Fireshape: a shape optimization toolbox for Firedrake

Alberto Paganini¹ · Florian Wechsung²

Received: 17 May 2020 / Revised: 22 October 2020 / Accepted: 4 December 2020 / Published online: 11 February 2021
© The Author(s) 2021

Abstract

We introduce Fireshape, an open-source and automated shape optimization toolbox for the finite element software Firedrake. Fireshape is based on the moving mesh method and allows users with minimal shape optimization knowledge to tackle with ease challenging shape optimization problems constrained to partial differential equations (PDEs).

Keywords Shape optimization · Finite elements · Open-source software

1 Introduction

One of the ultimate goals of structural optimization is the development of fully automated software that allows users to tackle challenging structural optimization problems in the automotive, naval, and aerospace industries without requiring deep knowledge of structural optimization theory. The scientific community is working actively in this direction, and recent years have seen the publication of educational material that simplifies the understanding of structural optimization algorithms and guides the development of related optimization software. These resources are based on different models, such as moving mesh methods (Allaire and Pantz 2006; Dapogny et al. 2018; Courtais et al. 2019; Allaire 2007), level-sets (Laurain 2018; Allaire et al. 2002; Bernland et al. 2019), phase fields (Garcke et al. 2015; Dondl et al. 2019; Blank et al. 2014), and SIMP¹ (Sigmund 2001; Talischi et al. 2012; Bendsoe 1989; Bendsoe and Sigmund 2004), and are implemented in various software

environments such as Matlab (Sigmund 2001; Talischi et al. 2012), FreeFem++ (Allaire and Pantz 2006; Dapogny et al. 2018; Allaire et al. 2002), OpenFOAM CFD (Courtais et al. 2019), FEMLAB (Liu et al. 2005), and FEniCS (Laurain 2018; Dokken et al. 2019), to mention just a few.

In this work, we introduce Fireshape: an automated shape optimization library based on the moving mesh approach that requires very limited input from the user. Shape optimization refers to the optimization of domain boundaries and plays an important role in structural design. For instance, shape optimization plays a crucial role in the design of airfoils (He et al. 2019; Wang et al. 2018; Schmidt et al. 2013) and boat hulls (Lombardi et al. 2012; Marinić-Kragić et al. 2016; Abramowski and Sugalski 2017). Shape optimization is also a useful refinement step to be employed after topology optimization (Bendsoe and Sigmund 2004, Ch. 1.4). Indeed, topology optimization allows more flexibility in geometric changes and it is a powerful tool to explore a large design space. However, topology optimization may return slightly blurred (grey-scale) and/or staircase designs (Bendsoe and Sigmund 2004). By adding a final shape optimization step, it is possible to post-process results computed with topology optimization and devise optimal designs with sharp boundaries and interfaces, when this is necessary.

Fireshape is based on the moving mesh shape optimization approach (Allaire 2007, Ch. 6). In this approach, geometries are parametrized with meshes that can be arbitrarily precise and possibly curvilinear. The mesh nodes and faces are then optimized (or “moved”) to minimize a chosen target function. Fireshape has been developed on the moving mesh approach because the latter has a very neat

¹The acronym SIMP stands for Solid Isotropic Material with Penalisation.

Responsible Editor: Gregoire Allaire

✉ Alberto Paganini
a.paganini@leicester.ac.uk

¹ School of Mathematics and Actuarial Science, University of Leicester, University Road, Leicester, LE1 7RH, UK

² Courant Institute of Mathematical Sciences, New York University, 251 Mercer St., New York, NY 10012, USA

interpretation in terms of geometric transformations and is inherently compatible with standard finite element software, as it relies on the canonical construction of finite elements via pullbacks (Paganini et al. 2018). The main drawback of the moving mesh approach is that it does not allow topological changes in a straightforward and consistent fashion. However, Fireshape has been developed to facilitate shape optimization, and topology optimization is beyond its scope.

Fireshape couples the finite element library Firedrake (Rathgeber et al. 2016; Luporini et al. 2017; Homolya et al. 2018; Homolya et al. 2017; Gibson et al. 2020) with the Rapid Optimization Library (ROL) (Ridzal and Kouri 2014). Fireshape allows decoupled discretization of geometries and state constraints and it includes all necessary routines to perform shape optimization (geometry updates, regularization terms, geometric constraints, etc.). To solve a shape optimization problem in Fireshape, users must describe the objective function and the eventual constraints using the Unified Form Language (UFL) (Alnæs et al. 2014), a language embedded in Python to describe variational forms and their finite element discretizations that is very similar to standard mathematical notation. Once objective functions and constraints have been implemented with UFL, users need only to provide a mesh that describes the initial design and, finally, select their favorite optimization algorithm from the optimization library ROL, which contains algorithms for unconstrained, bound constrained, and (in-)equality constrained optimisation problems. In particular, users do not need to spend time deriving shape derivative formulas and adjoint equations by hand, nor worry about their discretization, because Fireshape automates these tasks.

Typically the computational bottleneck of PDE constrained optimization code lies in the solution of the state and the adjoint equation. While Fireshape and Firedrake are both written in Python, to assemble the state and adjoint equations the Firedrake library automatically generates optimized kernels in the C programming language. The generated systems of equations are then passed to the PETSc library (also written in C) and can be solved using any of the many linear solvers and preconditioners accessible from PETSc (Balay et al. 2019; Balay et al. 1997; Dalcin et al. 2011; Chevalier and Pellegrini 2008; Amestoy et al. 2000; Hypré 2020). This combination of Python for user facing code and C for performance critical parts is well established in scientific computing as it provides highly performant code that is straightforward to develop and use. Finally, we mention that Fireshape, just as Firedrake, PETSc and ROL, supports MPI parallelization and hence can be used to solve even large-scale three-dimensional shape optimization problems. In particular, Firedrake has been used to solve systems with several billion degrees of freedom on supercomputers with tens of thousands of cores (see,

e.g., Mitchell and Müller 2016; Farrell et al. 2019), and we emphasize that any PDE solver written in Firedrake can be used for shape optimization problems with Fireshape.

Fireshape is an open-source software licensed under the GNU Lesser General Public License v3.0. Fireshape can be downloaded at <https://github.com/Fireshape/Fireshape>.

Its documentation contains several tutorials and is available at <https://fireshape.readthedocs.io/en/latest/>. To illustrate Fireshape's capabilities and ease of use, in Section 2 we provide a tutorial and solve a three-dimensional shape optimization problem constrained by the nonlinear Navier-Stokes equations. The shape optimization knowledge required to understand this tutorial is minimal. In Section 3, we provide another tutorial based on a shape optimization problem constrained to the linear elasticity equation. In Section 4, we describe in detail the rich mathematical theory that underpins Fireshape. In Section 5, we describe Fireshape's main classes and Fireshape's extended functionality. Finally, in Section 6, we provide concluding remarks.

2 Example: shape optimization of a pipe

For this hands-on introduction to Fireshape, we first consider a viscous fluid flowing through a pipe Ω (see Fig. 1) and aim to minimize the kinetic energy dissipation into heat by optimizing the design of Ω . This is a standard test case (Zhou and Li 2008; Villanueva and Maute 2017; Feppon et al. 2020; Dilgen et al. 2018; Dapogny et al. 2017). To begin with, we need to describe this optimization problem with a mathematical model.

We assume that the fluid velocity $\mathbf{u}$ and the fluid pressure p satisfy the incompressible Navier-Stokes equations (Elman et al. 2014, Eq. (8.1)), which read

$$-2\nu\nabla\cdot\varepsilon\mathbf{u}+(\mathbf{u}\cdot\nabla)\mathbf{u}+\nabla p=0\text{ in }\Omega, \quad (1a)$$

$$\operatorname{div}\mathbf{u}=0\text{ in }\Omega, \quad (1b)$$

$$\mathbf{u}=\mathbf{g}\text{ on }\partial\Omega\setminus\Gamma, \quad (1c)$$

$$p\mathbf{n}-2\nu\varepsilon\mathbf{u}\mathbf{n}=0\text{ on }\Gamma, \quad (1d)$$

Fig. 1 Viscous fluid flows through a pipe Ω from the left to the right. A poor pipe design can lead to an excessive amount of kinetic energy being dissipated into heat

where ν denotes the fluid viscosity, $\varepsilon \mathbf{u} = \varepsilon(\mathbf{u}) = \frac{1}{2}(\nabla \mathbf{u} + \nabla \mathbf{u}^\top)$, $\nabla \mathbf{u}$ is the derivative (Jacobian matrix) of $\mathbf{u}$ and $\nabla \mathbf{u}^\top$ is the derivative transposed, Γ denotes the outlet, and the function $\mathbf{g}$ is a Poiseuille flow velocity (Elman et al. 2014, p. 122) at the inlet and zero on the pipe walls. In our numerical experiment, the inlet is a disc of radius 0.5 in the xy -plane, and $\mathbf{g}(x, y, z) = (0, 0, 1 - 4(x^2 + y^2))^\top$ on the inlet.

To model the kinetic energy dissipation, we consider the diffusion term in (1a) and introduce the function

$$J(\Omega) = \int_{\Omega} \nu \varepsilon \mathbf{u} : \varepsilon \mathbf{u} \, dx, \quad (2)$$

where the colon symbol $:$ denotes the Frobenius inner product, that is, $\varepsilon \mathbf{u} : \varepsilon \mathbf{u} = \text{trace}((\varepsilon \mathbf{u})^\top \varepsilon \mathbf{u})$.

Now, we can formulate the shape optimizations problem we are considering. This reads:

$$\text{Find } \Omega^* \text{ such that } J(\Omega^*) = \inf J(\Omega) \text{ subject to (1).} \quad (3)$$

To make this test case more interesting, we further impose a volume equality constraint on Ω . Otherwise, without this volume constraint, the solution to this problem would be a pipe with arbitrarily large diameter.

In the next subsections, we explain step by step how to solve this shape optimization problem in Fireshape. To this end, we need to create: a mesh that approximates the initial guess Ω (Section 2.1), an object that describes the PDE constraint (1) (Section 2.2), an object that describes the objective function (2) (Section 2.3), and, finally, a “main file” to set up and solve the optimization problem (3) with the additional volume constraint (Section 2.4).

