NACA 4/5 series aerofoils are described in terms of thickness and maximum camber position together with underlying polynomials, but if this parameterisation is used in a numerical procedure it restricts the optimised shapes to be describable by that particular family of polynomials, and since the global optimum is unlikely to be a member of this family it effectively limits the quality of the result. This is not the only difficulty with parameterisation; even if the chosen design variables cover the complete design space very effectively, the number of steps required and the path taken to find the optimum are both intimately linked to the parameterisation. Slow convergence to an optimum, or detection of multiple local minima rather than the global minimum, can be caused by a poor parameterisation. The goal is therefore to cover the complete physical design space using as few design variables as possible and in a manner that maximises the chance of finding a global minimum. The volumetric parameterisation presented in this work is capable of representing both different topologies/design configurations (which in this work means different numbers of objects, but more generally this could also imply designs containing holes) as well as geometries (ie. the precise size, shape and position of each of the separate objects), a feature not available in many parameterisation schemes, which usually only represent differing geometries with the same topology. It also allows for large and small scale geometric changes, and for design constraints to be incorporated within the parameterisation itself rather than imposed mathematically on the optimiser (broadly, this is also true of any volume-based parameterisation).

Usage of SPH for numerical optimisation has been relatively restricted. Ha and Cho (2010) have shown how an SPH code utilising the projection technique for incompressibility can be differentiated to yield design sensitivities which allow for gradient based optimisation. This method for sensitivity calculation is utilised by Ha et al. (2011) to minimise the movement of points on a flexible sheet under the influence of a jet of water, with both the water jet and flexible sheet modelled with SPH, and the strain energy of a corrugated panel is also minimised. However, the method used to calculate the sensitivity to the design variables is not suitable for application with ‘black box’ SPH solvers as modifications must be made at a source code level to enable the solver to calculate the sensitivity. Additionally, it is doubtful if the sensitivity calculation method could be so easily applied to the common weakly compressible form of SPH. Work in structural optimisation has been accomplished by Yamada et al. (2013) using the moving particle semi-implicit (MPS) method (similar to SPH but with a different differential operator), where a level set parameterisation was applied to optimise a structure under transverse loading to arrive at a truss-like configuration. Son et al. (2015) used a structural SPH model for optimisation of a shield to protect against space debris, using panel thicknesses as design variables and using a design of experiment/response surface technique to find an optimum.

The three optimisation case studies considered in this work are selected to demonstrate both the capabilities of the parameterisation and the flexibility that SPH possesses for free-surface problems. The first problem considered is optimisation of a nozzle to raise the discharge coefficient at low Reynolds numbers, which is relevant for a range of industrial and manufacturing processes, and is primarily a steady case. The second problem, to improve the damping of a sloshing box, has many analogues with motion of fuel tanks or fluid-filled dampers and broadens the physics to include unsteadiness and coupled fluid-structural behaviour.

The final problem of minimising overtopping for a seawall, which is again unsteady, also demands access to wide variety of different shapes to be successful.

2 Volume fraction parameterisation

The majority of existing geometry parameterisations used in aerodynamics either seek to represent the surface directly by building it from splines or other functions (Hicks and Henne 1978; Kulfan and Bussoletti 2006; Leung and Zingg 2012) or deforming an existing geometry (Gagnon and Zingg 2015; Morris et al. 2008). The parameterisation presented here is inspired by front reconstruction techniques used in the (VOF) method (Hirt and Nichols 1981) for free surface simulation and the successful usage of the level set approach in optimisation (van Dijk et al. 2013; Kreissl et al. 2011; Yaji et al. 2014). In contrast to the VOF method this parameterisation method uses the volume fraction of solid to construct surfaces; the design variables being the volume fractions of solid in the cells making up a parameterisation mesh or grid in this work. Compared to a level set route based on the Hamilton-Jacobi equation (HJE), the technique used here constructs surfaces from the scalar function directly rather than by evolving a front. The principle of using a scalar field is also known as a ‘material distribution’ or ‘phase-field’ method within fluids (Kreissl et al. 2011), and as a ‘density-based’ (Deaton and Grandhi 2014; van Dijk et al. 2013) method in terms of structural optimisation.

The mechanism by which multiple topologies can be represented can be understood by imagining two regions of cells in the parameterisation separated by some distance, a situation illustrated in Fig. 2. Cells in these regions are set as solid by having a volume fraction of solid equal to one, while the cells separating the two regions have a volume fraction set to zero in order to be completely void. It is then clear that two separate surfaces should be constructed. By subsequently changing the volume fraction of solid in the red cells in Fig. 2 these two regions can later be joined if desired. Once the shape of the desired boundary is known, a set of discrete points is placed on the boundary in order to enforce the solid wall boundary condition in the subsequent SPH calculation.

Variations on the level set method have also been developed that extract the surface directly from the volume function rather than by evolving the HJE, although these methods can become quite distant to the original HJE basis and might better be considered as contour finding methods (Prilepov et al. 2013), which have a useful background in computer science. The method used here resides in this wider landscape of level set and contour approaches. However, in contrast to the radial basis functions used in some explicit level set methods (Kreissl et al. 2011), the shape

Fig. 2 Volumetric parameterisation grid with volume fraction value shown in greyscale

functions used here are Hermite splines defined on mesh edges, which provide excellent smoothness for fluids analysis. The interpretation of the volume of solid function is also more literal; although not enforced to a high numerical tolerance, matching the volume of created objects closely to the volumes defined on the parameter grid is what determines the resulting shape. Furthermore, because a detailed interpolation of the volume function is built within the cell, geometry may be extracted on a scale significantly smaller than the parameter cell and surface detail retained. This allows detailed shapes to be built with fewer design variables (which are equivalent to the parameter grid cells), thereby improving the efficiency of the design space coverage.

Kreissl et al. (2011) have applied a non-HJE level set approach together with an XFEM (Kreissl and Maute 2012) formulation for optimising low Reynolds number flows. XFEM was used to allow precise enforcement of the wall boundary condition for the finite-element Navier-Stokes model. In this sense the objective of the work presented here is similar, however, a different philosophy is taken. Conventional boundary fitted meshes have a long and successful history in fluids analysis, and it would be preferable to retain this route in order to ensure current numerical methods may be reused, without adopting XFEM methods. The goal is therefore to extract exact boundaries, that may either be discretised for use with a method such as SPH, or to make meshes for finite-volume methods.

For the VOF method (Hirt and Nichols 1981) the conservation of fluid volume is of vital importance and adequate results can be produced without smooth surfaces, however when parameterising a geometry for optimisation the volume fraction can be considered merely a convenient design variable and it is not necessary to ensure that any constructed surface strictly conserves the volume fraction in the

Fig. 3 Greyscale plot of underlying volume fraction data and reconstructed surface

cell. A far more important consideration is the smoothness of the surface, in particular ensuring that it is continuous, a property that the reconstructions used for the VOF method often lack.

