

# Shape optimisation for faster washout in recirculating flows

J.G. Williams<sup>1†</sup>, F. Wechsung<sup>2</sup>, B.W. Turney<sup>3</sup>, S.L. Waters<sup>4</sup>, and D.E. Moulton<sup>4</sup>

<sup>1</sup>Harvard-MIT Biomedical Engineering Center, Massachusetts Institute of Technology, 77 Massachusetts Ave, Cambridge, MA 02139

<sup>2</sup>Courant Institute of Mathematical Sciences, New York University, 251 Mercer Street, New York, N.Y. 10012-1185

<sup>3</sup>Nuffield Department of Surgical Sciences, University of Oxford, OX3 9DU

<sup>4</sup>Mathematical Institute, University of Oxford, Woodstock Road, OX2 6GG, UK

(Received xx; revised xx; accepted xx)

How to design an optimal biomedical flow device to minimise trapping of undesirable biological solutes/debris and/or enhance their washout is a pertinent but complex question. While biomedical devices often utilise externally-driven flows to enhance washout, the presence of vortices – arising as a result of fluid flows within cavities – hinder washout by trapping debris. Motivated by this, we solve the steady, incompressible Navier–Stokes equations for flow through channels into and out of a 2D cavity. In endourology, the presence of vortices – enhanced by flow symmetry breaking – has been linked to long washout times of kidney stone dust in the renal pelvis cavity, with dust transport modelled via advection and diffusion of a passive tracer (Williams *et al.* 2020). Here we determine the inflow and outflow channel geometries that minimise washout times.

For a given flow field  $\mathbf{u}$ , vortices are characterised by regions where  $\det \nabla \mathbf{u} > 0$  (Jeong & Hussain 1995). Integrating a smooth form of  $\max(0, \det \nabla \mathbf{u})$  over the domain provides an objective to minimise recirculation zones (Kasumba & Kunisch 2012). We employ adjoint-based shape optimisation to identify inflow and outflow channel geometries that reduce this objective. We show that a reduction in the vortex objective correlates with reduced washout times. We additionally show how multiple solutions to the flow equations lead to solution branch-switching during the optimisation routine by characterising the change in solution bifurcation structure with the change in scope tip geometry.

**Key words:**

## 1. Introduction

Many biomedical devices are designed to minimise local flow recirculation to avoid the trapping of undesirable biological solutes/agents and debris and/or enhance washout of such solutes. For example, vortices resulting from blood flow through arterial stents can serve as accumulation sites for inflammatory elements, leading to arterial thrombosis, and a streamlined stent strut design is recommended to reduce regions of closed streamlines (Jiménez & Davies 2009). Additionally, inducing 3D spiral flow through the stented region of an artery has been shown to enhance washout through the clearance of recirculation

---

† Email address for correspondence: jgw@mit.edu

zones (Coppola & Caro 2009; Xhang *et al.* 2019). The trapping of debris within regions of vortical flow is also undesirable in medical imaging procedures such as endoscopy or optical coherence tomography. In these procedures, flushing with saline or a contrast agent is used to remove biological debris and blood, respectively, and consequently to enhance visualisation.

Recirculating vortices and the associated mass transport also has a number of applications in cleaning and decontamination processes (see Min *et al.* (2020) and references therein). During the removal of the chemical agents used in these processes, agents can be trapped in the neighbourhood of, or adsorbed onto, the rough surface of the object being cleaned. To investigate the role of rough surface topography on the boundary layer flow of contaminant, Min *et al.* (2020) considered the fluid flow over a backward facing step and the associated mass transport properties of a passive scalar. Recirculation regions were found adjacent to the vertical face of the step, and the study revealed a subtle interplay between the advective and diffusion transport mechanisms of a passive scalar.

In many biomedical devices, recirculation regions of flow arise as a result of flow into or out of cavities (Clavica *et al.* 2014; Williams *et al.* 2020). Here we draw particular motivation from urological examples. In *ureteroscopy*, a minimally invasive procedure for kidney stone removal, a long thin *ureteroscope*, which contains a hollow tube called the *working channel*, is inserted through the urinary tract and into the *renal pelvis*, the hollow cavity within the kidney. Dust-like fragments of kidney stones – which are broken up using thermal energy introduced via high frequency laser pulses – are cleared from the kidney by a continuous flow of *irrigation* fluid through the working channel tip and into the kidney. The fluid exits the body in the opposite direction, along the outside of the scope shaft. In ureteroscopy, visualisation is provided to the clinician via a small light and camera embedded in the scope tip, and thus a rapid and thorough clearing of kidney stone dust is important for maintaining a clear field of view. As another example, when the kidney stone blocks the ureter connecting the bladder to the renal pelvis, a urinary stent – a hollow tube with sideholes in the wall to permit the transport of fluid between the intra and extraluminal spaces – is deployed to promote drainage. In the neighbourhood of a stone the extraluminal region takes the form of a closed cavity with flow entering via a stent sidehole and exiting along the extraluminal outflow. The development of recirculating regions/vortices causes particles (stone crystals and bacteria) to become trapped before they are washed away, resulting in stent encrustation and potential stent failure (Clavica *et al.* 2014). Percutaneous nephrolithotomy (PCNL) provides yet another urological example. In this procedure, kidney stones too large to be removed using minimally invasive procedures such as ureteroscopy, are manually extracted with a nephroscope which is inserted into the patients kidney via an abdominal incision. As with ureteroscopy, continuous irrigation is required to wash residual stone particles from the kidney. A previous numerical and experimental study examined the effects of irrigation flow speed and nephroscope angle (with respect to the bottom of the treated calyx) on particle clearance rate (Paster *et al.* 2019). It was found that clearance was most efficient with the nephroscope positioned perpendicularly to the bottom of the calyx using the fastest tested flow speed.

Consideration of such cavity flows motivates the question of how best to design the device and/or the procedure to minimise the washout time of biological debris. We recently considered cavity flow and washout in a two-dimensional cavity with inlet and outlet channels modelling fluid flow with the steady incompressible Navier–Stokes equations (Williams *et al.* 2020). We demonstrated the presence of symmetry breaking in the flow structure with increasing Reynolds number, such that the asymmetric flow observed at higher Reynolds numbers is dominated by a large central vortex. When

incorporating washout of a passive tracer via advection and diffusion, we found a strong link between washout time and the presence, size, and location of vortices. This result fits a general trend: the time required for a passive tracer to exit regions of closed streamlines in a steady flow scales both with the size of the vortex and decreasing diffusion coefficient of the tracer (Rhines & Young 1983). Thus, for fixed diffusive properties, decreasing the size of vortices is paramount for reduction in washout time (Abergel & Temam 1990).

In this paper we reveal how alternative shapes for the inflow channel can be used to enhance the clearance of a passive tracer, and determine the inflow and outflow channel geometries for which the washout time of an initial concentration of passive tracer is minimised. We approach this question using shape optimisation techniques, which rely on the identification of an appropriate objective to minimise. Ideally, the objective would be the washout time for a given initial concentration of tracer. However, such an objective depends upon features of both the flow and concentration fields, and thus optimisation presents a computationally prohibitive task: even if we consider a steady flow field which is unaffected by the motion of the tracer, computing the washout for any given shape requires solving the fluid equations and subsequently solving a transient advection-diffusion problem for the motion of the tracer. Instead, to define a tractable shape optimisation problem, we determine an objective that is purely based on properties of the steady flow field, but which correlates with washout time.