2.1 Step 1: Provide an initial guess

The first step is to provide a mesh that describes an initial guess of Ω . For this tutorial, we create this mesh using the software Gmsh (Geuzaine and Remacle 2009). The initial guess employed is sketched in Fig. 1. The only mesh detail necessary to understand this tutorial is that, in Listing 1, the number 10 corresponds to the inlet, whereas the boundary flags 12 and 13 refer to the pipe’s boundary. For the remaining geometric details, we refer to the code archived on Zenodo (2020).

2.2 Step 2: Implement the PDE constraint

The second step is to implement a finite element solver of the Navier-Stokes equations (1). To derive the weak formulation of (1a), we multiply (1a) with a test (velocity) function $\mathbf{v}$ that vanishes on $\partial\Omega \setminus \Gamma$, and (1b) with a test (pressure) function q . Then, we integrate over Ω , integrate by parts (Elman et al. 2014, Eq. (3.18)), and impose the

```

1 from firedrake import *
2 from fireshape import PdeConstraint

3 class NavierStokesSolver(PdeConstraint):
4     """Incompressible Navier-Stokes equations."""

5     def __init__(self, mesh_m, viscosity):
6         super().__init__()
7         self.mesh_m = mesh_m
8         self.failed_to_solve = False

9         # Setup problem, Taylor-Hood finite elements
10        self.V = VectorFunctionSpace(mesh_m, "CG", 2) \
11            * FunctionSpace(mesh_m, "CG", 1)

12        # Preallocate soln variables for state eqn
13        self.solution = Function(self.V, name="State")
14        self.testfunction = TestFunction(self.V)

15        # Define viscosity parameter
16        self.viscosity = viscosity

17        # Weak form of incompressible Navier-Stokes equations
18        u, p = split(self.solution)
19        v, q = split(self.testfunction)
20        nu = viscosity # shorten notation
21        self.F = nu*2*inner(sym(grad(u)), sym(grad(v))) *
22            - p*div(v)*dx + div(u)*q*dx \
23            + inner(dot(grad(u), u), v)*dx

24        # Dirichlet boundary conditions
25        X = SpatialCoordinate(mesh_m)
26        rsq = X[0]**2+X[1]**2 # squared radius
27        uin = as_vector([0, 0, 1-4*rsq])
28        bc1 = DirichletBC(self.V.sub(0), 0., [12, 13])
29        bc2 = DirichletBC(self.V.sub(0), uin, [10])
30        self.bcs = [bc1, bc2]

31        # PDE-solver parameters
32        self.params = {"snes_max_it": 100,
33                      "mat_type": "aij", "pc_type": "lu",
34                      "pc_factor_mat_solver_type": "mumps"}

35    def solve(self):
36        super().solve()
37        self.failed_to_solve = False
38        u_old = self.solution.copy(deepcopy=True)
39        try:
40            solve(self.F == 0, self.solution,
41                  bcs=self.bcs,
42                  solver_parameters=self.params)
43        except ConvergenceError:
44            self.failed_to_solve = True
45            self.solution.assign(u_old)
    
```

Listing 1 Incompressible Navier-Stokes solver based on Taylor-Hood finite elements. In particular, the weak formulation described in (4) is implemented in lines 21–23

boundary condition (1d). The resulting weak formulation of equations (1) reads:

Find $\mathbf{u}$ such that $\mathbf{u} = \mathbf{g}$ on $\partial\Omega \setminus \Gamma$ and

$$\int_{\Omega} 2\nu \varepsilon \mathbf{u} : \varepsilon \mathbf{v} - p \operatorname{div}(\mathbf{v}) + \mathbf{v} \cdot (\nabla \mathbf{u}) \mathbf{u} + q \operatorname{div}(\mathbf{u}) \, dx = 0 \quad (4)$$

for any pair $(\mathbf{v}, q)$.

To implement the finite element discretization of (4), we create a class `NavierStokesSolver` that inherits from the `Fireshape`'s class `PdeConstraint`, see Listing 1. To discretize (4), we employ P2-P1 Taylor-Hood finite elements, that is, we discretize the trial and test velocity functions $\mathbf{u}$ and $\mathbf{v}$ with piecewise quadratic Lagrangian finite elements, and the trial and test pressure functions p and q with piecewise affine Lagrangian finite elements. It is well known that this is a stable discretization of the incompressible Navier-Stokes equations (Elman et al. 2014, pp. 136–137).

To address the nonlinearity in (4), we use PETSc's Scalable Nonlinear Equations Solvers (SNES) (Balay et al. 2019; Balay et al. 1997; Dalcin et al. 2011; Chevalier and Pellegrini 2008; Amestoy et al. 2000; Hype 2020). In each iteration, SNES linearizes (4) and solves the resulting system with a direct solver. In general, it is possible that the SNES solver fails to converge sufficiently quickly (in our code, we allow at most 100 SNES iterations). For instance, this happens if the finite element mesh self-intersects, or if the initial guess used to solve (4) is not sufficiently good. Most often, these situations happen when the optimization algorithm takes an optimization step that is too large. In these cases, we can address SNES' failure to solve (4) by reducing the optimization step. In practice, we deal with these situations with Python's `try: ... except: ... block`. If the SNES solver fails with a `ConvergenceError`, we catch this error and set the boolean flag `self.failed_to_solve` to `True`. This flag is used to adjust the output of the objective function `J`, see Section 2.3.

2.3 Step 3: Implement the objective function

The third step is to implement a code to evaluate the objective function `J` defined in (2). For this, we create a class `EnergyDissipation` that inherits from `Fireshape`'s class `ShapeObjective` (see Listing 2). Of course, evaluating `J` requires access to the fluid velocity $\mathbf{u}$. This access is implemented by assigning the variable `self.pde_solver`. This variable gives access also to the variable `NavierStokesSolver.failed_to_solve`, which can be used to control the output of `J` when the Navier-Stokes' solver fails to converge. Here, we decide that the value of `J` is `NaN` ("not a number") if the Navier-Stokes' solver fails to converge.

2.4 Final step: Set up and solve the problem

At this stage, we have all the necessary ingredients to tackle the shape optimization problem (3). The final step is to create a "main file" that loads the initial mesh, sets up the optimization problem, and solves it with an optimization

```

1 from firedrake import *
2 from fireshape import ShapeObjective
3 from PDEconstraint import NavierStokesSolver
4 import numpy as np

5 class EnergyDissipation(ShapeObjective):
6     """Kinetic energy dissipation."""

7     def __init__(self, pde_solver: NavierStokesSolver,
8                 *args, **kwargs):
9         super().__init__(*args, **kwargs)
10        self.pde_solver = pde_solver

11    def value_form(self):
12        """Evaluate energy functional."""
13        nu = self.pde_solver.viscosity

14        if self.pde_solver.failed_to_solve:
15            return np.nan * dx(self.pde_solver.mesh_m)
16        else:
17            z = self.pde_solver.solution
18            u, p = split(z)
19            return nu * inner(sym(grad(u)), sym(grad(u)))

```

Listing 2 Objective function `J` (2), which quantifies the kinetic energy dissipation of the fluid. Note that `J` returns `NaN` ("not a number") if the solver of the PDE constraint (4) fails to converge

algorithm. The rest of this section contains a line-by-line description of the "main file", which is listed in Listing 3.

- In lines 1–6, we import all necessary Python libraries and modules.
- In lines 8–13, we load the mesh of the initial design and we specify to discretize perturbations of the domain Ω using quadratic B-splines with support in a box that does not include the inlet and the outlet. In this example, we also impose some boundary regularity on B-splines and limit perturbations to be only in direction of the y -axis. Sections 5.1 and 5.2 provide more details about the `Fireshape` classes `ControlSpace` and `InnerProduct`.
- In lines 15–16, we initiate the Navier-Stokes finite element solver.
- In lines 18–19, we tell `Fireshape` to store the finite element solution to Navier-Stokes' equations in the folder `solution` every time the domain Ω is updated. We can visualize the evolution of $\mathbf{u}$ along the optimization process by opening the file `u.pvd` with `Paraview` (Ahrens et al. 2005).
- In lines 21–22, we initiate the objective function `J` and the associated reduced functional, which is used by `Fireshape` to define the appropriate Lagrange functionals to shape differentiate `J`. We refer to Section 4.2 for more details about shape differentiation using Lagrangians. We stress that `Fireshape` does not require users to derive (and implement) shape derivative formulas by hand. The whole shape differentiation

```

1 from firedrake import *
2 from fireshape import *
3 import fireshape.zoo as fsz
4 import ROL
5 from PDEconstraint import NavierStokesSolver
6 from objective import EnergyDissipation

7 # setup problem
8 mesh = Mesh("pipe.msh")
9 bbox = [(-0.5, +0.5), (-0.5, +5.5), (2, 12)]
10 orders = [3, 3, 3]; levels = [2, 2, 4]
11 Q = BsplinesControlSpace(mesh, bbox, orders, levels,
12     boundary_regularity=[0, 0, 2], fixed_dims=[0, 2])
13 q = ControlVector(Q, LaplaceInnerProduct(Q))

14 # setup PDE constraint
15 viscosity = Constant(1/10.)
16 e = NavierStokesSolver(Q.mesh_m, viscosity)

17 # save state variable evolution in file u.pvd
18 out = File("solution/u.pvd")
19 def cb(): return out.write(e.solution.split()[0])

20 # create PDE-constrained objective functional
21 J_ = EnergyDissipation(e, Q, cb=cb)
22 J = ReducedObjective(J_, e)

23 # add regularization to improve mesh quality
24 Jq = fsz.MoYoSpectralConstraint(10, Constant(0.5), Q)
25 J = J + Jq

26 # Set up volume constraint
27 vol = fsz.VolumeFunctional(Q)
28 vol0 = vol.value(q, None)
29 C = EqualityConstraint([vol], target_value=[vol0])
30 M = ROL.StdVector(1)

31 # ROL parameters
32 params_dict = {
33     'General': {'Secant': {'Type': 'Limited-Memory BFGS',
34         'Maximum Storage': 20}},
35     'Step': {'Type': 'Augmented Lagrangian',
36         'Augmented Lagrangian':
37             {'Subproblem Step Type': 'Trust Region',
38                 'Subproblem Iteration Limit': 50,
39                 'Penalty Parameter Growth Factor': 2}},
40     'Status Test': {'Gradient Tolerance': 1e-2,
41         'Step Tolerance': 1e-2,
42         'Constraint Tolerance': 1e-2,
43         'Iteration Limit': 10}}
44 params = ROL.ParameterList(params_dict, "Parameters")
45 problem = ROL.OptimizationProblem(J, q, econ=C, emul=M)
46 solver = ROL.OptimizationSolver(problem, params)
47 solver.solve()
    
```

Listing 3 Minimize the kinetic energy dissipation of a fluid using Fireshape

process is automated using pyadjoint (Ham et al. 2019; Dokken et al. 2020), see Remark 1.