The volumetric parameterisation used here builds on the work of Prilepov et al. (2013) in the field of computer graphics and is described in detail by Hall et al. (2015). The goal of the approach is to construct a contour of a scalar function, defined on a fixed ‘parameter grid’, which then becomes a closed body or group of bodies to be analysed to determine the objective function. Given a grid of cells, each containing a value of this volume fraction between zero and one, the method finds the contour. This process is made more complicated as a consequence of the properties commonly required for numerical fluids analysis; the surface must not contain discontinuities in either position or first or second derivatives. The requirements for optimisation are also important. Principally, if creation of new objects is to be permitted, the parameterisation should smoothly create a new object in any parameter cell where there is a non-zero value of the volume function, while the created object may potentially be much smaller than the size of the parameter cell that contains it. Objects of finite size appearing or disappearing in a discontinuous manner would be likely to interfere with convergence of any gradient-based optimiser, as the gradients would be ill-defined, and should be avoided. These requirements drive the design of the contour-finding method.

Figure 3 gives an example showing the volume cells as a greyscale image with the constructed surface as a smooth curve overlaid, while Fig. 4 illustrates the capability for fitting more complicated shapes, including those with holes. The steps in producing a surface from a set of volume fractions can be summarised as follows:

1. Nodally average the cell volume fractions ϕ , shown on the parameter grid in Fig. 5, to volume cell vertices.

Fig. 4 Example images and fitted contours (note: target contour was extracted manually from image)

- This permits later interpolation of the volume fraction between the cell vertices and within the interior of the cell, which is necessary in order to determine if a query point is inside or outside the object
2. Construct approximate gradients $\nabla\phi$ of the volume fraction (shown red in Fig. 5) at the vertices using a volume/boundary integral via Green’s theorem. This gradient is required to provide an accurate volume fraction interpolation within the cells
3. Perform an interpolation, using the vertex values of volume fraction ϕ and volume fraction gradient $\nabla\phi$, over each volume cell to construct a continuous volume fraction function. In this case a Hermite spline (Farin 1993) (shown blue in Fig. 5) is used along edges, followed by a Coons patch (Coons 1967) across each parameter cell. The Coons patch blends smoothly between the Hermite splines that have been constructed along opposing edges of the cell, so that a value of $\phi(x, y)$ may be found at any point within the cell
4. In each cell, points where the value of the volume fraction function $\phi(x, y)$ is higher than a threshold are considered to be within a surface while lower values are outside. For each cell choose the threshold value ϕ_{thresh} such that the fraction of the cell area inside the surface is equal to the (fixed) volume fraction ϕ_n for that cell. The selection process is accomplished by using a Cartesian grid for the sampling points (shown as crosses in Fig. 5),

Fig. 5 Stages in construction of solid object from volume of solid parameter grid

followed by sorting these points based on their volume fraction value and choosing a threshold fraction value that places the correct proportion of them within the object, ie. select the value of ϕ_{thresh} that most closely solves $\frac{n_{inside}(\phi_{thresh})}{n_{total}} \approx \phi_n$ where n_{total} is the total number of sample points used in that cell. The value of ϕ_{thresh} is somewhat truncated due to the integer values of n_{inside} and n_{total} , but precise recovery of ϕ is not required

- Construct surfaces using either a marching squares or triangles method based on the sample points that are now known to reside inside the object. This takes the (inside) sampled points (crosses circled red in Fig. 5) and draws a boundary along their outer extremity, which is the final edge of the object

Figure 3 shows an example wall geometry constructed during optimisations from the greyscale VOS shown. The technique produces a smooth shape with relatively few design variables, but the potential to achieve substantial changes is clear, especially as the number of cells is raised.

3 SPH method

There are two approaches for producing SPH formulations of the governing equations. The first, more generally used, views SPH as a method of representing any given function

by summation over a set of particles, while the second starts from a Lagrangian (Price 2012). Liu and Liu (2005) give details of the former approach, where the process is to take a function that approximates the Dirac delta function and convolve it with the function to be represented, giving

$$\langle f(\mathbf{x}) \rangle = \int_{\Omega} f(\mathbf{x}') W(\mathbf{x} - \mathbf{x}', h) d\mathbf{x}' \quad (1)$$

The angle brackets denote the integral representation, and the objective is to approximate $\frac{\nabla p}{\rho}$ in order to be able to integrate forwards (in inviscid flow) $\frac{D\mathbf{u}}{Dt} = -\frac{\nabla p}{\rho}$. $f(\mathbf{x})$ is a function to be represented by integration over volume Ω , $W(\mathbf{x} - \mathbf{x}', h)$ is a smoothing function that depends on the difference between the location of interest and the points over which the integral is evaluated. h is a characteristic length of W and is known as the smoothing length.

This integral representation is then approximated by summation over N nearby particles, leading to

$$\langle f(\mathbf{x}) \rangle \approx \sum_{j=1}^N f(\mathbf{x}_j) W(\mathbf{x} - \mathbf{x}_j, h) \Delta V_j \quad (2)$$

This summation is possible because W is usually chosen to be a compact function which is non-zero only for a small fraction of the total domain. The radius over which W is non-zero is known as the support radius, thus N is the

number of particles within the smoothing length. If each particle is assigned a mass and a density then (2) can be written as

$$\langle f(\mathbf{x}) \rangle \approx \sum_{j=1}^N \frac{m_j}{\rho_j} f(\mathbf{x}_j) W(\mathbf{x} - \mathbf{x}_j, h) \quad (3)$$

Derivatives can be approximated in the same way

$$\langle \nabla f(\mathbf{x}) \rangle = \int_{\Omega} \nabla f(\mathbf{x}') W(\mathbf{x} - \mathbf{x}', h) d\mathbf{x}' \quad (4)$$

However, while the derivative of f may not be known the derivative of the smoothing function is, so the following identity may be applied

$$\nabla[f(\mathbf{x}') W(\mathbf{x} - \mathbf{x}', h)] = \nabla f(\mathbf{x}') W(\mathbf{x} - \mathbf{x}', h) + f(\mathbf{x}') \nabla W(\mathbf{x} - \mathbf{x}', h) \quad (5)$$

Hence

$$\langle \nabla f(\mathbf{x}) \rangle = \int_{\Omega} \nabla[f(\mathbf{x}') W(\mathbf{x} - \mathbf{x}', h)] - f(\mathbf{x}') \nabla W(\mathbf{x} - \mathbf{x}', h) d\mathbf{x}' \quad (6)$$

And by applying the divergence theorem the first volume integral becomes a surface integral over the problem domain S

$$\langle \nabla f(\mathbf{x}) \rangle = \int_S f(\mathbf{x}') W(\mathbf{x} - \mathbf{x}', h) d\mathbf{S} - \int_{\Omega} f(\mathbf{x}') \nabla W(\mathbf{x} - \mathbf{x}', h) d\mathbf{x}' \quad (7)$$

Usually the support domain is entirely within the problem domain such that the surface integral is zero (as $W = 0$ at the edge of the support) and the representation can be expressed entirely in terms of the volume integral. This is not true in the case where the support domain overlaps the boundary (either a free surface or a fixed boundary Ferrand et al. 2013) and has been the focus of substantial development effort in the field.