As noted, we expect the size of vortices to serve as a useful proxy for washout time, and thus base our objective on the size of vortical regions, though quantifying the presence of vortices is in itself a nontrivial pursuit. A logical choice would be to consider vorticity magnitude (Jeong & Hussain 1995). This has been previously used as an objective in shape optimisation to reduce recirculation zones (Abergel & Temam 1990), although more recently objectives which focus on properties of the determinant of the velocity gradient have been shown to achieve superior results (Kasumba & Kunisch 2010, 2012). For a given flow field  $\mathbf{u}$ , vortices are characterised by regions where  $\det \nabla \mathbf{u} > 0$  (Jeong & Hussain 1995). Thus, integrating a smooth form of  $\max(0, \det \nabla \mathbf{u})$  over the desired domain provides an objective that can be used with a shape optimisation algorithm to determine the optimal inflow and outflow channel shapes to minimise recirculation zones (Kasumba & Kunisch 2012). We employ such an objective and compute optimal shapes over a range of fluid properties, characterised by Reynolds number. We then confirm that this objective correlates well with washout time across the range of parameters tested, including the diffusivity of tracer particles, and show that the optimised channel shapes can provide a significant decrease in washout time.

## 2. Problem description

We consider a simplified, two-dimensional representation of inflow and outflow to/from a rectangular cavity, though we also briefly explore channel optimisation for different cavity geometries in Section 4. This setup is motivated by an idealisation of a ureteroscopy tip within the renal pelvis, but more generally provides a toy geometry for considering shape optimisation on flows characterised by large vortical structures. We employ a Cartesian coordinate system  $(x^*, y^*)$  (stars indicate dimensional variables throughout) with corresponding coordinate directions  $\mathbf{i}$  and  $\mathbf{j}$ .

The flow domain is represented by the light pink and light orange regions in Figure 1, i.e.  $\Omega_t^* \cup \Omega_c^*$ . The  $\Omega_t^*$  domain encompasses both inflow and outflow channels. The inflow boundary to the working channel is  $\Gamma_{\text{in}}^*$ , the upper and lower outflow boundaries are  $\Gamma_{\text{out,u}}^*$  and  $\Gamma_{\text{out,l}}^*$ , respectively,  $\Gamma_f^*$  is the solid boundary we will allow to vary in the shape optimisation, and  $\Gamma_{\text{wall}}^*$  denotes all other (fixed) solid boundaries. Following on




Figure 1: Problem set-up. Fixed boundaries are  $\Gamma_{\text{out,u}}^*$ ,  $\Gamma_{\text{out,l}}^*$ ,  $\Gamma_{\text{in}}^*$ ,  $\Gamma_{\text{wall}}^*$ ; the red boundary  $\Gamma_{\text{f}}^*$  is the control boundary that we allow to vary in the shape optimisation.

our previous experimental and numerical work in a similar geometry (Williams *et al.* 2020), we use the parameter values  $l_1 = 1.2743$  cm,  $l_2 = 1$  cm for the channel and cavity lengths, respectively, and set  $a = 0.06$  cm,  $b = 0.25701$  cm, and  $h = 0.11598$  cm in the initial geometry.

The dimensional velocity field is given by  $\mathbf{u}^* = u^*\mathbf{i} + v^*\mathbf{j}$  and pressure  $p^*$ . The flow is governed by the steady, incompressible Navier–Stokes and continuity equations, subject to the following boundary conditions. The flow is driven by an imposed fully-developed parabolic profile at the inlet boundary with maximum velocity  $U$ . As outflow conditions, we impose zero normal stress and parallel velocity. On all other boundaries (solid black and red lines in Figure 1) we assume no-slip and no normal velocity conditions.

A passive tracer of concentration  $c^*(\mathbf{x}^*, t^*)$  over time  $t^*$  is advected by the flow and also diffuses with constant diffusion coefficient  $\mathcal{D}$ . We impose no total flux (which has both advective and diffusive components) of the tracer through the inflow or solid walls ( $\Gamma_{\text{in}}^*$ ,  $\Gamma_{\text{wall}}^*$ ,  $\Gamma_{\text{f}}^*$ ) and no diffusive flux of tracer through the outflows ( $\Gamma_{\text{out,u}}^*$ ,  $\Gamma_{\text{out,l}}^*$ ). We assume that initially the passive tracer has constant concentration  $C$  that is uniform within the cavity and zero elsewhere.

### 2.1. Dimensionless governing equations

We nondimensionalise as follows

$$\mathbf{x} = \frac{\mathbf{x}^*}{a}, \quad t = \frac{Ut^*}{a}, \quad \mathbf{u} = \frac{\mathbf{u}^*}{U}, \quad p = \frac{a(p^* - p_{\text{ref}})}{U\mu}, \quad c = \frac{c^*}{C}, \quad (2.1a-e)$$

where  $p_{\text{ref}}$  is the pressure at the outflow boundaries of the truncated domain  $\Gamma_{\text{out,u}}^*$ ,  $\Gamma_{\text{out,l}}^*$ , and we note that time has been non-dimensionalised using the advective timescale. The steady, incompressible Navier–Stokes equations are

$$\text{Re}(\mathbf{u} \cdot \nabla) \mathbf{u} = -\nabla p + \nabla \cdot \mathcal{E}(\mathbf{u}), \quad \nabla \cdot \mathbf{u} = 0, \quad \text{in } \Omega_c \cup \Omega_t, \quad (2.2a,b)$$



where Reynolds number  $\text{Re} = \rho U a / \mu$ ,  $\mathcal{E}(\mathbf{u}) = \nabla \mathbf{u} + (\nabla \mathbf{u})^\top$ , and  $\nabla = (\partial/\partial x, \partial/\partial y)$ .<sup>†</sup> Flow boundary conditions are

$$u = [1 - y^2], \quad v = 0, \quad \text{on } \Gamma_{\text{in}}, \quad (2.3\text{a,b})$$

$$2 \frac{\partial u}{\partial x} - p = 0, \quad v = 0, \quad \text{on } \Gamma_{\text{out,l}} \cup \Gamma_{\text{out,u}}, \quad (2.3\text{c,d})$$

$$u = v = 0, \quad \text{on } \Gamma_{\text{wall}} \cup \Gamma_{\text{f}}. \quad (2.3\text{e})$$

The advection-diffusion equation for tracer concentration is

$$\frac{\partial c}{\partial t} = \text{Re}^{-1} \text{Sc}^{-1} \nabla^2 c - \mathbf{u} \nabla c, \quad (2.4)$$

where the Schmidt number  $\text{Sc} = \mu / (\rho \mathcal{D})$  represents the ratio of momentum diffusivity to tracer diffusivity. Equation (2.4) is solved subject to the following boundary conditions

$$\mathbf{J} \cdot \mathbf{n} = 0, \quad \text{on } \Gamma_{\text{in}}, \Gamma_{\text{wall}}, \Gamma_{\text{f}}, \quad \nabla c \cdot \mathbf{n} = 0, \quad \text{on } \Gamma_{\text{out}}, \quad (2.5\text{a,b})$$

where  $\mathbf{J} = -\text{Sc}^{-1} \text{Re}^{-1} \nabla c + \mathbf{u} c$ . Additionally, equation (2.4) is solved subject to the initial condition

$$c(\mathbf{x}, 0) = \begin{cases} 1, & \text{in } \Omega_{\text{c}}, \\ 0, & \text{in } \Omega_{\text{t}}. \end{cases} \quad (2.6)$$