- In lines 24–25, we add an additional regularization term to J to promote mesh quality of domain updates. This regularization term controls the pointwise maximum singular value of the geometric transformation used to update Ω (Wechsung 2019). See Section 4.2, Fig. 9, and Remark 3 for more information about the role of geometric transformations in Fireshape.

- In lines 27–30, we set up an equality constraint to ensure that the volume of the initial and the optimized domains are equal.
- In lines 32–43, we select an optimization algorithm from the optimization library ROL. More specifically, we use an augmented Lagrangian approach (Nocedal and Wright 2006, Ch. 17.3) with limited-memory BFGS Hessian updates (Nocedal and Wright 2006, Ch. 6.1) to deal with the volume equality constraint, and a Trust-Region algorithm (Nocedal and Wright 2006, Ch. 4.1) to solve the intermediate models generated by the augmented Lagrangian algorithm.
- In lines 44–47, finally, we gather all information and solve the problem.

Remark 1 A reader may wonder why Fireshape does not require information about shape derivative formulas and adjoint equations to solve a PDE-constrained shape optimization problem. This is possible because Fireshape employs UFL’s and pyadjoint’s automated shape differentiation and adjoint derivation capabilities (Ham et al. 2019; Dokken et al. 2020). In particular, UFL combines the automated construction of finite element pullbacks with symbolic differentiation to automate the evaluation of directional shape derivatives along vector fields discretized by finite elements (Ham et al. 2019). This process mimics the analytical differentiation of shape functions and is equivalent to the “optimize-then-discretize” paradigm, i.e. it yields exact (up to floating point accuracy) derivatives.

2.5 Results

Running the code contained in Listing 3 optimizes the pipe design when the fluid viscosity ν is 0.1, which corresponds to Reynolds number $\Re = 1/\nu = 10$. Of course, the resulting shape depends on the fluid viscosity. A natural question is how the optimized shape depends on this parameter. To answer this question, we can simply run Listing 3 for different Reynolds numbers (by modifying line 13 with the desired value) and inspect the results. Here we perform this comparison for $\Re \in \{1, 250, 500, 750, 1000\}$.²

In Fig. 2, we show the initial design and plot the magnitude of the fluid velocity on a cross section of the pipe for different Reynolds numbers. In Fig. 3, we show the resulting optimized shapes, and in Fig. 4 the corresponding magnitudes of the fluid velocity. Qualitatively, we observe that, as the Reynolds number increase, we obtain increasingly S-shaped designs that avoid high curvature at the two

²For this computation we extend the code in Listing 1 and include analytic continuation in the Reynolds number to re-compute good initial guesses for the SNES solver when this fails to converge. This modification is necessary to solve (4) at high Reynolds numbers. The code used to obtain these results can be found on Zenodo (2020).

Fig. 2 Initial design of the pipe and velocity solutions for $\Re = 1, 250, 500, 750, 1000$ (from top to bottom) inside the pipe. Red corresponds to high velocity, blue corresponds to low velocity

fixed ends. Finally, we remark that the objective is reduced by approximately 6.9%, 9.5%, 10.8%, 11.7% and 12.1% in at most 10 augmented Lagrangian steps, respectively. Each numerical experiment was run in parallel on a server with 36 cores.

3 Example: shape optimization of a cantilever

In this second tutorial, we consider another shape optimization classic (Allaire and Pantz 2006; Allaire 2007):

Fig. 3 Optimized shapes for $\Re = 1, 250, 500, 750, 1000$ (from top to bottom)

Fig. 4 Velocity field for optimized shapes for $\Re = 1, 250, 500, 750, 1000$ (from top to bottom). Red corresponds to high velocity, blue corresponds to low velocity

the minimization of the compliance of a cantilever Ω (see Fig. 5) subject to a dead surface load $\mathbf{g}$. To make this example reproducible on a basic laptop we consider a two dimensional version of this problem. Adapting the code presented in this section to simulate a three-dimensional cantilever is straightforward.

We model the cantilever elastic response using the linear elasticity equations (Braess 2007, Ch. VI, Sect. 3). More precisely, we consider the displacement formulation: the displacement field $\mathbf{u}$ satisfies the variational problem

$$\int_{\Omega} \sigma(\mathbf{u}) : \epsilon(\mathbf{v}) \, d\mathbf{x} = \int_{\Gamma_2} \mathbf{g} \cdot \mathbf{v} \, dS \quad \text{for every } \mathbf{v}, \quad (5)$$

where $\sigma(\mathbf{u}) = \lambda \operatorname{div}(\mathbf{u})\mathbf{I} + 2\mu\epsilon(\mathbf{u})$, λ and μ are the Lamé constants, $\mathbf{I}$ is the identity matrix, $\epsilon(\mathbf{u}) = \frac{1}{2}(\nabla\mathbf{u} + \nabla\mathbf{u}^T)$, $\nabla\mathbf{u}$ is the derivative (Jacobian matrix) of $\mathbf{u}$ and $\nabla\mathbf{u}^T$ is the derivative transposed, the colon symbol $:$ denotes the Frobenius inner product, that is $\sigma(\mathbf{u}) : \epsilon(\mathbf{v}) = \operatorname{trace}(\sigma(\mathbf{u})^T \epsilon(\mathbf{v}))$, and Γ_2 denotes the part of the boundary that is subject to the surface load $\mathbf{g}$. Additionally, although

Fig. 5 A cantilever is attached to a wall on the left (Γ_1) and pulled downwards on the right (Γ_2) by a dead surface load. A poor cantilever design may fail to offer sufficient support

not explicitly indicated in (5), the displacement $\mathbf{u}$ vanishes on Γ_1 , which is the part of the boundary $\partial\Omega$ that corresponds to the wall (see Fig. 5).

To model the compliance of the structure Ω , we consider the functional

$$J(\Omega) = \int_{\Omega} \sigma(\mathbf{u}) : \varepsilon(\mathbf{u}) \, dx, \quad (6)$$

where $\mathbf{u}$ is the displacement and satisfies (5).

As typical for this shape optimization test case, we also impose a volume equality constraint on Ω to prescribe the amount of material used to fabricate the cantilever.

Similarly to Section 2, in the following subsections we explain how to solve this shape optimization test case in Fireshape.

3.1 Step 1: Provide an initial guess

To create the mesh of the initial guess, we use again the software Gmsh (Geuzaine and Remacle 2009). The initial guess employed is sketched in Fig. 5. The only mesh detail necessary to understand this tutorial is that, in Listing 4, the number 1 corresponds to Γ_1 , whereas the boundary flag 2 refers to Γ_2 . For the remaining geometric details, we refer to the code archived on Zenodo (2020).

3.2 Step 2: Implement the PDE constraint

In this example, the PDE constraint (5) is already in weak form. To implement its finite element discretization, we create a class `LinearElasticitySolver` that inherits from Fireshape's class `PdeConstraint` (see Listing 4). To discretize (5), we employ the standard piecewise affine Lagrangian finite elements (Braess 2007, Ch. II, Sect. 5). To solve the resulting linear system, we employ a direct solver. This is sufficiently fast for the current 2D example. For readers interested in creating a three-dimensional version of this example, we suggest replacing the direct solver with a Krylov method preconditioned by one of the algebraic multigrid methods that can be accessed via PETSc.

3.3 Step 3: Implement the objective function

We implement the objective function J defined in (6) in a class `Compliance` that inherits from Fireshape's class `ShapeObjective` (see Listing 5). This functional accesses the displacement field $\mathbf{u}$ via the variable `self.pde_solver`.

To verify that the domain used to compute $\mathbf{u}$ is feasible, that is, that the mesh is not tangled, we create a piecewise constant function `detDT` and interpolate the determinant of the derivative of the transformation $\mathbf{T}$ used to generate the

```

1 from firedrake import *
2 from fireshape import PdeConstraint
3
4 class LinearElasticitySolver(PdeConstraint):
5     """Linear elasticity equation."""
6
7     def __init__(self, mesh_m, viscosity):
8         super().__init__()
9         self.mesh_m = mesh_m
10
11         # Setup problem, linear finite elements
12         self.V = VectorFunctionSpace(mesh_m, "CG", 1)
13
14         # Preallocate soln variables for state eqn
15         self.solution = Function(self.V, name="State")
16         self.testfunction = TestFunction(self.V)
17
18         # Define surface dead load and Lamé parameters
19         self.g = as_vector([0, -1])
20         E = 15
21         nu = 0.35
22         self.mu = E/(2*(1+nu))
23         self.lbd = E*nu/((1+nu)*(1-2*nu))
24
25         # Weak form of linear elasticity
26         u = self.solution
27         v = self.testfunction
28         self.F = inner(self.sigma(u), self.epsilon(v))*dx
29             - dot(self.g, v)*ds(2)
30
31         # Dirichlet boundary conditions
32         u_fix = as_vector([0, 0])
33         self.bcs = DirichletBC(self.V, u_fix, [1])
34
35         # PDE-solver parameters
36         self.params={"ksp_type": "preonly",
37                     "pc_type": "lu"}
38
39     def epsilon(self, w):
40         return 0.5*(grad(w) + grad(w).T)
41
42     def sigma(self, w):
43         Id = Identity(self.mesh_m.geometric_dimension())
44         return (return self.lbd*div(w)*Id
45                 + 2*self.mu*self.epsilon(w))
46
47     def solve(self):
48         super().solve()
49         solve(self.F == 0, self.solution, bcs=self.bcs,
50               solver_parameters=self.params)

```

Listing 4 Linear elasticity solver

updated domains (see Sections 4.2 and 4.3 for more details about the role of the transformation $\mathbf{T}$ in Fireshape). If the minimum of `detDT` is a small number, this indicates that the mesh has poor quality (see also Remark 3), and this negatively affects the accuracy of the finite element method. In this case, we decide that the value of J is NaN ("not a number") to ensure that only trustworthy simulations are used.³

³In Section 2.3, we did not include this extra test because the Navier-Stokes' solver usually fails to converge on poor quality meshes.

```

1 from firedrake import *
2 from fireshape import ShapeObjective
3 from PDEconstraint import LinearElasticitySolver
4 import numpy as np

5 class Compliance(ShapeObjective):
6     """Compliance functional for linear Elasticity."""
7
8     def __init__(self,
9                 pde_solver: LinearElasticitySolver,
10                 Q, cb, *args, **kwargs):
11         super().__init__(Q, cb, *args, **kwargs)
12         self.pde_solver = pde_solver
13         Vdet = FunctionSpace(Q.mesh_r, "DG", 0)
14         self.detDT = Function(Vdet)
15
16     def value_form(self):
17         """Evaluate compliance functional."""
18
19         self.detDT.interpolate(det(grad(self.Q.T)))
20         if min(self.detDT.vector()) > 0.05:
21             u = self.pde_solver.solution
22             sigma_u = self.pde_solver.sigma(u)
23             eps_u = self.pde_solver.epsilon(u)
24             return inner(sigma_u, eps_u)*dx
25         else:
26             return np.nan * dx(self.pde_solver.mesh_m)

```

Listing 5 Objective function J (6), which quantifies the compliance of the structure. Note that J returns NaN (“not a number”) if the mesh has poor quality

3.4 Final step: Set up and solve the problem

Finally, to shape optimize the initial desing we create a “main file” that loads the initial mesh, sets up the optimization problem, and solves it with an optimization algorithm. The rest of this section contains a line-by-line description of the “main file”, which is listed in Listing 6.