The particle approximation then gives

$$\langle \nabla f(\mathbf{x}_i) \rangle \approx - \sum_{j=1}^N \frac{m_j}{\rho_j} f(\mathbf{x}_j) \nabla W(\mathbf{x}_i - \mathbf{x}_j, h) \quad (8)$$

With this information, and substituting (9) (note that alternative identities may be used, see for example Price (2012)) in (8)

$$\frac{\nabla p}{\rho} = \nabla \left(\frac{p}{\rho} \right) + \frac{p}{\rho^2} \nabla \rho \quad (9)$$

it is possible to construct the governing equations in the SPH formulation to give

$$\begin{aligned} \frac{\nabla p}{\rho} &\approx - \sum_{j=1}^N m_j \left[\frac{p_j}{\rho_j^2} \right] \nabla_i W_{ij} - \sum_{j=1}^N m_j \left[\frac{p_i}{\rho_i^2} \right] \nabla_i W_{ij} \\ &= - \sum_{j=1}^N m_j \left[\frac{p_j}{\rho_j^2} + \frac{p_i}{\rho_i^2} \right] \nabla_i W_{ij} \end{aligned} \quad (10)$$

where on the right hand side, $f = \frac{p}{\rho}$ and $f = \rho$ can be identified and used to construct the approximation from (8). When alternative identities are employed different formulations arise for (8) (Price 2012).

The final SPH formulation used in this work is

$$W_{ij} = \frac{7}{4\pi h^2} \left(1 - \frac{|\mathbf{x}_i - \mathbf{x}_j|}{2h} \right)^4 \left(2 \frac{|\mathbf{x}_i - \mathbf{x}_j|}{h} + 1 \right), \quad \frac{|\mathbf{x}_i - \mathbf{x}_j|}{h} \leq 2 \quad (11)$$

$$\begin{aligned} \frac{D\mathbf{u}_i}{Dt} &= - \sum_{j=1}^N m_j \left[\frac{p_j}{\rho_j^2} + \frac{p_i}{\rho_i^2} \right] \nabla_i W_{ij} + \\ &\sum_{j=1}^N m_j \nu \frac{\rho_i + \rho_j}{\rho_i \rho_j} \frac{\mathbf{x}_{ij} \cdot \nabla_i W_{ij}}{|\mathbf{x}_{ij}|^2 + 0.001h^2} \mathbf{u}_{ij} + \mathbf{g} \end{aligned} \quad (12)$$

$$\frac{D\rho_i}{Dt} = \sum_{j=1}^N m_j (\mathbf{u}_i - \mathbf{u}_j) \cdot \nabla_i W_{ij} \quad (13)$$

$$p = B \left(\left(\frac{\rho}{\rho_0} \right)^{\gamma} - 1 \right) \quad (14)$$

where a Newmark integration is applied to advance particle positions.

Equation (11) is the quintic kernel suggested by Wendland (1995). The equation of state, (14), is commonly used in SPH to remove the need to solve a pressure Poisson equation, and an equation of this form was first suggested by Cole (1948) for water, before Monaghan (1994) then suggested setting B such that the speed of sound in the fluid is an order of magnitude greater than the fastest particle velocity. γ here is set to 7 following Cole's (1948) work. This is known as the weakly compressible method, and is able to predict free surface motion (Monaghan 1994; Lee et al. 2008). The second term in (12) is a viscosity term proposed by Morris et al. (1997) which helps prevent non-physical particle interpenetration.

An alternative to the weakly compressible form of SPH is incompressible SPH (ISPH), implemented either through solution of a pressure Poisson equation (Lind et al. 2012; Skillen et al. 2013; Shao and Lo 2003) or by fixing fluid volumes to be constant through Lagrange multipliers (Ellero et al. 2007). Indeed, this was the route taken by Ha et al. (2011) in their optimisation. Exact enforcement

of incompressibility can bring improvements in timestep and smoother pressure fields (especially when combined with shifting approaches that beneficially adjust particle positions (Lind et al. 2012)), avoiding the noise-polluted fields that commonly result when using weakly compressible SPH. For future optimisation work, and especially any that required a noise-free pressure field, ISPH would be a suitable route to follow.

The final task is to enforce the correct solid wall boundary condition on the fluid. There is a significant body of work addressing boundary conditions (see for example Bierbrauer et al. 2009) in SPH and approaches of varying accuracy and complexity are available. Here, wall boundary conditions are set using fixed (ie. non-moving) SPH particles distributed along the edge of the shape calculated from the volume of solid parameterisation. This has the advantage of simplicity, but comes at an expense in terms of accuracy, as the particles will ‘stand off’ the solid wall by a distance on the order of the smoothing length. An advantage of SPH is that no special treatment of the free surface is required, with the pressure naturally dropping to zero at this interface (gas pressures are not modelled here).

4 Optimisation

Optimisation is the process of finding a maximum or minimum in an objective function with respect to the design variables. From this point on, minimisation problems shall be discussed but it should be understood that any maximisation problem can be easily converted to a minimisation problem and vice versa. The function to be minimised is often referred to as the objective function and this convention will be followed here; other names include the ‘fitness function’ or ‘cost function’. Optimisation problems are frequently constrained, that is to say that not all inputs or outputs are allowable (‘feasible’), and this means that the objective function must now be minimised while satisfying the imposed conditions which are often of the form of an inequality or equality. An example constrained problem would be to minimise the drag, the objective function, over a surface in a flowing fluid while ensuring that the volume enclosed by the surface is greater than a given value, an inequality constraint.

The conjugate gradient (CG) method (Hestenes and Stiefel 1952) used in this work incorporates information from previous steps with the current gradient to form the search direction. Each new search direction must be conjugate to all previous directions. The merits of a CG method are its simplicity and relatively good performance, crucially it does not require the evaluation of the Hessian. The precise algorithm used is that of Press et al. (1992) with adaptation to include barrier functions for the imposition of constraints

on the optimisation problem and the addition of a variable size initial guess for the line searching algorithm. Constraints are applied to input design variables by projecting them onto the allowable space. Although the conjugate gradient method of optimisation is well known it is presented here for clarity.

The conjugate gradient method can solve exactly, linear systems of equations, $\mathbf{Ax} = \mathbf{b}$ by minimisation of the quadratic form:

$$f(\mathbf{x}) = \frac{1}{2} \mathbf{x}^T \mathbf{Ax} - \mathbf{x}^T \mathbf{b} \quad (15)$$

If the objective function can be approximated by an equation of this form then the conjugate gradient method can be applied to find the minimum, with $\mathbf{x}$ the vector of volume of solid values in every parameter cell. An iterative process is followed where an initial guess of the minimum solution, or indeed any an arbitrary point, is selected $\mathbf{x}_0$. The minimum along the local steepest descent direction, $\mathbf{p}_0$, is then found by evaluating the gradient of $f(\mathbf{x})$, $\mathbf{p}_0 = -\nabla f(\mathbf{x}_0)$. This minimum is found by a line search algorithm detailed below. After this point the line minimum is searched for, not along the direction of steepest descent but rather along a direction conjugate to all previous search directions. Once the line minimum is found, a new search direction is calculated at the point given by the line minimisation; each new search direction is calculated according to

$$\mathbf{p}_i = -\nabla f(\mathbf{x}_i) + \beta_i \mathbf{p}_{i-1} \quad (16)$$

$$\beta_i = \frac{\nabla f(\mathbf{x}_i) \cdot (\nabla f(\mathbf{x}_i) - \nabla f(\mathbf{x}_{i-1}))}{\nabla f(\mathbf{x}_{i-1}) \cdot \nabla f(\mathbf{x}_{i-1})} \quad (17)$$

β_i can take a number of forms but in this work the suggestion of Polak and Ribière (1969) in (17) is followed.