## 2.2. Washout time

Following Williams *et al.* (2020), to quantify the rate at which the passive tracer exits the entire domain  $\Omega_{\text{c}} \cup \Omega_{\text{t}}$ , we define the percentage loss at time  $t$  as

$$\% \text{ loss} = \frac{\iint_{\Omega_{\text{c}} \cup \Omega_{\text{t}}} [c(\mathbf{x}, 0) - c(\mathbf{x}, t)] \, \text{d}\Omega}{\iint_{\Omega_{\text{c}} \cup \Omega_{\text{t}}} c(\mathbf{x}, 0) \, \text{d}\Omega}, \quad (2.7)$$

and we define a metric for *washout time*,  $T_{90}$  corresponding to the time required for 90% of the tracer to exit the domain, i.e.,

$$\frac{\iint_{\Omega_{\text{c}} \cup \Omega_{\text{t}}} [c(\mathbf{x}, 0) - c(\mathbf{x}, T_{90})] \, \text{d}\Omega}{\iint_{\Omega_{\text{c}} \cup \Omega_{\text{t}}} c(\mathbf{x}, 0) \, \text{d}\Omega} = 0.9. \quad (2.8)$$

## 2.3. Optimisation objectives

As discussed in Section 1, vortices within a flow field are associated with regions where  $\det \nabla \mathbf{u} > 0$  (Jeong & Hussain 1995). Therefore,

$$\iint_{\Omega_{\text{c}}} \max(0, \det \nabla \mathbf{u}) \, \text{d}\Omega \quad (2.9)$$

is a logical objective to associate with recirculation zones. However, to ensure differentiability of the objective and enhance optimisation performance we utilise the smoothing

<sup>†</sup> The long formulation of equation (2.2a), rather than reducing  $\nabla \cdot \mathcal{E}(\mathbf{u})$  to  $\nabla^2 \mathbf{u}$  using equation (2.2b), ensures zero-stress as the natural boundary condition in the weak formulation of the equations.

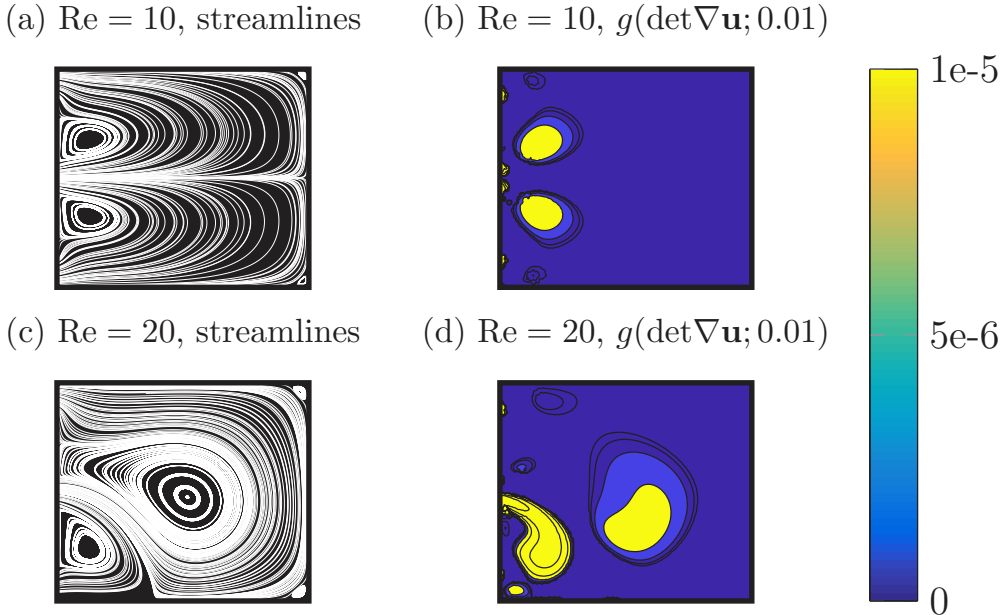


Figure 2: Figures (a), (c) Streamlines in the cavity for  $Re = 10$ ,  $Re = 20$ , respectively, and (b), (d) corresponding contours of  $g(\det \nabla \mathbf{u}; 0.01)$ , where  $g(s; \epsilon)$  is given by equation (2.10).

function

$$g(s; \epsilon) = \begin{cases} 0, & s \leq 0, \\ s^3 / (s^2 + \epsilon), & s > 0, \end{cases} \quad (2.10)$$

and consider minimising

$$\iint_{\Omega_c} g(\det \nabla \mathbf{u}; \epsilon) d\Omega, \quad (2.11)$$

as introduced by Kasumba & Kunisch (2010, 2012). We choose  $\epsilon = 0.01$  – as opposed to  $\epsilon = 1$  used by Kasumba & Kunisch (2010, 2012) – to well-capture vortices within the flow. Contours of  $g(\det \nabla \mathbf{u}; 0.01)$  for  $Re = 10$  and  $Re = 20$  are plotted in Figures 2b and 2d, respectively. Comparison with corresponding streamlines for  $Re = 10$  and  $Re = 20$  are shown in Figures 2a and 2c, respectively, demonstrates strong qualitative correlation between  $g(\det \nabla \mathbf{u}; 0.01)$  and areas of closed streamlines.

## 2.4. Numerical details

### 2.4.1. State equation (Navier–Stokes) solver

Our implementation is based on the Firedrake finite-element library (Rathgeber *et al.* 2016; Luporini *et al.* 2017; Homolya *et al.* 2017b,a) which uses the Unified Form Language (UFL) (Alnæs *et al.* 2014). UFL is a domain specific language embedded in Python that enables the expression of partial differential equations in their weak form in close to mathematical notation. We discretise the equations using continuous, piecewise quadratic Lagrange finite elements for the velocity field, and discontinuous, piecewise affine Lagrange finite elements for the pressure. This element pair is also referred to as the quadratic Scott-Vogelius element, with the velocity space denoted by  $V_h$  and the pressure

space by  $Q_h$ . While the Scott-Vogelius element pair is not known to be inf-sup stable on arbitrary triangulations (see (Guzmn & Scott 2018) for a recent discussion), it has been shown to be inf-sup stable on meshes with a particular structure. To obtain an apposite mesh structure, we perform a barycentric refinement of an unstructured mesh generated by GMSH (Geuzaine & Remacle 2009); this guarantees inf-sup stability (Qin 1994). We choose the Scott-Vogelius element as it enforces the incompressibility constraint exactly, a property that more recently has been understood to be of importance for the accuracy of numerical solutions to incompressible flow (John *et al.* 2017).