- In lines 1–6, we import all necessary Python libraries and modules.
- In lines 8–11, we load the mesh of the initial design and we specify to discretize perturbations of the domain Ω using piecewise affine Lagrangian finite elements. We also employ a metric inspired by the linear elasticity equations and specify that the boundaries Γ_1 and Γ_2 cannot be modified. Sections 5.1 and 5.2 provide more details about the Fireshape classes `ControlSpace` and `InnerProduct`.
- In line 13, we initiate the linear elasticity solver.
- In lines 15–16, we tell Fireshape to store the finite element approximation of the displacement field $\mathbf{u}$ in the folder `solution` every time the domain Ω is updated. We can visualize the evolution of $\mathbf{u}$ along the optimization process by opening the file `u.pvd` with Paraview (Ahrens et al. 2005).
- In lines 18–19, we initiate the objective function J and the associated reduced functional, which is used by

```

1 from firedrake import *
2 from fireshape import *
3 import fireshape.zoo as fsz
4 import ROL
5 from PDEconstraint import LinearElasticitySolver
6 from objective import Compliance

7 # setup problem
8 mesh = Mesh("cantilever.msh")
9 Q = FeControlSpace(mesh)
10 inner = ElasticityInnerProduct(Q, fixed_bids=[1, 2])
11 q = ControlVector(Q, inner)

12 # setup PDE constraint
13 e = LinearElasticitySolver(Q.mesh_m)

14 # save state variable evolution in file u.pvd
15 out = File("solution/u.pvd")
16 def cb(): return out.write(e.solution)

17 # create PDE-constrained objective functional
18 J_ = Compliance(e, Q, cb=cb)
19 J = ReducedObjective(J_, e)

20 # add regularization to improve mesh quality
21 Jq = fsz.MoYoSpectralConstraint(10, Constant(0.5), Q)
22 J = J + Jq

23 # Set up volume constraint
24 vol = fsz.VolumeFunctional(Q)
25 vol0 = vol.value(q, None)
26 C = EqualityConstraint([vol], target_value=[vol0])
27 M = ROL.StdVector(1)

28 # ROL parameters
29 params_dict = {
30     'General': {'Secant': {'Type': 'Limited-Memory BFGS'}},
31     'Step': {'Type': 'Augmented Lagrangian',
32             'Augmented Lagrangian':
33                 {'Subproblem Step Type': 'Trust Region',
34                  'Subproblem Iteration Limit': 10,
35                  'Penalty Parameter Growth Factor': 2}},
36     'Status Test': {'Gradient Tolerance': 1e-2,
37                    'Step Tolerance': 1e-3,
38                    'Constraint Tolerance': 1e-2,
39                    'Iteration Limit': 15}}
40 params = ROL.ParameterList(params_dict, "Parameters")
41 problem = ROL.OptimizationProblem(J, q, econ=C, emul=M)
42 solver = ROL.OptimizationSolver(problem, params)
43 solver.solve()

```

Listing 6 Minimize the kinetic energy dissipation of a fluid using Fireshape

Fireshape to define the appropriate Lagrange functionals to shape differentiate J . We refer to Section 4.2 for more details about shape differentiation using Lagrangians. Note that, as mentioned in Remark 1, Fireshape does not require users to derive (and implement) shape derivative formulas by hand. The whole shape differentiation process is automated using pyadjoint (Ham et al. 2019; Dokken et al. 2020).

- In lines 21–22, we add an additional regularization term to J to promote mesh quality of domain updates. This regularization term controls the pointwise maximum singular value of the geometric transformation used to

update Ω (Wechsung 2019). See Section 4.2, Fig. 9, and Remark 3 for more information about the role of geometric transformations in Fireshape.

- In lines 24–27, we set up an equality constraint to ensure that the volume of the initial and the optimized domains are equal.
- In lines 29–39, we select an optimization algorithm from the optimization library ROL. More specifically, we use an augmented Lagrangian approach (Nocedal and Wright 2006, Ch. 17.3) with limited-memory BFGS Hessian updates (Nocedal and Wright 2006, Ch. 6.1) to deal with the volume equality constraint, and a Trust-Region algorithm (Nocedal and Wright 2006, Ch. 4.1) to solve the intermediate models generated by the augmented Lagrangian algorithm.
- In lines 44–47, finally, we gather all information and solve the problem.

3.5 Results

Figure 6 shows the initial design (top) and the design optimized using Listing 6 (middle). At first look, the result looks reasonable. However, after an inspection of ROL's output, it becomes apparent that the optimization algorithm stops after 8 (augmented Lagrangian outer) steps because the optimization step length becomes too small. This is due to a triangle on the top left corner of the cantilever that becomes almost flat (see Fig. 7, left). This deterioration of the mesh quality is detected by the implementation of the objective function (Listing 5, line 21), which forces the trust-region optimization algorithm to reduce the trust-region radius. After 8 iterations, this leads to an optimization step length that is smaller than 10^{-4} , and the optimization algorithm terminates. This is a positive result: the algorithm detects that something is wrong and stops the simulation. It is the task of the user to understand what originates the problem and provide a remedy. After a look at the iterates on Paraview, it becomes apparent that optimization algorithms tries to move a node above the boundary Γ_1 , which is fixed, and this leads to the generation of an elongated triangle. A difficulty with this part of the design was already experienced in Allaire and Pantz (2006), where the authors suggest to “arbitrarily set the shape gradient to zero near the corners of the shape”. Here, we present an alternative solution.

As mentioned in the previous paragraph, the issue is that the shape updates lead to the creation of an elongated triangle. To overcome this problem, we can simply increase from 10 to 100 the coefficient of the regularization term in line 21 Listing 6. The increased regularization sufficiently penalizes this unwanted behavior, and the algorithm can further optimize the design (see Fig. 8).

Fig. 6 *Top:* Initial cantilever design. The absolute value of the displacement field $\mathbf{u}$ is colored from blue (minimum) to red (maximum). *Middle:* Cantilever designed optimized using Listing 6. *Bottom:* Cantilever designed optimized by increasing the regularization of J in line 21 of Listing 6

4 Shape optimization via diffeomorphisms

In this section, we describe the theory that underpins Fireshape. We begin with a brief introduction to PDE-constrained optimization to set the notation and mention the main idea behind optimization algorithms (Section 4.1). Then, we continue with an introduction to shape calculus (Section 4.2). Finally, we conclude with a discussion about the link between shape optimization and parametric finite elements (Section 4.3).

4.1 Optimization with PDE constraints

The basic ingredients to formulate a PDE-constrained optimization problem are the following: a control variable $\mathbf{q}$

Fig. 7 *Left:* Mesh detail of upper left corner of the cantilever design optimized with Listing 6. The leftmost triangle has almost collapsed to a line. *Right:* Same mesh detail when the regularization of J in line 21 of Listing 6 is increased. The elongated triangle is no longer present

that lives in a control space $\mathbf{Q}$, a state variable $\mathbf{u}$ that lives in a state space $\mathbf{U}$ and that solves a (possibly nonlinear) state equation $\mathbf{A}(\mathbf{q}, \mathbf{u}) = \mathbf{0}$, and a real function $J : \mathbf{Q} \times \mathbf{U} \rightarrow \mathbb{R}$ to be minimized.

Example 1 Equation (3) is a PDE-constrained optimization problem. The control variable $\mathbf{q}$ corresponds to the domain Ω , the state variable $\mathbf{u}$ represents the pair velocity-pressure $(\mathbf{u}, p)$, the nonlinear constraint $\mathbf{A}$ represents the Navier-Stokes equations (1), and J denotes the objective function (2). The state space $\mathbf{U}$ corresponds to the space of pairs $(\mathbf{a}, b)$ with weakly differentiable velocities $\mathbf{a}$ that satisfy $\mathbf{a} = \mathbf{g}$ on $\partial\Omega \setminus \Gamma$ and square integrable pressures b (Elman et al. 2014, Ch. 8.2). The control space $\mathbf{Q}$ is specified in Section 4.2 (see (7)).

Fig. 8 Convergence history of the compliance. When the regularization of J in line 21 of Listing 6 is insufficient, the algorithm stops after eight iterations because a triangle becomes very elongated. However, if sufficient regularization is provided, the algorithm does not experience this issue and further optimizes the design. Finally, we point out that the small value of the compliance after the first iteration is obtained by excessively violating the volume equality constraint. This is subsequently corrected by the augmented Lagrangian algorithm

Most⁴ numerical methods for PDE-constrained optimization attempt to construct a sequence of controls $\{\mathbf{q}^{(k)}\}_{k \in \mathbb{N}}$ and corresponding states $\{\mathbf{u}^{(k)}\}_{k \in \mathbb{N}}$ such that

$$\lim_{k \rightarrow \infty} J(\mathbf{q}^{(k)}, \mathbf{u}^{(k)}) = \inf_{\mathbf{q}, \mathbf{u}} \{J(\mathbf{q}, \mathbf{u}) : \mathbf{A}(\mathbf{q}, \mathbf{u}) = \mathbf{0}\}.$$

Often, the sequence $\{\mathbf{q}^{(k)}\}_{k \in \mathbb{N}}$ is constructed using derivatives of the function J and of the constraint $\mathbf{A}$. Common approaches are steepest descent algorithms and Newton methods (in their quasi, Krylov, or semi-smooth versions) (Hinze et al. 2009; Kelley 2003). These algorithms ensure that the sequence $\{J(\mathbf{q}^{(k)}, \mathbf{u}^{(k)})\}_{k \in \mathbb{N}}$ decreases. In special cases, it is even possible to show that the sequence $\{\mathbf{q}^{(k)}\}_{k \in \mathbb{N}}$ converges (Hinze et al. 2009). Although it may be difficult to ensure these assumptions are met in industrial applications, these optimization algorithms are still powerful tools to improve state-of-the-art designs and are widely used to perform shape optimization.