Along each search direction the objective function must be minimised and this is accomplished using the method suggested by Brent (1973), which can be applied as each line search is a one-dimensional problem with the objective function depending on the scalar that multiplies the search direction. First the minimum must be bracketed by two values which requires evaluating the objective function at a minimum of three points, but likely at rather more. The bracketing will yield two points which are guaranteed not to be the minimum and one point which is, at worst, an approximation to the true minimum. The search procedure uses a mixture of a golden section search and parabolic interpolation between previously evaluated points along the line to find the minimum which may be local rather than global.

Penalty functions are employed to enforce constraints as the conjugate gradient method does not naturally allow for this. Constraints here may for example be the design variable range (0 to 1) for each parameter cell, or physical

Fig. 6 Paramaterisation for optimisation of discharge nozzle. Filled cells constrained to be solid in order to fix total width

constraints (such as the total volume of the object). The form used here is

$$g(s) = As^2 \quad (18)$$

and the value of A in (18) is chosen dependent on the optimisation problem and the constraint applied. s is the distance from the edge of the valid region to the actual value of the constrained variable. The penalty function is applied to the evaluated value of the objective function as a scaling factor

$$f(x) = f(x) \times (g(s) + 1) \quad (19)$$

In the event of multiple constraint violations the maximum possible value of the penalty functions is used.

Gradients were evaluated in this work using central finite differences to avoid direction bias, and where a central difference was not feasible (for example when filling a new region, which can only use a positive increment), one-sided differences were used instead. The perturbation size required was selected by running a sweep in perturbation size, and selecting a region where the resulting gradient is independent of the perturbation size. Finite-differencing for gradients is computationally costly, but this is mitigated by a very simple parallelisation process (all perturbation runs here were independent). Since the problems addressed in this work were unsteady, a single perturbation is a relatively costly complete unsteady simulation, which motivates the use of an efficient parameterisation to maximise the topologies and shapes available while limiting the number of design variables needed. Fluids-based objectives are purely a function of the design variables via the SPH solver,

Fig. 7 Initial and optimised nozzle geometries

but are of course time-consuming implicit functions to evaluate, and geometric objectives (such as the sectional area) are explicitly dependent on the design variables.

5 Results

To demonstrate the capabilities of the parameterisation method optimisation cases focussing on design of a nozzle, damped sloshing box (fluid-structure coupled) and coastal defence will now be explored.

Fig. 8 Optimiser history for discharge nozzle

Fig. 9 Zoomed view of initial, intermediate and final nozzle shapes

5.1 Nozzle optimisation

This test cases explores increasing the mass flow rate out of a two-dimensional tank by modifying a nozzle geometry. This two-dimensional problem has been chosen to demonstrate the potential that optimisation, especially when combined with the geometry parameterisation presented in this work, has to improve performance while still being a manageable problem from a computational cost perspective. The parameterisation method is particularly useful here, as it is a volume-based method, and therefore allows imposition of constraints on the nozzle geometry without the need to apply those constraints within the optimiser.

The objective function to maximise for this problem is the steady state mass flow rate through the nozzle, $\dot{m}$, which is a function of the vector of volumes of solid in the n parameter cells, $\mathbf{x} \in \mathbb{R}^n$. The optimisation is unconstrained, except by the extent of the parameterisation grid and the restriction to have a fixed width nozzle (imposed using the fixed parameterisation cells shown filled) as illustrated in Fig. 6. The mass flow rate is calculated by a two-dimensional SPH simulation with pressure applied by a constant force piston in order to generate a constant applied driving pressure, with the mass flow rate calculated from an average flow rate taken once the transient behaviour died away. So, the problem is

$$\begin{aligned} &\text{maximise}_{\mathbf{x} \in \mathbb{R}^n} \dot{m}(\mathbf{x}) \\ &\text{Unconstrained} \end{aligned}$$

The initial tank geometry is shown in Fig. 6; the parameterisation grid, consisting of a grid of seven rows by five columns, is shown overlaid on the schematic representation of the tank. The central row of cells, shown filled in solid black in Fig. 6, is forced to have a volume fraction of 1 at all times in the optimisation and in this way the width of the nozzle constrained at a constant value throughout the optimisation. The rows of cells immediately above and below the fixed cells are initially set to solid, so that the initial

(a) Settled fluid volume

(b) Calculation of contained settled fluid

(c) Final filling using settled fluid particles

Fig. 10 Fluid initialisation procedure before each objective function evaluation

geometry is as shown in Fig. 7a, but they are subsequently free to change. Any geometry created on the left side is mirrored on the right hand edge of the nozzle, but only one parameterisation grid used in order to reduce the number of design variables.

By examining the objective function percentage change shown in Fig. 8 it can be seen that the optimisation process has produced a geometry that offers significant improvements over the base design. This has been accomplished by smoothing the transition between the tank and the nozzle, which can be seen in Figs. 7b and 9. Here the nozzle geometry has been modified by the removal of the sharp corners at the inlet and exit of the nozzle; these have been replaced by geometry that extends slightly above and below the initial width of the nozzle and which forms a smooth curve

Fig. 11 Paramaterisation for optimisation of sloshing box. Note: parameterisation grid is mirrored about a vertical axis positioned at the horizontal mid-point of the box

towards the minimum nozzle diameter point, set by the parameterisation grid. The replacement of the 90° corners in the initial design reduces pressure losses associated with such abrupt flow turning.

5.2 Sloshing tank optimisation

5.2.1 Problem set-up

Fuel slosh has been shown to be important in the dynamics of a number of vehicle types; spacecraft (Abramson 1966; Slabinski 1978; Vreeburg 2005) tanker lorries (Sankar et al. 1992) and ships (Kim et al. 2007). Also, tuned liquid dampers are tanks, partially full of liquid, whose purpose is to increase the damping of a structure by using the free surface motion of the fluid to counter any excitation. These devices have found application in reducing the vibration in tall towers such as the Hobart Tower in Tasmania and the Gold Tower in Japan (Kareem and Kijewski 1999). A related class of devices are the anti-roll tanks that can be employed on ships (Gawad et al. 2001; Bulian et al. 2010). There is therefore a strong case for numerically exploring the design space available for achieving fluid damping, and this work considers the effect of modifying the liquid damping behaviour of a pivoting tank by optimising the external boundary.

In order to explore how the damping characteristics of a pivoting tank can be altered, a model problem previously developed and experimentally validated by the authors (Hall et al. 2015) was used, consisting of (in the undeformed case) a square, partially filled thin (ie. two-dimensional) box pivoting about the middle of its upper edge. The model was fully fluid-structure coupled in a strong manner, with the aim of optimisation being to modify the external geometry of the tank in order to maximise mechanical damping of the

Fig. 12 Optimised shapes and flows for sloshing box

complete system. The initial angle of the tank to the vertical was 20° and the tank $\frac{1}{4}$ full. This initial condition was selected to give representative non-linear flow features such as breaking waves. In contrast to the experimental validation presented in Hall et al. (2015), no damping was used at the pivot so the forces on the tank are entirely due to gravity and the internal fluid forces; this allows the effect of the fluid to be isolated from any other effects.