Since the Navier–Stokes equations are non-linear, we employ a Newton-scheme. The Jacobian required for this is derived automatically and symbolically by the UFL library. To assemble the linear systems arising in each step of the Newton iteration, Firedrake generates highly optimised C code at runtime. The system is then solved using the MUMPS (Amestoy *et al.* 2000) sparse direct solver. For higher accuracy, we perform a few (usually two) iterations of PETSc’s (Balay *et al.* 2016, 2018, 1997; Dalcin *et al.* 2011) GMRES implementation using the LU factorisation as a preconditioner, until the residual is reduced to  $10^{-10}$ .

Starting from a zero initial guess, the solver finds a symmetric solution. However, for large enough Reynolds number, this solution is not stable. For these cases we perform a first solve with a no-slip boundary condition applied to one of the outlets, and use this solution as the initial guess, to guarantee that the stable, asymmetric solution branch is found and subsequently used for the initial step of the optimisation routine (Williams *et al.* 2020).

#### 2.4.2. Optimisation

For the optimisation we rely on the Fireshape shape optimisation library (Paganini & Wechsung 2020). Fireshape is an open source library that connects the Firedrake finite element library with the Trilinos Rapid Optimisation Library (Ridzal & Kouri 2014). Fireshape employs the perturbation of identity approach (Micheletti 1972; Murat & Simon 1976; Simon 1980; Delfour & Zolésio 2011), which is based on deforming an initial domain  $\Omega_t$  using deformations of the form  $\mathbf{T} = \mathbf{Id} + \mathbf{X}$ , where  $\mathbf{Id}(\mathbf{x}) = \mathbf{x}$  is the identity mapping. The optimisation is then formulated as finding an optimal deformation  $\mathbf{X}$  and at each iteration, the shape is given by  $\Omega_t^{(k)} = (\mathbf{Id} + \mathbf{X}^{(k)})(\Omega_t)$ . Denoting by  $X_h$  the space of continuous, piecewise affine vector fields that vanish on  $\partial\Omega_t \setminus \Gamma_f$ , the vanilla shape optimisation problem can then be expressed as

$$\begin{aligned} \underset{(\mathbf{u}, p) \in V_h \times Q_h, \mathbf{X} \in X_h}{\text{minimise}} \quad & J_e(\Omega_c) = \iint_{\Omega_c} g(\det \nabla \mathbf{u}, \epsilon) d\Omega \\ \text{subject to} \quad & (\mathbf{u}, p) \text{ satisfy (2.2) in } (\mathbf{Id} + \mathbf{X})(\Omega_t) \cup \Omega_c. \end{aligned} \tag{2.12}$$

We remark that while the integration domain  $\Omega_c$  in the objective stays fixed, it is the velocity field in the integrand that changes as the channel domain  $\Omega_t$  is deformed.

Both analysing and solving shape optimisation problems is known to be difficult for several reasons. First, shape objectives are usually highly nonlinear and non-convex in the control  $\mathbf{X}$ , hence a unique optimal shape often does not exist. Second, when deforming meshes one needs to ensure that the mesh does not degenerate or overlap. Fireshape has several capabilities that allow us to mitigate these problems.

To avoid the first issue, we add a Tikhonov type term to the objective, that is a term of the form  $\alpha \|\mathbf{X}\|_{X_h}^2$  (we choose  $\alpha = 0.01$  in our simulations). Such a term is commonly used for ill-posed optimisation problems as it is strictly convex. The norm  $\|\cdot\|_{X_h}^2$  on the

space  $X_h$  is chosen to penalise nonsmoothness of the deformations

$$\begin{aligned} \|\mathbf{X}\|_{X_h}^2 &:= (\mathbf{X}, \mathbf{X})_{X_h}, \\ (\mathbf{X}, \mathbf{Y})_{X_h} &:= \iint_{\Omega_t} \mathcal{E}(\mathbf{X}) : \mathcal{E}(\mathbf{Y}) + 0.001 \Delta \mathbf{X} : \Delta \mathbf{Y} d\Omega. \end{aligned} \quad (2.13)$$

Since standard finite elements only admit a first (weak) derivative, we employ a standard interior penalty method (Brenner & Sung 2005; Wells *et al.* 2006). To counteract the second issue, we add a term  $\beta \|\max(0, \sigma_{\max}(\nabla \mathbf{X}) - 0.9)\|_{L^2}^2$  to the objective, that penalises large values of the largest singular value of the gradient of  $\mathbf{X}$  (we choose  $\beta = 10^3$ ). This modification is motivated by a result that guarantees that for a convex domain, the deformed mesh does not overlap if the largest singular value of  $\nabla \mathbf{X}$  is smaller than 1 (Delfour & Zolésio 2011, §3, Theorem 2.15). Finally, we note that any shape objective (such as  $J_e$ ), is invariant with respect to deformations of the interior mesh nodes: deforming the interior of the mesh does not change the overall shape and hence does not have an impact on the objective. To remove this invariance, we require that the deformation satisfies

$$(\mathbf{X}, \mathbf{Y})_{X_h} = 0 \quad \text{for all } \mathbf{Y} \in X_{h,0} = \{\mathbf{Y} \in X_h : \mathbf{Y}|_{\partial\Omega_t} = 0\}. \quad (2.14)$$

This is the weak formulation of requiring  $\nabla \cdot \mathcal{E}(\mathbf{X}) + \gamma \Delta^2 \mathbf{X} = 0$  inside the domain that is free to move. We refer to Wechsung (2019, §8.2) for more details on these modifications. To summarise, the minimisation problem that we solve is given by

$$\begin{aligned} &\underset{(\mathbf{u}, p) \in V_h \times Q_h, \mathbf{X} \in X_h}{\text{minimise}} && J_e(\Omega_c) + \alpha \|\mathbf{X}\|_{X_h}^2 + \beta \|\max(0, \sigma_{\max} \nabla \mathbf{X} - 0.9)\|_{L^2}^2 \\ &\text{subject to} && (\mathbf{u}, p) \text{ satisfy (2.2) in } (\mathbf{Id} + \mathbf{X})(\Omega_t) \cup \Omega_c \\ &&& \mathbf{X} \text{ satisfies (2.14).} \end{aligned} \quad (2.15)$$

In order to solve the optimisation problem efficiently the shape derivative is required. Fireshape relies on UFL to automatically derive the required adjoint equations and shape derivatives (Alnæs *et al.* 2014; Ham *et al.* 2019). Finally, the resulting optimisation problem is solved using the L-BFGS implementation of the Rapid Optimization Library (Nocedal & Wright 2006; Ridzal & Kouri 2014).

### 2.4.3. Tracer washout equations

We approximated the concentration function with piecewise linear elements on the same mesh as the velocity field. The velocity field was solved for and subsequently restored to drive the advective term in equation (2.4). Equation (2.4) was discretised in time using a Crank-Nicolson scheme, and we added a mesh-dependent SUPG stabilisation term, as described in Franca *et al.* (1992) as standard Galerkin discretisations of advection-diffusion equations are oscillatory in the advection-dominated regime (Silvester *et al.* 2014). We validated our advection-diffusion code using the method of manufactured solutions, with details presented in Williams *et al.* (2020, §B.2).