4.2 Shape optimization and shape calculus

Shape optimization with PDE constraints is a particular branch of PDE-constrained optimization where the control space $\mathbf{Q}$ is a set of domain shapes. The space of shapes is notoriously difficult to characterize uniquely. For instance, one could describe shapes through their boundary regularity, or as level-sets, or as local epigraphs (Delfour and Zolésio 2011, Ch. 2). The choice of the shapes' space characterization plays an important role in the concrete implementation of a shape optimization algorithm, and it can also affect formulas that result by differentiating J with respect to perturbations of the shape $\mathbf{q}$. To be more precise,

⁴An alternative approach is to employ so-called “one-shot” methods, where the optimality system of the problem is solved directly (Schulz and Gherman 2009). Note that, since the optimality system is often nonlinear, one-shot methods still involve iterative algorithms.

different methods generally lead to the same first order shape derivative formula, but differ on shape derivatives of higher order (Delfour and Zolésio 2011, ch. 9)

In this work, we model the control space $\mathbf{Q}$ as the images of bi-Lipschitz geometric transformations $\mathbf{T}$ applied to an initial set $\Omega \subset \mathbb{R}^d$ (Delfour and Zolésio 2011, ch. 3), that is,

$$\mathbf{Q} := \{\mathbf{q} = \mathbf{T}(\Omega) : \mathbf{T} : \mathbb{R}^d \rightarrow \mathbb{R}^d \text{ is bi-Lipschitz}\}, \quad (7)$$

see Fig. 9. We choose this model of $\mathbf{Q}$ because it provides an explicit description of the domain boundaries via $\partial(\mathbf{T}(\Omega)) = \mathbf{T}(\partial\Omega)$, and because it is compatible with higher order finite elements (Paganini et al. 2018), as explained in detail in the Section 4.3. Henceforth, we use the term diffeomorphism to indicate that a geometric transformation $\mathbf{T}$ is bi-Lipschitz.

In this setting, the shape derivative of J corresponds to the classical Gâteaux derivative in the Sobolev space $W^{1,\infty}(\mathbb{R}^d, \mathbb{R}^d)$ (Hinze et al. 2009, Def. 1.29). To see this, let us momentarily remove the PDE constraint, and only consider a shape functional of the form $\Omega \mapsto J(\Omega)$. The shape derivative of J at Ω is the linear and continuous operator $dJ(\Omega, \cdot) : W^{1,\infty}(\mathbb{R}^d, \mathbb{R}^d) \rightarrow \mathbb{R}$ defined by

$$dJ(\Omega, \mathbf{V}) := \lim_{t \searrow 0} \frac{J(\Omega_t) - J(\Omega)}{t}, \quad (8)$$

where Ω_t is the set $\Omega_t = \{\mathbf{x} + t\mathbf{V}(\mathbf{x}) : \mathbf{x} \in \Omega\}$ (Allaire 2007, Def. 6.15). By replacing Ω with $\mathbf{q} = \mathbf{I}(\Omega)$, where $\mathbf{I}$ denotes the identity transformation defined by $\mathbf{I}(\mathbf{x}) = \mathbf{x}$ for any $\mathbf{x} \in \mathbb{R}^d$, (8) can be equivalently rewritten as

$$dJ(\mathbf{q}, \mathbf{V}) := \lim_{t \searrow 0} \frac{J(\mathbf{q} + t\mathbf{V}) - J(\mathbf{q})}{t}.$$

We highlight that this interpretation immediately generalizes to any $\mathbf{q}$ in $\mathbf{Q}$ and can be used to define higher order shape derivatives.

The same definition (8) of shape derivative holds in the presence of PDE constraints: the shape derivative at $\mathbf{q}$ of the

function $J : \mathbf{Q} \times \mathbf{U} \rightarrow \mathbb{R}$ subject to $\mathbf{A}(\mathbf{q}, \mathbf{u}) = \mathbf{0}$ is the linear and continuous operator defined by

$$dJ(\mathbf{q}, \mathbf{u}, \mathbf{V}) := \lim_{t \searrow 0} \frac{J(\mathbf{q} + t\mathbf{V}, \mathbf{u}_t) - J(\mathbf{q}, \mathbf{u})}{t}, \quad (9)$$

where $\mathbf{u}_t$ is the solution to $\mathbf{A}(\mathbf{q} + t\mathbf{V}, \mathbf{u}_t) = \mathbf{0}$. Computing shape derivative formulas using (9) may present the difficulty of computing the shape derivative of $\mathbf{u}$, which intuitively arises by the “chain rule” formula. The shape derivative of $\mathbf{u}$, which is sometimes called “material derivative” of $\mathbf{u}$, can be eliminated with the help of adjoint equations (Hinze et al. 2009, Section 1.6.2). This process can be automated by introducing the Lagrange functional (Hinze et al. 2009, Section 1.6.3)

$$\mathcal{L}(\mathbf{q}, \mathbf{u}, \mathbf{p}) := J(\mathbf{q}, \mathbf{u}) + \langle \mathbf{A}(\mathbf{q}, \mathbf{u}), \mathbf{p} \rangle. \quad (10)$$

The term $\langle \mathbf{A}(\mathbf{q}, \mathbf{u}), \mathbf{p} \rangle$ stems from testing the equation $\mathbf{A}(\mathbf{q}, \mathbf{u}) = \mathbf{0}$ with a test function $\mathbf{p}$ in the same way it is usually done when writing a PDE in its weak form.

Example 2 If $\mathbf{A}$ denotes the Navier-Stokes (1), then $\langle \mathbf{A}(\mathbf{q}, \mathbf{u}), \mathbf{p} \rangle$ corresponds to the weak formulation (4).

The advantage of introducing the Lagrangian (10) is that, by choosing $\mathbf{p}$ as the solution to the adjoint equation

$$\langle \partial_{\mathbf{u}} \mathbf{A}(\mathbf{q}, \mathbf{u}), \mathbf{p} \rangle = -\partial_{\mathbf{u}} J(\mathbf{q}, \mathbf{u}) \quad \text{for all } \mathbf{u} \text{ in } \mathbf{U},$$

where $\partial_{\mathbf{u}}$ denotes the partial differentiation with respect to the variable $\mathbf{u}$, the shape derivative of J can be computed as

$$dJ(\mathbf{q}, \mathbf{u}, \mathbf{V}) = \partial_{\mathbf{q}} \mathcal{L}(\mathbf{q}, \mathbf{u}, \mathbf{p}, \mathbf{V}),$$

where $\partial_{\mathbf{q}}$ denotes partial differentiation with respect to the variable $\mathbf{q}$. This is advantageous because partial differentiation does not require computing the shape derivative of $\mathbf{u}$.

The Lagrangian approach to compute derivatives of PDE-constrained functionals is well established (Ito and Kunisch 2008), and its steps can be replicated by symbolic computation software. Probably, the biggest success in this direction is the pyadjoint project (Farrell et al. 2013; Mitusch et al. 2019), which derives “the adjoint and tangent-linear equations and solves them using the existing solver methods in FEniCS/Firedrake” (Mitusch et al. 2019). Thanks to the shape differentiation capabilities of UFL introduced in Ham et al. (2019), pyadjoint is also capable of shape differentiating PDE-constrained functionals along directions discretized by finite elements (Dokken et al. 2020).

Remark 2 Optimization algorithms are usually based on steepest descent directions to update the control variable $\mathbf{q}$. A steepest descent direction is a direction $\mathbf{V}^*$ that minimizes the derivative dJ . Since dJ is linear, it is necessary to introduce a norm $\|\cdot\|$ on the space of possible perturbations $\{\mathbf{V}\}$ (the tangent space of $\mathbf{Q}$ at $\mathbf{q}$), and to restrict the search

Fig. 9 To construct the control space $\mathbf{Q}$, we select an initial set Ω and a set of diffeomorphisms $\{\mathbf{T}\}$ (black arrows), and collect every image $\mathbf{T}(\Omega)$

Fig. 10 Simplest approach to PDE-constrained shape optimization. An initial mesh (*left*) is used to compute state and adjoint variables. This information, together with shape derivatives, is used to devise an update of the nodes' coordinates (*center*). A new and improved

initial guess is obtained by updating the nodes' coordinates and retaining the initial mesh connectivity (*right*). This process is repeated until convergence

of a steepest descent direction to directions $\mathbf{V}$ of length $\|\mathbf{V}\| = 1$. A natural choice would be to select the $W^{1,\infty}$ -norm,⁵ with respect to which a minimizer has been shown to exist for most functionals (Paganini et al. 2018, Prop. 3.1). However, in practice it is more convenient to employ a norm that is induced by an inner product, so that the steepest descent direction corresponds to the Riesz representative of dJ with respect to the inner product (Hinze et al. 2009, p. 98).

4.3 Geometric transformations, moving meshes, and parametric finite elements

To solve a PDE-constrained optimization problem iteratively, it is necessary to employ a numerical method that is capable of solving the constraint $\mathbf{A}(\mathbf{q}, \mathbf{u}) = \mathbf{0}$ for any feasible control $\mathbf{q}$. In shape optimization, this translates into the requirement of a numerical scheme that can solve a PDE on a domain that changes at each iteration of the optimization algorithm. There are several competing approaches to construct a numerical scheme with this feature, such as the level-set method (Allaire et al. 2002) or the phase field approach (Blank et al. 2014), among others.