The objective function to be minimised by the optimiser is the total variation of angular position, given below. Here t_2 is five seconds after the tank was released, $t_1 = 0$ and ω is the angular velocity of the tank. The optimisation is constrained to not allow the tank area A to change more than 5% from the initial value A_0 with an additional constraint

Fig. 13 Objective function history for sloshing box

on the gravitational potential energy of the fluid E that it not reduce by more than 10% of the initial value E_0 , giving

$$\begin{aligned} &\text{minimise } \int_{t_1}^{t_2} |\omega| dt \\ &0.95A_0 \leq A \leq 1.05A_0 \\ &0.9E_0 \leq E \end{aligned}$$

5.2.2 Fluid initialisation

SPH is primarily an unsteady method. Without sacrificing the Lagrangian framework and moving to an arbitrary Lagrangian-Eulerian technique where some ‘slip’ is allowed between particles and the flow, there is no steady way to resolve even a fundamentally steady flow (such as that past an aerofoil, or a ship hull in the absence of breaking waves). Even hydrostatic equilibrium represents a challenge as the particle positions that satisfy this equilibrium must be found by marching an unsteady calculation to a steady state (usually accelerated with roughly critical damping). During the optimisation a large number of different tank designs are produced, each with a different geometry that must be filled to the selected depth with fluid in hydrostatic equilibrium (‘settled’). However, producing this settled state in a timely manner for an arbitrary tank geometry is challenging.

The usual method for initialising (by running an unsteady computation until a steady solution is reached) is computationally time consuming and a new settling process is required every time the geometry is changed. This is particularly expensive for optimisation problems because as Fig. 1 shows every new shape, and every new perturbation from that shape, will require an individual settling computation *in addition* to the unsteady runs that are already required. Colagrossi et al. (2012) present a method of initialising particles based upon performing an initial evolution of an equation that is not physical but rather designed to reduce anisotropy. Although Colagrossi et al. (2012) use a more suitable equation for settling the fluid particles than the governing fluid equations the method still requires the evolution in time of a solution with all the associated calculation this requires.

Here a novel method is used which relies upon a pre-initialised background state which can then be applied to rapidly initialise a new state with a different geometry. A large tank of fluid is initialised once by time marching, having an extent greater than any anticipated geometry that may be initialised from it, as in Fig. 10a. Figure 10b shows how when a new settled state is needed the boundary geometry of the new state is moved into the settled tank until a specified volume of the new object is filled. The particles within the new geometry are then retained while all others are deleted, including those internal to the new geometry or within a smoothing length of walls, as illustrated in Fig. 10c. Finally, the particles making up the boundary of the new geometry must be initialised appropriately with suitable values of density and, through the equation of state (14), pressure. This can be done by a weighted average over nearby fluid particles. The desired unsteady simulation can then be run. This initialisation procedure does not produce precisely hydrostatic solutions but the generated particle distributions are very close to equilibrium and there is minimal particle movement once particles are allowed to evolve under the influence of the governing equations.

The advantage of this approach is that it requires a single, larger settling calculation instead of $\mathcal{O}(100)$ marginally smaller ones. The settling calculation must be for a slightly larger volume of fluid so that even the largest shape can be initialised, but this small extra cost is negligible compared

Fig. 14 Schematic of optimisation problem. Design region shown hatched

Fig. 15 Diagram of reflected wave interaction

to the repeated savings within the gradient and optimisation loops.

5.2.3 Results

Figure 11 shows how the parameterisation scheme was used starting from a four by four grid of square volume cells located in the bottom corner of the tank. These design variables are mirrored around the vertical centreline of the tank in order to create a symmetric design and reduce the number of design variables needed. This means only the lower half of the box is optimised, with the constraint that it must be left-right symmetric. The reconstructed surface is constrained to lie within the volume cells and as such the VOS parameterisation grid limits the extent of the geometry. This feature is a strength, as it allows geometric constraints to be applied without needing to implement constraints in the optimisation routine. The upper half of the box in Fig. 11 was also constructed from a VOS grid, but these VOS cells contained fixed values in order to produce a suitable upper shape and were not part of the optimisation. Hence, the upper half of the box is not precisely square as reconstruction of exactly straight edges requires adjusting the VoS values iteratively, or further refinement of the parameter grid, and was not considered necessary for this case. Similarly, corners are sharp to a degree that depends on the refinement level of the parameter grid.

Figure 12a shows the evolution of the tank geometry throughout the optimisation process. The geometry changes consist mostly of moving the bottom corners of the tank outwards. The optimisation increases the width of the tank

higher up the sides of the tank as well as near the corners. This widening of the tank will change the time that waves take to travel from one side of the tank to the other.

Figure 12b shows the behaviour of the fluid at a chosen instant just as the tank begins to pivot back towards the right. In the initial configuration the fluid hits the left hand wall at this point whereas in the final configuration the wave is still some distance from hitting the wall, due to the increased width of the tank that the optimisation process has produced.

The time history of tank angle is shown in Fig. 12c and objective function changes in Fig. 13, showing how the optimised tank geometry improves behaviour. The different behaviour between the two designs is clearly seen at the first negative peak in the time history where the final design changes direction faster and then has a lower amplitude positive angle peak. This is due to the later wave impact on the left hand wall, meaning that the final design can change direction without having to act against the wave impacting on the wall. The wave then impacts the wall later, thereby reducing the amplitude of the next peak.

5.3 Coastal defence

A coastal defence structure can take many forms and serve many purposes such as the prevention of erosion, the protection of land and structures behind the defence or the creation of areas of calmer water. To demonstrate optimisation, the type of coastal defence considered in this work is of the seawall type, with the goal being to produce a defence minimising the volume of over-topping water. Although a specific type of coastal defence is considered here, the

Fig. 16 Plot of all particles at the start of the recording period. Particles to be recorded lie within the highlighted x locations

Fig. 17 Recorded particles in red with remaining simulated particles

methods employed are generalisable to a variety of situations by variation of the objective function, wave action and constraints.

The conjugate gradient method described in Section 4 was used to perform optimisation with the VOS parameterisation method employed to generate the coastal defence geometry from the design parameters. The objective function was then evaluated for each design using the (SPH) method described in Section 3.

A number of methods exist to generate oncoming waves but here this is achieved by using a piston type wave maker with the same geometry as used by Monaghan (1994), shown schematically in Fig. 14, where the left most wall acts as the piston wave maker. A wave maker type arrangement was chosen over a simpler method of wave generation such as a dam break as it would provide conditions closer to those likely to be encountered in reality. The parameterisation is allowed to generate the geometry of the coastal defence in the region of Fig. 14 shaded with hatched lines.

The geometry of the problem consists of a beach with a one in ten slope and a still water height, D , of five metres. Between the start of the beach and the wave maker is a region with a flat bed ten metres long. The piston amplitude is 5m and the period of the paddle motion is 5s.