## 2.5. Bifurcation diagrams

We computed the bifurcation diagram as a function of  $\text{Re}$  in Figure 6a using an implementation of numerical deflation techniques (Farrell *et al.* 2015) within Firedrake.

### 3. Results

#### 3.1. Washout time for initial shape

The streamlines associated with each Reynolds number are shown in Figures 3a–e. Based on our previous study of the solution structure to the incompressible, steady Navier–Stokes equations in this geometry we anticipate a pitchfork bifurcation near  $\text{Re} = 15$  (Williams *et al.* 2020). Indeed, we find that for  $\text{Re} = 10$  a unique symmetric solution is found, whereas for  $\text{Re} = 20, 30, 40$ , and  $50$ , three solutions exist, with the mirror pair of stable, asymmetric solutions constituting the flow patterns of physical relevance. The large central vortex, which emerges as a consequence of flow symmetry-breaking, grows in size as Reynolds number increases from  $20$  to  $50$  (see Figures 3b–e). In Figures 3a–e we observe that the vortical flow region is surrounded by streamlines that trace a path from the inlet to the outlet without recirculating. We refer to this surrounding region as the ‘racetrack’ region in the discussion below. As the Reynolds number increases, in addition to the size of the vortical region increasing, the speed of the racetrack also increases.

In Figure 3f we plot washout time as a function of the Schmidt number  $\text{Sc}$  for five different values of  $\text{Re}$ . The dimensionless washout time,  $T_{90}$ , is defined in equation (2.8). In many biomedical applications, properties of the fluid such as its viscosity and density are fixed. Additionally, the size of the device is likely constrained by physical anatomy, and hence we consider the lengthscale  $a$  to be fixed. Thus, varying the Reynolds number corresponds to changing the inflow rate characterised by  $U$ . In equation (2.1b), we non-dimensionalise time on the advection timescale  $a/U$ , so that the dimensional washout time,  $T_{90}^* = aT_{90}/U$ . In Figure 3f we present  $T_{90}^*$  normalised with respect to the *fixed* timescale for viscous diffusion of momentum,  $a^2/\nu$ ; this corresponds to  $T_{90}^*\mu/(\rho a^2) = T_{90}/\text{Re}$ .

For  $\text{Sc} = 1$ , corresponding to a balance between momentum and tracer diffusivities, washout time decreases monotonically with Reynolds number; i.e. the faster the flow (characterised by larger  $\text{Re}$ ), the quicker the washout. This can be seen by looking at the symbols on the far-left of Figure 3f. In this regime, the tracer diffusivity is sufficiently high that the tracer can diffuse out of the vortex region making its way onto the racetrack. As the speed of the racetrack increases with  $\text{Re}$ , the washout time monotonically decreases with increasing  $\text{Re}$ . As  $\text{Sc}$  increases – and tracer diffusivity decreases with respect to momentum diffusivity – the vortex-trapping effect becomes more significant, and hence the fastest flow (characterised by  $\text{Re}$ ) no longer provides the quickest washout. Since the proportion of the cavity domain occupied by a vortex also increases with  $\text{Re}$ , for  $\text{Sc} = 5$  we in fact see a reversal in the ordering of washout times as  $\text{Re}$  increases from  $20$  to  $50$ . This is indicated by the lines crossing in Figure 3f between  $\text{Sc} = 1$  and  $\text{Sc} = 5$ . Nevertheless, all these washout times beat the washout time for  $\text{Re} = 10$ , as although the vortex trapping effect at this Reynolds and Schmidt number is minimal due to small vortex size, the speed of the racetrack for this Reynolds number is comparatively slow, particularly near the rear wall of the cavity. The washout time ordering with Reynolds number remains the same for  $\text{Sc} = 10$  as  $\text{Sc} = 5$ . As  $\text{Sc}$  increases further, however, and the tracer diffusivity is smaller still compared to momentum diffusivity, tracer trapping within the large central vortex – and even the smaller vortex near the exit of the inflow channel – is now so pronounced that for  $\text{Re} = 40$  and  $\text{Re} = 50$  which have comparatively faster racetrack speeds than the lower Reynolds numbers, washout is slower, as these flow fields also contain the largest central vortices.

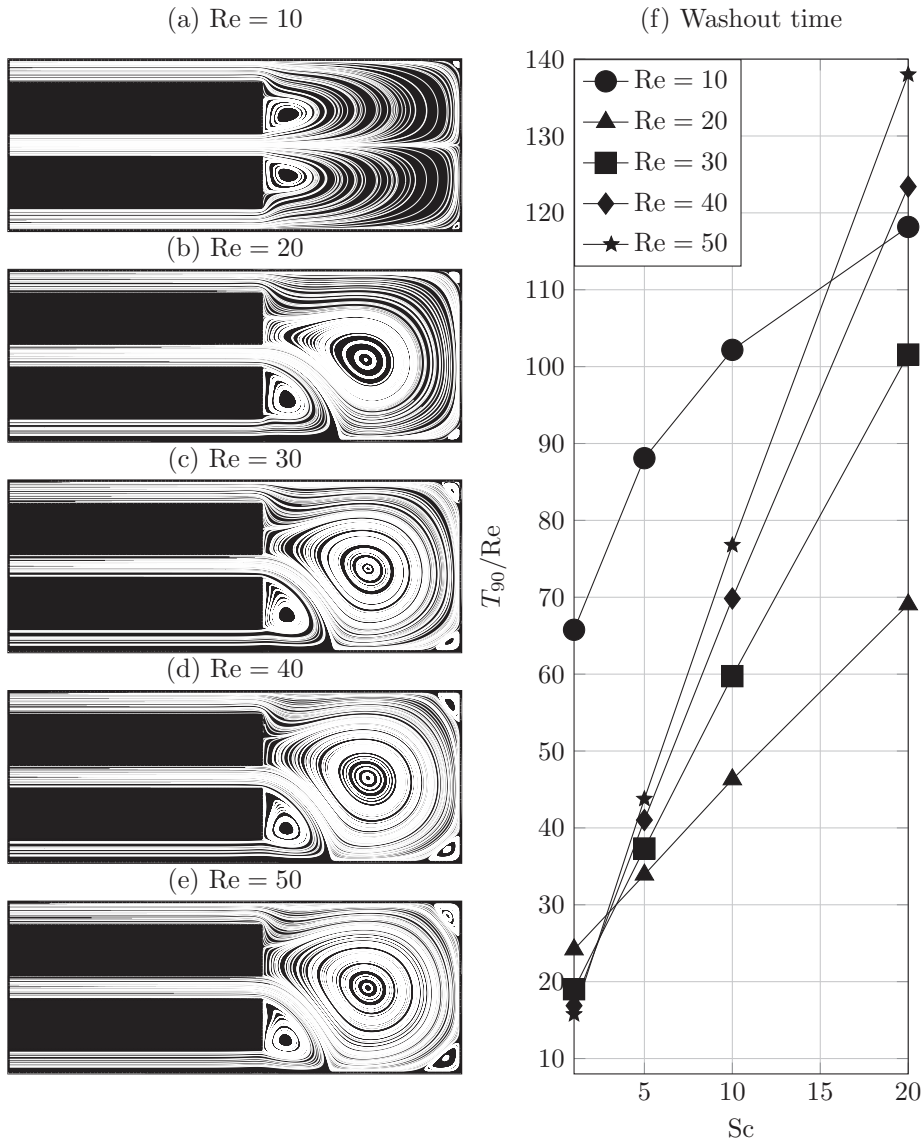


Figure 3: Washout time for initial shape. Different symbols represent different Reynolds numbers with associated streamlines shown in Figures a-e.