In Fireshape, we employ the approach sometimes known as “moving mesh” method (Paganini et al. 2018; Allaire 2007). In its simplest version (see Fig. 10), this method replaces (or approximates) the initial domain Ω with a polygonal mesh Ω_h . On this mesh, the state and adjoint equations are solved with a standard finite element method, whose construction on polygonal meshes is immediate. For instance, depending on the nature of the state constraint, one may solve the state and adjoint equations using linear Lagrangian or P2-P1 Taylor-Hood finite elements (see Section 2.2). With the state and adjoint solutions at hand, one employs shape derivatives (see Remark 2 and Section 5.2) to update the coordinates of mesh nodes while retaining the mesh connectivity of Ω_h . This leads to a

new mesh that represents an improved design. This update process is repeated until some prescribed convergence criteria are met.

The moving mesh method we just described is a simple and yet powerful method. However, in its current formulation, it requires polyhedral meshes, which limits the search space $\mathbf{Q}$ to polyhedra. In the remaining part of this section, we describe an equivalent interpretation of the moving mesh method that generalizes to curved domains. Additionally, this alternative interpretation allows approximating state and adjoint variables with arbitrarily high-order finite elements without suffering from reduction of convergence order due to poor approximation of domain boundaries (Ciarlet 2002, Ch. 4.4).

We begin by recalling the standard construction of parametric finite elements. For more details, we refer to (Ciarlet 2002, Sections 2.3 and 4.3). To construct a parametric finite element space, one begins by partitioning the domain Ω into simpler geometric elements $\{K_i\}$ (usually triangles or tetrahedra, as depicted in the first row of Fig. 11). Then, one introduces a reference element $\hat{K}$ and a collection of diffeomorphisms $\mathbf{F}_i$ that map the reference element $\hat{K}$ to the various K_i s, that is, $\mathbf{F}_i(\hat{K}) = K_i$ for every value of the index i (as depicted in the second row of Fig. 11). Finally, one considers a set of local reference basis functions $\{\hat{b}_j\}$ defined on the reference element $\hat{K}$, and defines local basis functions $\{b_j^i\}$ on each K_i via the pullback $b_j^i(\mathbf{x}) := \hat{b}_j(\mathbf{F}_i^{-1}(\mathbf{x}))$ for every $\mathbf{x}$ in K_i . These local basis functions are used to construct global basis functions that span the finite element space. An important property of this construction is that the diffeomorphisms $\mathbf{F}_i$ can be expressed in terms of local linear Lagrangian basis functions⁶ $\{\hat{\beta}_m\}$, that is, there are some coefficient vectors $\{\mu_m^i\}$ such that

$$\mathbf{F}_i(\mathbf{x}) = \sum_m \mu_m^i \hat{\beta}_m(\mathbf{x}) \quad \text{for every } \mathbf{x} \text{ in } \hat{K}. \quad (11)$$

⁵The norm $\|\mathbf{V}\|_{1,\infty}$ is the maximum between the essential supremum of $\mathbf{V}$ and of its derivative $D\mathbf{V}$.

⁶For linear Lagrangian finite elements, the two set of reference local basis functions $\{\hat{b}_j\}$ and $\{\hat{\beta}_m\}$ coincide.

Fig. 11 To construct finite element basis functions, one usually begins by triangulating a domain Ω (first row) and then introducing diffeomorphisms $\mathbf{F}_i$ that map the reference element $\hat{K}$ to the various K_i in the triangulation (second row)

Finite elements constructed following this procedure are usually called parametric, because they rely on the parametrization $\{\mathbf{F}_i\}$. Note that the most common finite elements families, such as Lagrangian, Raviart–Thomas, or Nédélec finite elements, are indeed parametric.

Keeping this knowledge about parametric finite elements in mind, we can revisit the mesh method (see Fig. 10). There, the main idea was to update only nodes' coordinates and keep the mesh connectivity unchanged, so that constructing finite elements on the new mesh is straightforward. In Paganini et al. (2018), it has been shown that the new finite element space can also be obtained by modifying the parametric construction of finite elements on the original domain Ω . In the next paragraphs, we give an extended explanation (with adapted notation) of the demonstration given in Paganini et al. (2018).

Let $\mathbf{T}$ denote the transformation employed to modify the mesh Ω_h on the left in Fig. 10 into the new and perturbed mesh $\mathbf{T}(\Omega_h)$ on the right in Fig. 10. Additionally, let $\{\mathbf{T}(K_i)\}$ denote the simple geometric elements that constitute the latter. Using the parametric approach, we can construct finite elements on the new mesh by introducing a collection of diffeomorphisms $\tilde{\mathbf{F}}_i$ that map the reference element $\hat{K}$ to the various $\mathbf{T}(K_i)$ s, that is, $\tilde{\mathbf{F}}_i(\hat{K}) = \mathbf{T}(K_i)$ for every value of the index i ; see Fig. 12.

The behavior of the transformation $\mathbf{T}$ in Fig. 10 is prescribed only the mesh nodes. Since its behavior on the interior of the mesh triangles can be chosen arbitrarily, we can decide that $\mathbf{T}$ is piecewise affine on each triangle. This

convenient choice implies that $\mathbf{T}$ can be written as a linear combination of piecewise affine Lagrangian finite elements defined on the first mesh, that is, $\mathbf{T}(\mathbf{x}) = \sum_{\ell} \mathbf{v}_{\ell} B_{\ell}(\mathbf{x})$ for every $\mathbf{x}$ in Ω , where $\{\mathbf{v}_{\ell}\}$ are some coefficient vectors and $\{B_{\ell}\}$ are global basis functions of the space of piecewise affine Lagrangian finite elements defined on the partition $\{K_i\}$. Since Lagrangian finite elements are constructed via pullbacks to the reference element, for every element K_i there are coefficient vectors $\{\mathbf{v}_m^i\}$ so that the restriction $\mathbf{T}|_{K_i}$ of $\mathbf{T}$ on K_i can be rewritten as

$$\mathbf{T}|_{K_i}(\mathbf{x}) = \sum_{\ell} \mathbf{v}_{\ell} B_{\ell}|_{K_i}(\mathbf{x}) = \sum_m \tilde{\mathbf{v}}_m^i \hat{\beta}_m(\mathbf{F}_i^{-1}(\mathbf{x})).$$

Therefore, the composition $\mathbf{T}|_{K_i} \circ \mathbf{F}_i$ is of the form

$$\mathbf{T}|_{K_i} \circ \mathbf{F}_i(\mathbf{x}) = \sum_m \tilde{\mathbf{v}}_m^i \hat{\beta}_m(\mathbf{x}) \quad \text{for every } \mathbf{x} \in \hat{K}, \quad (12)$$

that is, of the same form of (11). This implies that, to construct finite elements on the perturbed geometry $\mathbf{T}(\Omega_h)$, we only need to replace the original coefficients $\{\mu_m^i\}$ in (11) with the new coefficients $\{\tilde{\mathbf{v}}_m^i\}$ from (12).

This alternative and equivalent viewpoint on the moving mesh method generalizes naturally to higher order finite element approximations of control and state (and adjoint) variables. Indeed, one of the key steps to ensure that higher order finite elements achieve higher order convergence on curved domains is to employ sufficiently accurate polynomial interpolation of domain boundaries (Ciarlet 2002, Ch. 4.4). This boundary interpolation can be encoded in the

Fig. 12 The standard construction of finite elements on the perturbed mesh $\{\mathbf{T}(K_i)\}$ employs diffeomorphisms $\tilde{\mathbf{F}}_i$ that map the reference element $\hat{K}$ to the various perturbed elements $\mathbf{T}(K_i)$

diffeomorphisms $\{\mathbf{F}_i\}$ by using higher order Lagrangian local basis functions. Therefore, simply employing higher order Lagrangian finite element transformations $\mathbf{T}$ leads to a natural extension of moving mesh method to higher order finite elements.

This alternative and equivalent viewpoint generalizes further to allow the use of any arbitrary discretization of the transformation $\mathbf{T}$ (for instance, using B-splines (Höllig 2003), harmonic polynomials, or radial basis functions (Wendland 2005)). The only requirement to be fulfilled to ensure the desired order of convergence p is that the maps $\mathbf{T} \circ \mathbf{F}_i : \mathbf{x} \mapsto \mathbf{T}(\mathbf{F}_i(\mathbf{x}))$ satisfy the asymptotic algebraic estimates

$$\|\mathbf{D}^\alpha(\mathbf{T} \circ \mathbf{F}_i)\| = \mathcal{O}(h^\alpha) \quad \text{for } 0 \leq \alpha \leq p,$$

where $\mathbf{D}^\alpha(\mathbf{T} \circ \mathbf{F}_i)$ denotes the α^{th} derivative of $\mathbf{T} \circ \mathbf{F}_i$. Using a different discretization of the transformation $\mathbf{T}$ can give several advantages, like increasing the smoothness $\mathbf{T}$ (because finite elements are generally only Lipschitz continuous) or varying how shape updates are computed during the optimization process (Eigel and Sturm 2018).

Remark 3 An issue that can arise with the moving mesh method is that it can lead to poor quality (or even tangled) meshes. In terms of geometric transformations, a mesh with poor quality corresponds to a transformation $\mathbf{T}$ for which the value $\max_\alpha \|\mathbf{D}^\alpha \mathbf{T}\|$ is large (and a tangled mesh to a transformation $\mathbf{T}$ that is not a diffeomorphism). To a certain extent, it is possible to enforce moderate derivatives by employing suitable metrics to extract descent directions from shape derivatives (for instance, by using linear elasticity based inner products with a Cauchy-Riemann augmentation (Iglesias et al. 2018)) and/or by adding penalty terms to the functional J . For example, in Sections 2.4 and 3.4 we add a penalization term based on the spectral norm $\|\mathbf{D}(\mathbf{T} - \mathbf{I})\|_2$ of the Jacobian matrix $\mathbf{D}(\mathbf{T} - \mathbf{I})$. We point out that if Ω is ζ -quasiconvex⁷ and $\|\mathbf{D}(\mathbf{T} - \mathbf{I})\|_2 < \zeta$, then $\mathbf{T} : \Omega \rightarrow \mathbb{R}^d$ is bi-Lipschitz, and thus injective.

5 Anatomy of Fireshape

In this section, we give more details about Fireshape's implementation and features. Fireshape is organized in a few core Python classes (and associated subclasses) that implement the control space $\mathbf{Q}$, the metric to be employed by the optimization algorithm, and the (possibly PDE-constrained) objective function J . The following subsections describe these classes.