The selection of objective function and constraints is of key importance in using optimisation in design tasks, as it is the method through which design preferences and physical limitations can be expressed. In this work three different optimisation cases are considered with varying objective functions, constraints and starting conditions for the optimisation in order to demonstrate the range of designs achievable by modifying the optimisation problem.

Table 1 Objective function reductions for coastal defence optimisations (for case 3, the quoted figure refers to the fluid component of the weighted objective function)

Case	% reduction
1	–37
2a	–37
2b	–70
3	–44

Two objective functions are considered in this work, the first is that of minimising the total amount of fluid passing a datum line over a set period of time from left to right (‘overtopping volume’), given by the overtopping function

$$\mathcal{O}_{so}(\mathbf{x}) = \int_{t=0}^{t=t_{\text{end}}} \left(\sum_{i=1}^N m_i \frac{\max(V_x, 0)}{V_x} \delta(x - x_{\text{datum}}) \right) dt \quad (20)$$

with subscript so for single objective. Here the datum line is placed at the rear of the area in which the parameterisation can create geometry, as shown in Fig. 14, where the datum line is dashed. The amount of fluid remaining within the coastal defence geometry at the end of the simulation period is added to the fluid that crosses the datum to prevent any benefit from ‘trapping’ fluid in the structure. The second objective function is a multi-objective optimisation made up of the overtopping fluid volume plus an additional objective to minimise the enclosed area of the design, to use the least possible material. A single objective function is calculated by taking an equally weighted sum of both the overtopping fluid volume and the material volume (area in 2D), each normalised by their initial value, to give

$$\mathcal{O}_{mo}(\mathbf{x}) = \frac{1}{2} \left(\frac{\mathcal{O}_{so}}{\mathcal{O}_{so0}} + \frac{V}{V_0} \right) \quad (21)$$

This is referred to below in case 3 as the ‘multi-objective function’, and balances the overtopping objective function equally with the volume of material used.

The applied forces are known but the effect on the structure is also important; practical coastal defence structures must be able to resist the forces applied during use. An accurate assessment of the stresses experienced throughout the structure can only come from a numerical structural analysis; this is, however, prohibitively expensive in an optimisation context as the objective function is evaluated many times in the process making the penalty of any additional computational cost severe. For this reason a simplified model for the stress is used where it is assumed that providing the structure meets a minimum thickness at any given point, it will be able to resist the applied loads. This thickness can be tuned based on the problem and the experience of the designer.

Starting condition (i) for the parameter grid is with the bottom three rows of the parameterisation grid cells set

Fig. 18 Convergence history for a conjugate gradient optimisation of a coastal defence for case 1

to have a volume fraction of 1 such that a solid block is constructed across the full width of the parameterisation grid, while starting condition (ii) is formed by setting the two rightmost cells in the bottom three rows of the parameterisation grid to be solid such that a solid block is constructed across only part of the parameterisation grid width (these are initial configurations only, and are subsequently free to change via the optimiser). In addition, although the width of the parameterisation cells is kept constant, the parameterisation cell height was increased by a factor of 1.4 in specified cases when compared to the height used for the first starting condition. These objective functions and starting conditions are combined to form the following optimisation cases:

Case 1 The overtopping objective function including any trapped fluid with a thickness constraint and starting from a full width block (starting condition (i)), giving

$$\begin{aligned} &\text{minimise } \mathcal{O}_{so} \\ &t_{min} \leq t \end{aligned}$$

Case 2a Same as case 1, but starting from a partial width defence (starting condition (ii)).

Case 2b Same as case 1, but starting from a partial width defence (starting condition (ii)) and with $1.4\times$ higher parameterisation cells.

Case 3 The multi-objective objective function with a thickness constraint and starting from a partial width

Fig. 19 Case 1 final design

block (starting condition (ii)) with higher parameterisation cells. This case strikes a compromise between preventing overtopping and using a small material volume for the structure.

$$\begin{aligned} &\text{minimise } \mathcal{O}_{mo} \\ &t_{min} \leq t \end{aligned}$$

A single wave impact was chosen as it dramatically reduces the computational time needed to evaluate the flow over each trial design, however the drawback of this is that the behaviour of the defence under successive impacts is not evaluated and there is a risk that the defence will be optimised to a form where overtopping is reduced by storing a certain amount of fluid in the defence structure before it crosses the datum line. If this design is used in a situation with multiple wave impacts, this storage space will not be available after the first impact so the performance of the defence will be greatly reduced on the subsequent impacts. In order that the optimisation process avoids these designs, the amount of fluid contained within the coastal defence structure that is not draining away at the end of the simulation period is added to the fluid that has crossed the datum.

5.3.1 Computational cost reduction

Although the wave maker method of generating waves will produce more accurate results than a simple dam break, it suffers from the problem that it requires far more SPH particles to implement as a large area is required between the wave maker and the coastal defence in order for the flow to develop realistically. The computational requirements are of particular importance in an optimisation problem as the flow solver must be run many times in order to produce an optimised design.

Here the computational cost is first reduced by only simulating from the time just before the wave impact on the coastal defence, with the state at that time found from a prior simulation. The second measure to reduce the computational cost is to reduce the total number of particles by implementing a new boundary condition which allows the particles nearest the wave maker to be removed.

In this application, the only details of the flow that are important are the interaction with the coastal defence;

Fig. 20 Comparison of final coastal defence shapes

regions further away are only important in so far as they influence the flow near the coastal defence. Indeed, far enough upstream from the coastal defence location the flow is unaffected by whatever coastal defence geometry the flow is interacting with, assuming the simulation time is sufficiently short to prevent reflected waves propagating far enough upstream. This can be seen in Fig. 15 where a wave that has just been produced by interaction with the coastal defence travels to the left. As this wave has a finite speed, the fluid around the datum line A-A will have a period of time where it is unaffected by the reflected wave. As the particles at this upstream location have behaviour independent of the coastal defence and the input from the wave maker is identical these particles can be forced such that they follow the behaviour imposed by the wave maker.

In practice this is accomplished by running a simulation without a coastal defence but with the same wave maker parameters, amplitude and frequency, as will be used in the optimisation simulations. Particles upstream of the coastal defence location, between the two initial x locations in Fig. 16, are then recorded in all their essential parameters (position, velocity, density and pressure) over the length of the simulation. The initial location of the particles to be recorded is shown in Fig. 16. For future simulations particles initially upstream of these recorded particles are deleted from the simulation and the recorded particles are forced to have their recorded properties at each time step over the

total simulation time; Fig. 17 shows the particles that were recorded and gives an idea of the number of particles that no longer need to be simulated.

5.3.2 Optimised designs

Case 1 Case 1 achieves a 37% reduction in the objective function from the initial value (see Table 1). The convergence history of the optimisation is shown in Fig. 18, indicating that the majority of the gains are realised after five optimisation steps.

Inspecting geometry in Fig. 19 reveals that the structural constraints have had the desired effect, as the structure is free of unrealistically narrow sections and the amount of fluid remaining trapped within the defence is small. However, the full width and height of the parameterisation grid have been used, suggesting that for this case there was little to be gained from adjusting the shape of the defence; rather, the determining factor was simply scale. As the maximum height of the defence is relatively low, the oncoming flow depth becomes higher than the maximum defence height shortly after the wave front hits the defence. This makes it very difficult for the optimiser to achieve much change in the objective function as the majority of the flow over the defence happens without direct interaction with the defence.