### 3.2. Washout time after vortex optimisation

In Figure 4a we plot the objective, equation (2.12) – normalised with respect to the objective value at iteration zero – as a function of iteration number. At each value of  $Re$  considered, we observe a broadly monotonic decrease in the objective as a function of iteration number. The optimisation for  $Re = 20$  and  $Re = 30$  terminates after only 21 and 50 iterations, respectively, caused by the flow problem switching solution branches, and we explain this behaviour in more detail in Section 3.3. We see that for  $Re = 10$ ,  $Re = 40$ , and  $Re = 50$  the iterative process continues until the objective curve flattens. In Figure 4b, we consider the relative washout time for  $Sc = 5$ , defined as the dimensionless washout

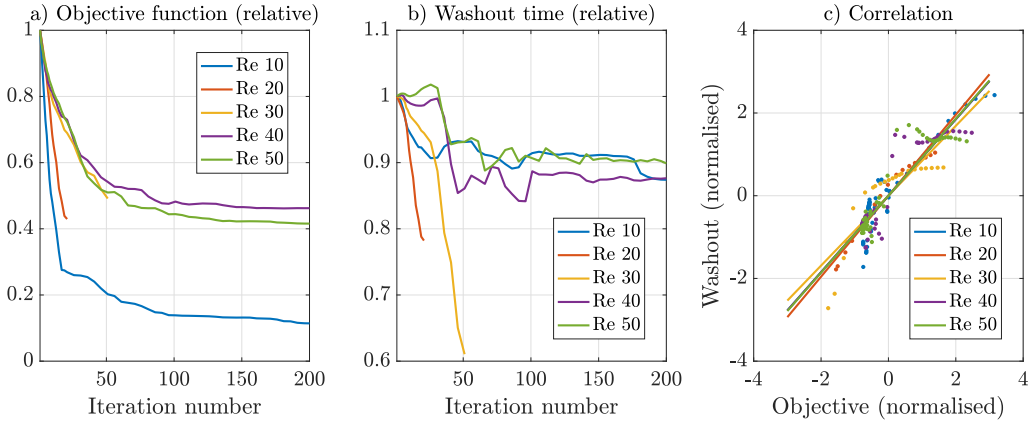


Figure 4: Left and Middle: Objective value and washout time ( $Sc = 5$ ) as a function of iteration number. Right: Normalised (to zero mean and unit standard deviation) objective value and washout and correlation. The correlation is 0.9245, 0.9786, 0.8449, 0.925, 0.9291 for  $Re \in \{10, 20, 30, 40, 50\}$ .

time (given by equation (2.8)) normalised with respect to the dimensionless washout time at iteration zero. We plot this relative washout time as a function of iteration number. For  $Re = 20$  and  $Re = 30$  washout time decreases monotonically with iteration number, correlating with the objective. For  $Re = 10$ ,  $Re = 40$ , and  $Re = 50$ , although the behaviour is not strictly monotonic, the washout time has reduced to approximately 90% of the initial value by iteration 200. Important to consider is that Figure 4b only indicates relative improvement of washout time for each value of  $Re$ , and not a comparison of washout time for different Reynolds numbers. We also note that the washout time in Figure 4b is only for a single Schmidt number –  $Sc = 5$  – and that correlation between objective and washout may improve with increasing  $Sc$ , although the long washout times with increasing Schmidt number render this comparison a computationally prohibitive task. We quantify the correlation between objective function and washout time in Figure 4c by plotting these metrics against each other (normalised to zero mean and unit standard deviation for each value of  $Re$ ). We observe that the data follows a roughly linear trend with correlation coefficients 0.9245, 0.9786, 0.8449, 0.925, 0.9291 for  $Re \in \{10, 20, 30, 40, 50\}$ .

In Figures 5a-e we show the channel shapes and associated streamlines after the optimisation routine. We emphasise that the optimised channel shape is different for each  $Re$  considered. For Reynolds numbers  $Re \in [20, 50]$  the optimised inflow channel is now curved, modifying the flow structure observed in the cavity region and decreasing the size of the central vortex. For  $Re = 10$  domain symmetry – and hence flow symmetry – have been broken, resulting in slightly reduced vortices near the channel entrance and a faster velocity near the back of the cavity compared to the initial shape (compare Figure 5a with Figure 3a). It is also worth noting the appearance of a small triangular indent in the bottom outflow of the optimised shapes for  $Re = 40$  and  $Re = 50$ . This is purely an artefact of the discretisation; although this feature plays a role in decreasing the objective function, likely due to singularities in  $\nabla \mathbf{u}$  at the outflow corners, its influence on  $T_{90}$  is negligible. We confirmed that the difference between  $T_{90}$  for  $Re = 50$  with and without this sharp geometric feature is only  $\mathcal{O} = 10^{-5}$ . To compare the washout times for the optimised shapes to the washout times for the initial shapes, in Figure 5f we plot  $T_{90}/Re$  as a function of  $Sc$  for the final channel geometries and flow patterns produced by the