⁷A domain Ω is ζ -quasiconvex if, for any $\mathbf{x} \in \Omega$ and $\mathbf{y} \in \Omega$, the length ℓ of the shortest path in Ω that connects $\mathbf{x}$ to $\mathbf{y}$ satisfies $\ell \leq \zeta \|\mathbf{x} - \mathbf{y}\|$.

5.1 The class `ControlSpace`

Fireshape models the control space $\mathbf{Q}$ of admissible domains using geometric transformations $\mathbf{T}$ as in (7) (see also Fig. 9). From a theoretical perspective, the transformations $\mathbf{T}$ can be discretized in numerous different ways as long as certain minimal regularity requirements are met. In Fireshape, the class `ControlSpace` allows the following options: (i) Lagrangian finite elements defined on the same mesh employed to solve the state equation, (ii) Lagrangian finite elements defined on a mesh coarser than the mesh employed to solve the state equation, and (iii) tensorized B-splines defined on a separate Cartesian grid (not to be confused with a spline or Bézier parametrization of the domain boundary, see Fig. 13). In this case, the user must specify the boundary and the refinement level of the Cartesian grid (`bbox` and `level`), as well as the order of the underlying univariate B-splines. The user can also set the regularity of B-splines on the boundary of the Cartesian grid (`boundary_regularities`), or restrict perturbations to specific directions (`fixed_dims`).

If discretization (i) can be considered to be the default option, discretization (ii) allows introducing a regularization by discretizing the geometry more coarsely (so-called “regularization by discretization”), whereas B-splines allow constructing transformation with higher regularity (Lagrangian finite elements are only Lipschitz continuous, whereas B-splines can be continuously differentiable and more).

The class `ControlSpace` can be easily extended to include additional discretization options, such as radial basis functions and harmonic polynomials.

Remark 4 The current implementation of the control space $\mathbf{Q}$ includes geometric transformations that can leave the value of a shape function unchanged. For instance, if $J(\Omega) = \int_\Omega \mathbf{d}\mathbf{x}$, then any diffeomorphism $\mathbf{T}$ of the form $\mathbf{T} = \mathbf{I} + \mathbf{V}$ with $\mathbf{V} \neq \mathbf{0}$ in Ω and $\mathbf{V} = \mathbf{0}$ on $\partial\Omega$ satisfy $\mathbf{T}(\Omega) = \Omega$, and thus $J(\mathbf{T}(\Omega)) = J(\Omega)$.

The presence of a shape function kernel in the control space $\mathbf{Q}$ does not represent a problem for steepest descent and quasi Newton optimization algorithms. However, this kernel leads to singular Hessians, and Newton's method cannot be applied in a straightforward fashion. The development of control space descriptions that exclude shape function kernels and allow direct applications of shape Newton methods in the framework of finite element simulations is a highly active research field (Sturm 2016; Paganini and Sturm 2019; Etling et al. 2020). We predict these new control spaces will be included in Fireshape as soon as their description has been established.

Fig. 13 Fireshape allows discretizing transformations $\mathbf{T}$ using tensorized B-splines. On the left, we display a Cartesian grid that covers the computational domain Ω . On the right, we display a quadratic tensorized B-spline

5.2 The class InnerProduct

To formulate a continuous optimization algorithm, we need to specify how to compute lengths of vectors (and, possibly, how to compute the angle between two vectors). The class `InnerProduct` addresses this requirement and allows selecting an inner product $(\cdot, \cdot)_{\mathcal{H}}$ to endow the control space $\mathbf{Q}$ with.⁸

The choice of the inner product affects how steepest-descent directions are computed. Indeed, a steepest-descent direction is a direction $\mathbf{V}$ of length $\|\mathbf{V}\|_{\mathcal{H}} = 1$ such that $dJ(\mathbf{q}, \mathbf{V})$ is minimal. Let $\alpha := \|dJ(\mathbf{q}, \cdot)\|_*$ denote the length of the operator $dJ(\mathbf{q}, \cdot)$ measured with respect to the dual norm. Then, as explained Hinze et al. (2009, p. 103), the steepest descent direction $\mathbf{V}$ satisfies the equation

$$\alpha(\mathbf{V}, \mathbf{W})_{\mathcal{H}} = -dJ(\mathbf{q}, \mathbf{W}) \quad \text{for all } \mathbf{W} \text{ in } \mathcal{H},$$

which clearly depends on the inner product $(\cdot, \cdot)_{\mathcal{H}}$.

The control space $\mathbf{Q}$ can be endowed with different inner products. In Fireshape, the class `InnerProduct` allows the following options: (i) an $H^1(\Omega)$ inner product based on standard Galerkin stiffness and mass matrices, (ii) a Laplace inner product based on the Galerkin stiffness matrix, and (iii) an elasticity inner product based on the linear elasticity mechanical model. These options can be complemented with additional homogeneous Dirichlet boundary conditions to specify parts of the boundary $\partial\Omega$ that are not to be modified during the shape optimization procedure.

Although all three options are equivalent from a theoretical perspective, in practice it has been observed that option (iii) generally leads to geometry updates that result in meshes of higher equality compared to options (i) and (ii). A thorough comparison is available in Iglesias et al. (2018), where the authors also suggest to consider complementing these inner products with terms stemming from Cauchy-Riemann equations to further increase mesh quality. This additional option is readily available in Fireshape.

⁸Note that specifying a norm would suffice to formulate an optimization algorithm. However, as mention in Remark 2, it is computationally more convenient to restrict computations to inner product spaces.

5.3 The classes Objective and PdeConstraint

In the vast majority of cases, users who aim to solve a PDE-constrained shape optimization problem are only required to instantiate the two classes `Objective` and `PdeConstraint`, where they can specify the formula of the function J to be minimized and the weak formulation of its PDE constraint $\mathbf{A}(\Omega, \mathbf{u}) = \mathbf{0}$ (see Sections 2.2 and 2.3, for instance). Since Fireshape is built on top of the finite element library Firedrake, these formulas must be written using the Unified Form Language (UFL). We refer to the tutorials on Firedrake's website (Firedrake-zenodo 2020) for more details about Firedrake and UFL.

5.4 Supplementary classes

Fireshape also includes a few extra classes to specify additional constraints, such as volume or perimeter constraints on the domain Ω , or spectral constraints to control the singular values of the transformation $\mathbf{T}$. For more details about these extra options, we refer to Fireshape's documentation and tutorials (Wechsung and Paganini 2020a).

6 Conclusions

We have introduced Fireshape: an open-source and automated shape optimization toolbox for the finite element software Firedrake. Fireshape is based on the moving mesh method and allows users with minimal shape optimization knowledge to tackle challenging shape optimization problems constrained to PDEs. In particular, Fireshape computes adjoint equations and shape derivatives in an automated fashion, allows decoupled discretizations of control and state variables, and gives full access to Firedrake's and PETSc's discretization and solver capabilities as well as to the extensive optimization library ROL.

Contributions Alberto Paganini and Florian Wechsung have contributed equally to the development of the software Fireshape and to this manuscript.

Funding The work of Florian Wechsung was partly supported by the EPSRC Centre For Doctoral Training in Industrially Focused Mathematical Modelling [grant number EP/L015803/1].

Compliance with ethical standards

Conflict of interest The authors declare that they have no conflict of interest.

Replication of results For reproducibility, we cite archives of the exact software versions used to produce the results in this paper. All major Firedrake components have been archived on Zenodo (2020). An installation of Firedrake with components matching those used to produce the results in this paper can be obtained following the instructions at <https://www.firedrakeproject.org/download.html>.