Starting from a design where the full width of the bottom of the parameterisation grid is solid makes finding a local minimum more likely. This is because when starting from a full width defence, moving the defence right requires removing material from the left hand side of the wall which will likely result in a (temporarily) worse design if the material is removed from the top of the wall.

Case 2a/2b In case 1, the design was dominated by moving to the largest possible structure, suggesting two modifications to the calculation in order to explore the design space. Firstly, starting condition (ii) was tested to give case 2a, which initialised the body from solid cells on the lower right corner of the parameter grid, and secondly the parameter grid height was increased by a factor of 1.4 to give case 2b.

Only altering the starting condition is enough to produce a shape comparable to common coastal defence designs. When starting only from the lower right corner case 2a

Fig. 21 Case 2b final design

Fig. 22 Convergence history for a conjugate gradient optimisation of a coastal defence for cases 2a and 2b

finds the geometry shown in Figs. 20 and 21, which is still recovered if the parameter grid height is increased and starting condition (ii) is still used. By initialising the geometry from the far right side of the parameterisation grid, geometry is built from the lower right corner in a manner that offers greater potential for modifying the geometry of the face that meets the oncoming wave, whilst avoiding local minima on the way. In effect, there is less of a penalty associated with adding material. The final design shown in Fig. 21 demonstrates the more complex left hand face which is now capable of affecting the flow to a greater extent than when starting from a full-width design. This curved front face geometry has the effect of redirecting the oncoming flow back on itself and in so doing reducing the amount of fluid flowing over the top of the defence.

Figure 22 shows the percentage reduction in the objective function for cases 2a and 2b. As shown in Table 1 case 2a gives a similar reduction to case 1 (but with a very different final design), while case 2b gives a reduction of 70%, which is larger than that achieved for either case 1 or case 2a due to the extra height. The design evolutions for case 2a are shown in Fig. 23, illustrating how the concave surface evolves. The convergence history shown in Fig. 22 shows convergence for case 2b after only a few steps, much fewer than seen in case 1 (Fig. 18) or case 2a and suggesting that this optimisation problem is better posed, or perhaps simpler, as a result of the taller parameter grid.

Case 3 In this multi-objective optimisation problem the overall percentage reduction in the composite objective function is less than that obtained in the single objective

Fig. 23 Design evolutions for case 2a

case 2; an expected result as both objectives must be balanced against each other. The convergence history is shown in Fig. 24 which shows that although the final reduction is lower than in case 2 it is reached in a smaller number of design evolutions. Comparing the separate objective functions, the area objective is reduced by 18% while the fluid flow objective is reduced by a larger 44% (Table 1), indicating that the area objective is the harder objective to improve and that, as the fluid objective does not reach the case 2 value, the overall limitation is on reducing the fluid objective while also reducing the area objective.

The final design shown in Fig. 25 shows a similar design philosophy to that seen in the case 2 design (Fig. 21) where a concave region is formed to redirect the flow. However in this case, the requirement to minimise the area of material used leads to the removal of material at the rear as well as a slightly different frontal region.

Fig. 24 Convergence history for a conjugate gradient optimisation of a coastal defence using the composite objective function for case 3

Fig. 25 Case 3 final design

6 Conclusions

A volume based parameterisation scheme has been coupled to SPH to demonstrate optimisation of a nozzle for maximum discharge, coupled fluid-structure sloshing behaviour for maximising damping and a coastal defence for minimising overtopping. The parameterisation scheme relieves any constraints on topology, and SPH is able to exploit this through its meshless structure. Despite the significant geometric changes in the coastal defence problem, the optimisation process is rapid and requires little intervention.

Initialisation of SPH for a wide range of shapes during optimisation can be achieved through use of a large reservoir of settled particles, followed by a containment search to retain only those particles within the interior of the geometry. Acceleration of the coastal optimisation process for SPH, which is primarily an unsteady method, can be achieved through use of a time-recorded boundary condition to eliminate the need for a sizeable number of particles.

The nozzle results indicate a smooth contouring of the inlet can produce an improvement in discharge, while damping from sloshing can be increased through widening of the base of the tank. The coastal defence problem illustrates three different designs that improve upon the initial design, showing how the objective functions, constraints and the starting conditions of the optimisation problem can be varied in order to meet the needs of specific design problems.

Future work will include applying more realistic stress constraints and extending the method to be three-dimensional, allowing more realistic designs that take account of flow parallel to the defence structure. The influence of increasing the design space resolution must also be explored to ensure convergence of the optimised geometries as the number of design variables is raised. Unfortunately, global refinement of the parameter mesh would lead to an unacceptable increase in cost (both from the finite difference gradients and from slower optimiser convergence), so it is likely a progressive refinement approach (Masters et al. 2006) will be followed instead.

Acknowledgments The authors would like to acknowledge the funding for James Hall from the EPSRC DTP EP/K502996/1 at the University of Bristol. Data access: the data necessary to support the conclusions are included in the paper.

The authors gratefully acknowledge the notable contributions made by the anonymous reviewers to the quality and interpretation of the test cases and results.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Abramson HN (1966) The dynamic behavior of liquids in moving containers, with applications to space vehicle technology. Report SP-106, NASA
- Bierbrauer F, Bollada PC, Phillips TN (2009) A consistent reflected image particle approach to the treatment of boundary conditions in smoothed particle hydrodynamics. *Comput Methods Appl Mech Eng* 198(41–44):3400–3410
- Brent RP (1973) Algorithms for minimization without derivatives. Prentice-Hall Inc
- Bulian G, Souto-Iglesias A, Delorme L, Botia-Vera E (2010) Smoothed particle hydrodynamics (SPH) simulation of a tuned liquid damper. *J Hydraul Res* 48(sup1):28–39
- Colagrossi A, Bouscasse B, Antuono M, Marrone S (2012) Particle packing algorithm for {SPH} schemes. *Comput Phys Commun* 183(8):1641–1653
- Cole RH (1948) Underwater explosions. Princeton Univ. Press, Princeton
- Coons SA (1967) Surfaces for computer-aided design of space forms. Technical report, Massachusetts Institute of Technology
- Deaton JD, Grandhi RV (2014) A survey of structural and multidisciplinary continuum optimisation: post 2000. *Struct Multidiscip Optim* 49:1–38
- Ellero M, Serrano M, Espaol P (2007) Incompressible smoothed particle hydrodynamics. *J Comput Phys* 226(1):1731–1752
- Farin G (1993) Curves and surfaces for CAGD a practical guide, 3rd edn. Academic Press Inc, London
- Ferrand M, Laurence DR, Rogers BD, Violeau D, Kassiotis C (2013) Unified semi-analytical wall boundary conditions for inviscid, laminar or turbulent flows in the meshless SPH method. *Int J Numer Methods Fluids* 71(4):446–472
- Gagnon H, Zingg DW (2015) Two-level free-form and axial deformation for exploratory aerodynamic shape optimization. *AIAA J* 53(7):2015–2026
- Gawad AF, Ragab SA, Nayfeh AH, Mook DT (2001) Abdel Roll stabilization by anti-roll passive tanks. *Ocean Eng* 28(5): 457–469
- Ha YD, Cho S (2010) Efficient design sensitivity analysis of incompressible fluids using SPH projection method. *Struct Multidiscip Optim* 40(1–6):307–318