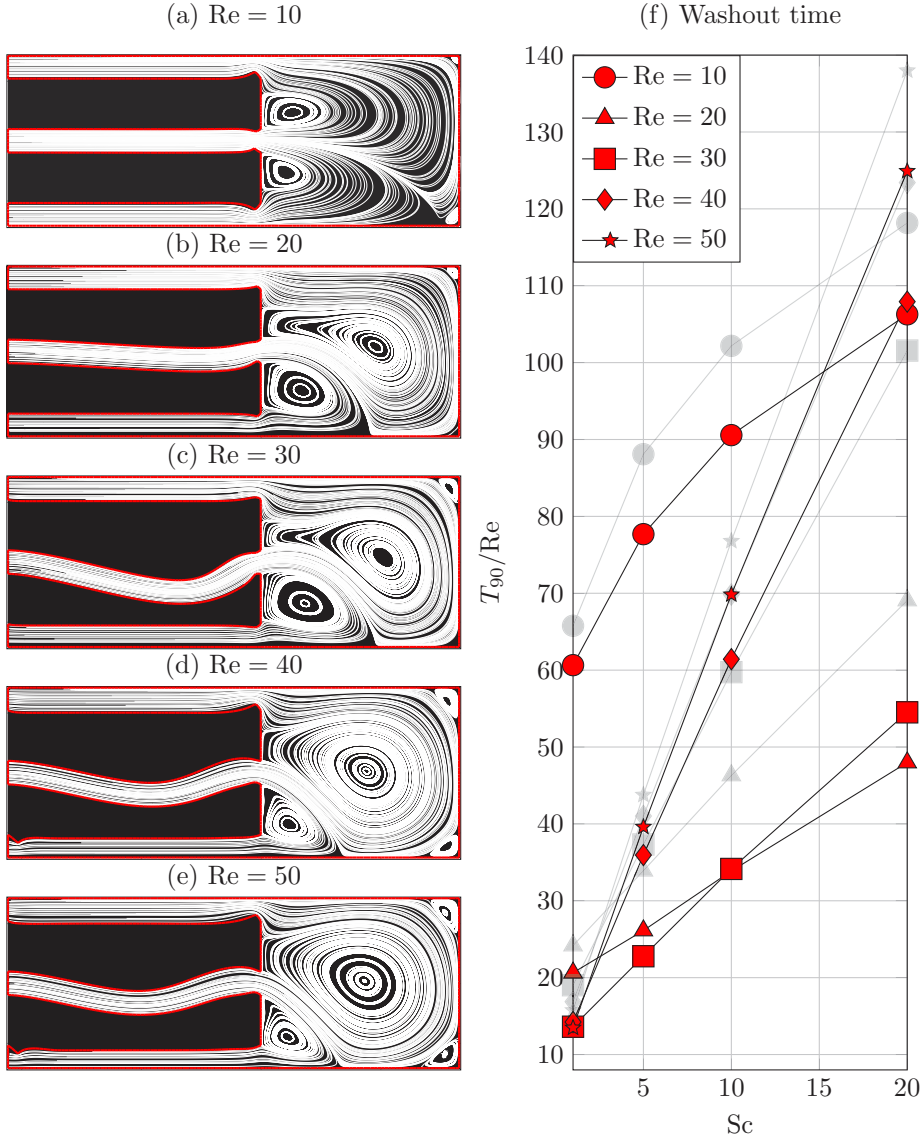


Figure 5: Washout times for final shapes after optimisation. Different symbols represent different Reynolds numbers with associated streamlines shown in Figures a-e. Red filled symbols correspond to values for optimised geometries, pale grey symbols correspond to values for initial shapes.

optimisation routine for each value of  $Re$ . For each  $Re$ , the corresponding washout times in Figure 5f have been reduced compared to the washout times for the initial shape at each value of  $Sc$  (the washout times for the initial shapes are reproduced from Figure 3 in grey). As with the initial shapes, crossings of the lines occur as tracer diffusion decreases, corresponding to the increasing significance – and competing effects – of flow structure and flow speed. Important to observe is that the difference between washout times for the initial and final shape for given Reynolds number increases with increasing Schmidt



number. This is to be expected: as  $Sc$  increases (diffusion decreases) vortex size becomes increasingly significant.

### 3.3. The effect of multiple solutions

At each iteration of the shape optimisation algorithm described in Section 2.4.2 there is an associated channel geometry  $\Omega_t^{(k)}$  and velocity and pressure solution  $(\mathbf{u}, p)^{(k)}$ . For critical Reynolds number  $Re_c$ , if  $Re < Re_c$ , then  $(\mathbf{u}, p)^{(k)}$  is unique and stable, whereas if  $Re > Re_c$ , then  $(\mathbf{u}, p)^{(k)}$  is one of three possible solutions, two of which are stable, and therefore physically relevant, and one of which is unstable.

For example, for  $\Omega_t^{(0)}$ , there is a pitchfork bifurcation – displayed in blue in Figure 6 – and  $Re_c \approx 12$ . For the initial solve with  $\Omega_t^{(0)}$  when  $Re > Re_c$ , we select one of the mirror asymmetric solution branches using the numerical method described at the end of Section 2.4.1 and used in Williams *et al.* (2020). Without loss of generality, we select for the solution on the upper branch, as indicated for  $Re = 20$  in Figure 6 by the blue circle, where we have defined the bifurcation functional as the flux difference between the lower and upper outflows, i.e.,

$$\Delta Q = \int_{\Gamma_{out,l}} \mathbf{u} \cdot \mathbf{n} \, ds - \int_{\Gamma_{out,u}} \mathbf{u} \cdot \mathbf{n} \, ds, \quad (3.1)$$

where  $\mathbf{n} = (-1, 0)^\top$  is the outward unit normal vector to  $\Gamma_{out,l}$ ,  $\Gamma_{out,u}$ .

With each shape iteration, the bifurcation structure, and hence the value of  $Re_c$  changes. If the domain deforms asymmetrically the once perfect pitchfork solution structure becomes a broken pitchfork, as demonstrated by the yellow and red lines in Figure 6, where the  $Re_c$  which characterises the transition point from one to three solutions is the location of the detached limit point.

This insight allows us to understand the branch-switching that causes the early termination of the  $Re = 20$  and  $Re = 30$  optimisations (red and yellow lines in Figure 4a, respectively). In Figure 6a we plot the broken pitchfork bifurcation diagrams associated with  $\Omega_t^{(11)}$  and  $\Omega_t^{(21)}$  in yellow and red, respectively, with corresponding coloured circles denoting the solutions  $(\mathbf{u}, p)^{(11)}$  and  $(\mathbf{u}, p)^{(21)}$  at  $Re = 20$  and solid and dashed lines indicating stable and unstable solution branches respectively. The black triangle, cross, and square on the red bifurcation diagram associated with  $\Omega_t^{(21)}$  denote the three possible solutions at  $Re = 20$ , and their corresponding streamlines are drawn in Figures 6b-d. As iteration number increases from 0 to 21,  $Re_c$  increases and approaches  $Re = 20$ . Iteration 22 then produces a candidate channel shape with  $Re_c > 20$ , and hence only a single solution (on the lower stable branch) exists at  $Re = 20$ . At this stage the Navier–Stokes solver diverges, as the chosen initial guess – Solution 1 in Figure 6 – is far away from the solution associated with the lower asymmetric branch, and hence is a poor initial guess for the Newton solver.

Of interest to note is that starting the optimisation routine with  $\Omega_t^{(21)}$  and the solution for  $(\mathbf{u}, p)^{(21)}$  given by the lower stable branch (Solution 3 in Figure 6) instead of the upper stable branch (Solution 1), results in a channel shape similar to that in Figures 6b-d, but flipped with respect to the  $x$ -axis of the domain. Notably, after several iterations of the optimisation routine the broken and unbroken branches of the pitchfork collide and switch stability – an example of a cusp catastrophe – so we are again on the stable, broken branch of the bifurcation diagram. After several more iterations, a candidate shape with  $Re_c > 20$  is proposed by the optimisation algorithm, and the Navier–Stokes solver diverges again.