The exact version of the Fireshape library used for these results has also been archived at (Wechsung and Paganini 2020b). The latest version of the Fireshape library can be found at <https://github.com/Fireshape/Fireshape>.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Abramowski T, Sugalski K (2017) Energy saving procedures for fishing vessels by means of numerical optimization of hull resistance. *Zeszyty Naukowe Akademii Morskiej w Szczecinie*
- Ahrens J, Geveci B, Law C (2005) 36 - ParaView: an end-user tool for large-data visualization. In: Hansen CD, Johnson CR (eds) *Visualization handbook*. Butterworth-Heinemann, Burlington, pp 717–731
- Allaire G (2007) *Conception optimale de structures*, Mathematics & Applications, vol 58. Springer, Berlin
- Allaire G, Jouve F, Toader AM (2002) A level-set method for shape optimization. *C R Math Acad Sci Paris* 334(12):1125–1130
- Allaire G, Pantz O (2006) Structural optimization with FreeFem++. *Struct Multidiscip Optim* 32(3):173–181
- Alnæs MS, Logg A, Ølgaard KB, Rognes ME, Wells GN (2014) Unified form language: a domain-specific language for weak formulations of partial differential equations. *ACM Trans Math Softw* 40(2):9:1–9:37
- Amestoy PR, Duff IS, L'Excellent JY, Koster J (2000) MUMPS: a general purpose distributed memory sparse solver. In: *International workshop on applied parallel computing*. Springer, pp 121–130
- Balay S, Abhyankar S, Adams MF, Brown J, Brune P, Buschelman K, Dalcin L, Eijkhout V, Gropp WD, Karpeyev D, Kaushik D, Knepley MG, May DA, McInnes LC, Mills RT, Munson T, Rupp K, Sanan P, Smith BF, Zampini S, Zhang H, Zhang H (2019) PETSC users manual. Tech. Rep. ANL-95/11 - Revision 3.11, Argonne National Laboratory
- Balay S, Gropp WD, McInnes LC, Smith BF (1997) Efficient management of parallelism in object oriented numerical software libraries. In: Arge E, Bruaset AM, Langtangen HP (eds) *Modern software tools in scientific computing*. Birkhäuser Press, pp 163–202
- Bendsøe MP (1989) Optimal shape design as a material distribution problem. *Struct Optim* 1(4):193–202
- Bendsoe MP, Sigmund O (2004) *Topology optimization: theory, methods and applications*. Springer, Berlin
- Bernland A, Wadbro E, Berggren M (2019) Shape optimization of a compression driver phase plug. *SIAM J Sci Comput* 41(1):B181–B204
- Blank L, Garcke H, Farshbaf-Shaker MH, Styles V (2014) Relating phase field and sharp interface approaches to structural topology optimization. *ESAIM Control Optim Calc Var* 20(4):1025–1058
- Braess D (2007) *Finite elements*, 3rd edn. Cambridge University Press, Cambridge
- Chevalier C, Pellegrini F (2008) PT-SCOTCH: a tool for efficient parallel graph ordering. *Parallel Comput* 34(6):318–331
- Ciarlet PG (2002) *The finite element method for elliptic problems*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA
- Courtais A, Lesage F, Privat Y, Frey P, Razak Latifi A (2019) Adjoint system method in shape optimization of some typical fluid flow patterns. In: Kiss AA, Zondervan E, Lakerveld R, Özkan L (eds) *29th European symposium on computer aided process engineering, computer aided chemical engineering*, vol 46. Elsevier, pp 871–876
- Dalcin LD, Paz RR, Kler PA, Cosimo A (2011) Parallel distributed computing using Python. *Adv Water Res* 34(9):1124–1139. *New Computational Methods and Software Tools*
- Dapogny C, Faure A, Michailidis G, Allaire G, Couvelas A, Estevez R (2017) Geometric constraints for shape and topology optimization in architectural design. *Comput Mech* 59(6):933–965
- Dapogny C, Frey P, Omnès F, Privat Y (2018) Geometrical shape optimization in fluid mechanics using FreeFem++. *Struct Multidiscip Optim* 58(6):2761–2788
- Delfour MC, Zolésio JP (2011) *Shapes and geometries. Metrics, analysis, differential calculus, and optimization*, Advances in Design and Control, vol 22, 2nd edn. Society for Industrial and Applied Mathematics (SIAM), Philadelphia
- Dilgen CB, Dilgen SB, Fuhrman DR, Sigmund O, Lazarov BS (2018) Topology optimization of turbulent flows. *Comput Methods Appl Mech Engrg* 331:363–393
- Dokken JS, Funke SW, Johansson A, Schmidt S (2019) Shape optimization using the finite element method on multiple meshes with nitsche coupling. *SIAM J Sci Comput* 41(3):A1923–A1948
- Dokken JS, Mitusch SK, Funke SW (2020) Automatic shape derivatives for transient PDEs in FEniCS and Firedrake
- Dondl P, Poh PSP, Rumpf M, Simon S (2019) Simultaneous elastic shape optimization for a domain splitting in bone tissue engineering. *Proc A* 475(2227):20180718, 17
- Eigel M, Sturm K (2018) Reproducing kernel hilbert spaces and variable metric algorithms in pde-constrained shape optimization. *Opt Methods Softw* 33(2):268–296
- Elman HC, Silvester DJ, Wathen AJ (2014) *Finite elements and fast iterative solvers: with applications in incompressible fluid dynamics*, 2nd edn. Oxford University Press, Oxford
- Etling T, Herzog R, Loayza E, Wachsmuth G (2020) First and second order shape optimization based on restricted mesh deformations. *SIAM J Sci Comput* 42(2):A1200–A1225
- Farrell PE, Ham DA, Funke SW, Rognes ME (2013) Automated derivation of the adjoint of high-level transient finite element programs. *SIAM J Sci Comput* 35(4):C369–C393
- Farrell PE, Mitchell L, Wechsung F (2019) An augmented lagrangian preconditioner for the 3d stationary incompressible navier-stokes equations at high reynolds number. *SIAM J Sci Comput* 41(5):A3073–A3096
- Feppon F, Allaire G, Dapogny C, Jolivet P (2020) Topology optimization of thermal fluid-structure systems using body-fitted

- meshes and parallel computing. *J Comput Phys* 417(109574): 30
- Firedrake-zenodo (2020) Software used in 'Fireshape: a shape optimization toolbox for Firedrake' (Version Firedrake_20200918.0). Zenodo. <https://doi.org/10.5281/zenodo.4036176>. Accessed 18 Sept 2020
- Garcke H, Hecht C, Hinze M, Kahle C (2015) Numerical approximation of phase field based shape and topology optimization for fluids. *SIAM J Sci Comput* 37(4):A1846–A1871
- Geuzaine C, Remacle JF (2009) Gmsh: a 3-d finite element mesh generator with built-in pre-and post-processing facilities. *Int J Numer Methods Eng* 79(11):1309–1331
- Gibson TH, Mitchell L, Ham DA, Cotter CJ (2020) Slate: extending Firedrake's domain-specific abstraction to hybridized solvers for geoscience and beyond. *Geosci Model Develop* 13(2):735–761
- Ham DA, Mitchell L, Paganini A, Wechsung F (2019) Automated shape differentiation in the Unified Form Language. *Struct Multidiscip Optim* 60(5):1813–1820
- He X, Li J, Mader CA, Yildirim A, Martins JR (2019) Robust aerodynamic shape optimization—from a circle to an airfoil. *Aerospace Sci Technol* 87:48–61
- Hinze M, Pinnau R, Ulbrich M, Ulbrich S (2009) Optimization with PDE constraints, Mathematical modelling: Theory and Applications, vol 23. Springer, New York
- Höllig K (2003) Finite element methods with B-splines. Society for Industrial and Applied Mathematics (SIAM), Philadelphia PA
- Homolya M, Kirby RC, Ham DA (2017) Exposing and exploiting structure: optimal code generation for high-order finite element methods. *arXiv:1711.02473*
- Homolya M, Mitchell L, Loporini F, Ham DA (2018) TSFC: a structure-preserving form compiler. *SIAM J. Sci. Comput.* 40(3):C401–C428
- Hypre (2020) High performance preconditioners. Lawrence Livermore National Laboratory. <https://computing.llnl.gov/projects/hypre-scalable-linearsolvers-multigrid-methods>. Accessed 25 Jan 2021
- Iglesias JA, Sturm K, Wechsung F (2018) Two-dimensional shape optimization with nearly conformal transformations. *SIAM J Sci Comput* 40(6):A3807–A3830
- Ito K, Kunisch K (2008) Lagrange multiplier approach to variational problems and applications, *Advances in Design and Control*, vol 15. Society for Industrial and Applied Mathematics (SIAM), Philadelphia
- Kelley CT (2003) Solving nonlinear equations with Newton's method, *Fundamentals of Algorithms*, vol 1. Society for Industrial and Applied Mathematics (SIAM), Philadelphia
- Laurain A (2018) A level set-based structural optimization code using FEniCS. *Struct Multidiscip Optim* 58(3):1311–1334
- Liu Z, Korvink J, Huang R (2005) Structure topology optimization: fully coupled level set method via femlab. *Struct Multidiscip Optim* 29(6):407–417
- Lombardi M, Parolini N, Quarteroni A, Rozza G (2012) Numerical simulation of sailing boats: dynamics, fsi, and shape optimization. In: *Variational analysis and aerospace engineering: mathematical challenges for aerospace design*. Springer, pp 339–377
- Loporini F, Ham DA, Kelly PHJ (2017) An algorithm for the optimization of finite element integration loops. *ACM Trans Math Softw* 44:3:1–3:26
- Marinić-Kragić I, Vučina D, Ćurković M (2016) Efficient shape parameterization method for multidisciplinary global optimization and application to integrated ship hull shape optimization workflow. *Comput-Aided Des* 80:61–75
- Mitchell L, Müller EH (2016) High level implementation of geometric multigrid solvers for finite element problems: applications in atmospheric modelling. *J Comput Phys* 327:1–18
- Mitusch SK, Funke SW, Dokken JS (2019) dolfin-adjoint 2018.1: automated adjoints for FEniCS and Firedrake. *J Open Source Softw* 4(38):1292
- Nocedal J, Wright SJ (2006) Numerical optimization, 2nd edn. Springer, New York
- Paganini A, Sturm K (2019) Weakly normal basis vector fields in rkhs with an application to shape newton methods. *SIAM J Numer Anal* 57(1):1–26
- Paganini A, Wechsung F, Farrell PE (2018) Higher-order moving mesh methods for PDE-constrained shape optimization. *SIAM J Sci Comput* 40(4):A2356–A2382
- Rathgeber F, Ham DA, Mitchell L, Lange M, Loporini F, McRae ATT, Bercea GT, Markall GR, Kelly PHJ (2016) Firedrake: automating the finite element method by composing abstractions. *ACM Trans Math Softw* 43(3):24:1–24:27
- Ridzal D, Kouri D (2014) Rapid optimization library, version 00. <https://www.osti.gov/servlets/purl/1232084>
- Schmidt S, Ilıc C, Schulz V, Gauger NR (2013) Three-dimensional large-scale aerodynamic shape optimization based on shape calculus. *AIAA* 51(11):2615–2627
- Schulz V, Gherman I (2009) One-shot methods for aerodynamic shape optimization. In: *MEGADESIGN and megaopt-german initiatives for aerodynamic simulation and optimization in aircraft design*. Springer, pp 207–220
- Sigmund O (2001) A 99 line topology optimization code written in matlab. *Struct Multidiscip Optim* 21(2):120–127
- Sturm K (2016) Convergence of newton's method in shape optimisation via approximate normal functions. *arXiv:1608.02699*
- Talischí C, Paulino GH, Pereira A, Menezes IF (2012) Polytop: a matlab implementation of a general topology optimization framework using unstructured polygonal finite element meshes. *Struct Multidiscip Optim* 45(3):329–357
- Villanueva CH, Maute K (2017) CutFEM topology optimization of 3D laminar incompressible flow problems. *Comput Methods Appl Mech Engrg* 320:444–473
- Wang H, Doherty J, Jin Y (2018) Hierarchical surrogate-assisted evolutionary multi-scenario airfoil shape optimization. In: *2018 IEEE congress on evolutionary computation (CEC)*, pp 1–8. IEEE
- Wechsung F (2019) Shape optimisation and robust solvers for incompressible flow. Ph.D. thesis, University of Oxford
- Wechsung F, Paganini A (2020a) Fireshape's Documentation. Read the Docs. <https://fireshape.readthedocs.io/en/latest/index.html>. Accessed 18 Sept 2020
- Wechsung F, Paganini A (2020b) Fireshape: a shape optimization toolbox for Firedrake (Version v0.3). Zenodo. <https://doi.org/10.5281/zenodo.4036213>. Accessed 18 Sept 2020
- Wendland H (2005) Scattered data approximation. Cambridge University Press, Cambridge
- Zenodo (2020) Software used in 'Fireshape: a shape optimization toolbox for Firedrake'. <https://doi.org/10.5281/zenodo.4036176>
- Zhou S, Li Q (2008) A variational level set method for the topology optimization of steady-state Navier-Stokes flow. *J Comput Phys* 227(24):10178–10195

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.