- Ha YD, Kim M-G, Kim H-S, Cho S (2011) Shape design optimization of sph fluid-structure interactions considering geometrically exact interfaces. *Struct Multidiscip Optim* 44(3):319–336
- Hall J, Poole DJ, Rendall T, Allen CB (2015) Volumetric shape parameterisation for combined aerodynamic geometry and topology optimisation. In: 16th AIAA/ISSMO multidisciplinary analysis and optimization conference, AIAA AVIATION Forum (AIAA 2015-3354). doi:[10.2514/6.2015-3354](https://doi.org/10.2514/6.2015-3354)
- Hall J, Rendall TCS, Peel H, Allen CB (2015) A multi-physics computational model of fuel sloshing effects on aeroelastic behaviour. *J Fluids Struct* 56:11–32
- Hestenes MR, Stiefel E (1952) Methods of conjugate gradients for solving linear systems. *J Res Natl Bur Stand* 49(6):409–436
- Hicks RM, Henne PA (1978) Wing design by numerical optimization. *J Aircr* 15(7):407–412
- Hirt CW, Nichols BD (1981) Volume of fluid (VOF) method for the dynamics of free boundaries. *J Comput Phys* 39:201–225
- Jameson A, Martinelli L, Pierce NA (1998) Optimum aerodynamic design using the Navier–Stokes equations. *Theor Comput Fluid Dyn* 10(1–4):213–237
- Kareem A, Kijewski T (1999) Mitigation of motions of tall buildings with specific examples of recent applications. *Wind And Structures* 2(3):201–251
- Kim Y, Nam BW, Kim DW, Kim YS (2007) Study on coupling effects of ship motion and sloshing. *Ocean Eng* 34:2176–2187
- Kreissl S, Maute K (2012) Levelset based fluid topology optimization using the extended finite element method. *Struct Multidiscip Optim* 46:311–326
- Kreissl S, Pingen G, Maute K (2011) An explicit level set approach for generalised shape optimization of fluids with the lattice-Boltzmann method. *Int J Numer Methods Fluids* 496–519
- Kulfan BM, Bussolotti JE (2006) “Fundamental” parametric geometry representations for aircraft component shapes. In: 11th AIAA/ISSMO multidisciplinary analysis and optimization conference: the modeling and simulation Frontier for multidisciplinary design optimization. Boeing Commercial Airplane Group Seattle, Washington. AIAA2006-6948
- Lee E-S, Moulinec C, Xu R, Violeau D, Laurance D, Stansby P (2008) Comparisons of weakly compressible and truly incompressible algorithms for the SPH mesh free particle method. *J Comput Phys* 227(18):8417–8436
- Leung TM, Zingg DW (2012) Aerodynamic shape optimization of wings using a parallel newton-krylov approach. *AIAA J* 50(3):540–550
- Lind SJ, Xu R, Stansby PK, Rogers BD (2012) Incompressible smoothed particle hydrodynamics for free-surface flows: a generalised diffusion-based algorithm for stability and validations for impulsive flows and propagating waves. *J Comput Phys* 231(4):1499–1523
- Liu GR, Liu MB (2005) Smoothed particle hydrodynamics: a meshfree particle method. World Scientific Publishing Co
- Masters DA, Taylor NJ, Rendall T, Allen CB (2006) Progressive subdivision curves for aerodynamic shape optimisation. In: 54th AIAA aerospace sciences meeting, AIAA SciTech Forum (AIAA 2016-0559). doi:[10.2514/6.2016-0559](https://doi.org/10.2514/6.2016-0559)
- Monaghan JJ (1994) Simulating free surface flows with SPH. *J Comput Phys* 110:399–406
- Morris AM, Allen CB, Rendall TCS (2008) CFD-based optimization of aerofoils using radial basis functions for domain element parameterization and mesh deformation. *Int J Numer Methods Fluids* 58:827–860
- Morris JP, Fox PJ, Yi Z (1997) Modeling low reynolds number incompressible flows using sph. *J Comput Phys* 136:214–226
- Polak E, Ribière (1969) Note sur la convergence de méthodes de directions conjuguées. *Revue française d’Informatique et de Recherche opérationnelle Série Rouge* 3(1):35–43
- Press WH, Teukolsky SA, Vetterling WT, Flannery BP (1992) Numerical Recipes in C. Cambridge University Press
- Price DJ (2012) Smooth particle hydrodynamics and magnetohydrodynamics. *J Comput Phys* 231:759–794
- Prilepov I, Obermaier H, Deines E, Garth C, Joy KI (2013) Cubic gradient-based material interfaces. *IEEE Trans Vis Comput Graph* 19(10):1687–1699
- Samareh J (2000) Multidisciplinary aerodynamic-structural shape optimization using deformation (massoud). In: Multidisciplinary analysis optimization conferences. American Institute of Aeronautics and Astronautics
- Sankar S, Ranganathan R, Rakheja S (1992) Impact of dynamic fluid slosh loads on the directional response of tank vehicles. *Veh Syst Dyn* 21(6):385–404
- Shao S, Lo EYM (2003) Incompressible SPH method for simulating newtonian and non-newtonian flows with a free surface. *Adv Water Resour* 26(7):787–800
- Skillen A, Lind SJ, Stansby PK, Rogers BD (2013) Incompressible Smoothed Particle Hydrodynamics (SPH) with reduced temporal noise and generalised Fickian smoothing applied to body-water slam and efficient wave-body interaction. *Comp Methods Appl Mech Eng* 265:163–173
- Slabinski VJ (1978) INTELSAT IV in-orbit liquid slosh tests and problems in the theoretical analysis of the data. In: Flight Mechanics/Estimation Theory Symposium, pp 183–221
- Son BJ, Yoo J, Lee M (2015) Optimal design of a composite space shield based on numerical simulations. *J Mech Sci Technol* 29(12):5299–5308
- van Dijk NP, Maute K, Lanelaar M, van Keulen F (2013) Level-set methods for structural topology optimization: a review. *Struct Multidiscip Optim* 34:437–472
- Vreeburg JPB (2005) Measured states of Sloshsat FLEVO. National Aerospace Laboratory NLR NLR-TP-2005-518. <http://hdl.handle.net/10921/551>
- Wang MY, Wang X, Guo D (2003) A level set method for structural topology optimization. *Comput Methods Appl Mech Eng* 192(1–2):227–246
- Wendland H (1995) Piecewise polynomial, positive definite and compactly supported radial functions of minimal degree. *Adv Comput Math* 4:389–396
- Yaji K, Yamada T, Yoshino M, Matsumoto T, Izui K, Nishiwaki S (2014) Topology optimization using the lattice Boltzmann method incorporating level set boundary expressions. *J Comput Phys* 274:158–181
- Yamada T, Manabe M, Izui K, Nishiwaki S (2013) A topology optimization method for geometrically nonlinear problems incorporating level set boundary expressions and a particle method. *J Adv Mech Design Syst Manuf* 7(4):630–643