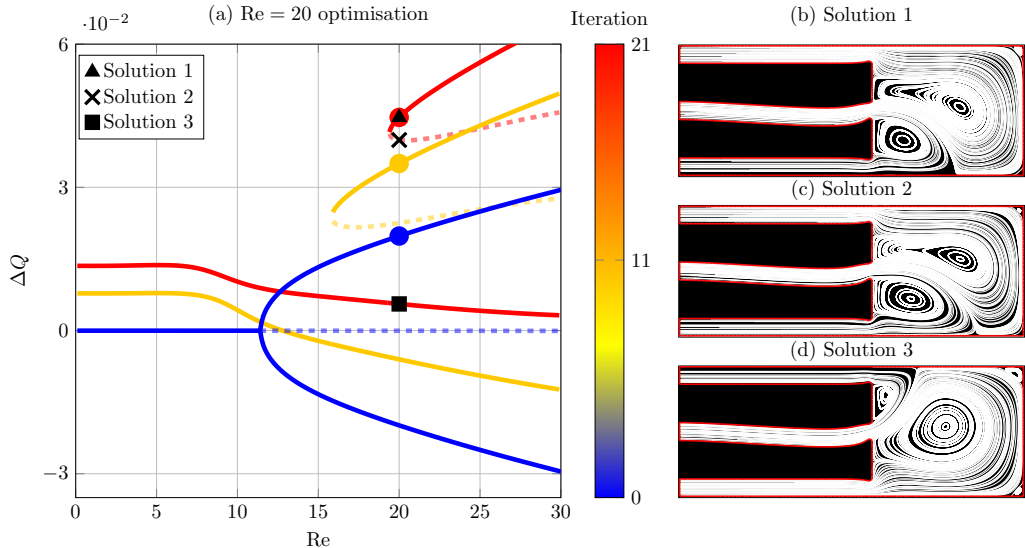


Figure 6: Bifurcation diagrams for  $Re = 20$  at shape iterations 0, 11, and 21. Dashed lines indicate unstable solutions and solid lines stable (physically attainable) ones. The coloured circles indicate the relevant solution during optimisation. The three black symbols indicate the three possible solutions at  $Re = 20$  and iteration 21, with corresponding streamlines presented in Figures b-d.

Of the five Reynolds numbers we explored, the branch-switching behaviour during the initial optimisation run occurs in only  $Re = 20$  and  $Re = 30$  at iterations 21 and 50, respectively (see Figure 4). For higher Reynolds numbers – e.g.  $Re = 50$  – the bifurcation diagram never evolves to the stage where  $Re_c > 50$  as would be required to observe the phenomenon described in this section.

Computing bifurcation diagrams for the optimised channel shape highlights an important point: although the optimised flow solution is stable, for some Reynolds numbers it is not the solution you would find by increasing flow rate in physical practice (or by numerical Reynolds number continuation in computation). To obtain the optimal flow field associated with the optimal shape in reality, techniques such as blocking and unblocking an outflow – previously used to obtain flow fields on a desired branch of a pitchfork bifurcation in experiment (Williams *et al.* 2020) – may be required.

#### 4. Summary and discussion

Many biomedical procedures require efficient clearance of debris, whether to improve visualisation in endoscopic surgeries, or to reduce accumulation of undesirable biological agents in a stented artery or ureter. Debris clearance typically relies on fluid flow – an injected saline solution, contrast agent, or blood circulation – and debris aggregation sites can arise as a result of recirculation zones in the fluid. Modelling debris dynamics as the advection and diffusion of a passive tracer, the time required for the tracer to escape a region of closed streamlines depends non-linearly on the velocity, diffusive properties of the debris, and vortex size (Rhines & Young 1983). This motivated the hypothesis that underpinned the work in this article: *for fixed flow speed and debris characteristics, a reduction in vortex size will promote more efficient debris clearance.*

In this paper we considered flow in a two-dimensional Cartesian geometry, consisting of a rectangular cavity and inlet and outlet channels, and used shape optimisation to determine channel geometries that reduce the presence of vortices. We characterised recirculation zones via  $\det(\nabla \mathbf{u}) > 0$  – i.e. locally circular or spiral streamlines – thereby choosing the objective for our optimisation to be a smoothed form of  $\max(0, \det \nabla \mathbf{u})$  integrated over the cavity domain (Kasumba & Kunisch 2012). We determined the relationship between washout time and the objective by solving an advection-diffusion equation for the motion of a passive tracer in both the initial and optimised flow fields and calculating the resulting washout time. We used  $T_{90}$  as the metric characterising washout time, corresponding to the time taken for 90% of the tracer to exit the domain. We note here that alternatives to  $T_{90}$  for the washout metric can be found in the literature: for example, Min *et al.* (2020) track the location and magnitude of the peak of the remaining concentration. A comparison of washout metrics in relation to our objective function would be an interesting future study, but here we chose to restrict attention to  $T_{90}$  as this is a commonly used standard measure of washout efficiency.

We first solved for washout time in the initial uniform channel geometry for Reynolds numbers  $Re = 10, 20, 30, 40$ , and  $50$  and confirmed that washout time depends non-trivially on Reynolds number, flow structure, and the diffusion coefficient associated with the passive tracer. When Schmidt number  $Sc = 1$  – tracer motion is governed by a balance of momentum diffusivity of the fluid and diffusivity of the concentration – washout time decreases monotonically with increasing Reynolds number. However, for decreased tracer diffusivity, we found that the case of fastest flow ( $Re = 50$ ) in fact provided the slowest washout time. This seemingly counterintuitive result is easily explained by the fact that the fastest flow also creates the largest vortex, and in an advection-dominated regime vortex size has a dominant impact on washout time.

Running the optimisation algorithm for each of the five Reynolds numbers, we computed the optimal shape and found that washout time is indeed lower in the optimised geometries compared to the initial geometries in each case. This partially confirmed the hypothesis that our objective serves as a proxy for washout time. To explore the correlation more thoroughly, we computed washout time for several geometries (and corresponding flow fields) along the optimisation pathway. With  $Re = 20$  and  $Re = 30$ , washout time decreases monotonically with each shape iteration. For  $Re = 10, 40$ , and  $50$ , washout is not monotonic with iteration number in the optimisation algorithm, although after 200 iterations the washout time is reduced to approximately 90% of its initial value. This strong, if imperfect, correlation suggests that minimising washout may be well accomplished via a proxy objective measuring the size of vortical structures, a choice that creates significant computational advantage. This reflects the intuitive notion that if tracer transport is not diffusion dominated, tracer can remain ‘stuck’ in vortices for long periods of time, and thus reducing vortical structures can significantly reduce washout time. By construction,  $Sc$  does not appear in the optimisation algorithm, but the connection between vortex size and washout suggests that the optimal shape for a given  $Re$  will not vary significantly with varying  $Sc$ , at least in the advection-dominated regime. The connection between washout reduction and the size of vortical structures is less strong in the high-diffusion regime; here maximising flow in the non-vortical ‘racetrack’ regions may be more important than minimising the size of vortices, and a different proxy metric may be more suitable. However, such a regime is less interesting from an optimisation point of view, as it is characterised by very low washout times in any case.

Also of note on the computational side is the fact that for two of the Reynolds numbers ( $Re = 20$  and  $Re = 30$ ), solution branch-switching occurred during the optimisation routine. This feature was shown to occur as an artefact of the evolving




Figure 7: Optimisation on different geometries. The top row shows the flow pattern and computed  $T_{90}$  before optimisation; the bottom row shows the same after optimisation. All computations are for  $Re = 20$ ,  $Sc = 20$ .

broken pitchfork bifurcation structure for the flow problem during the evolution of the channel shape following the optimisation pathway. Such a feature may naturally arise in an optimisation scheme involving multi-solution flow fields, and highlights an important point, particularly when translating to a clinical or industrial setting: the flow field associated with a computed optimal shape may not be the only stable solution, thus care may be needed in physically realising the desired optimal flow.

We also explored the washout time corresponding to the optimised channel shapes as a function of the diffusion of the passive tracer, captured by Schmidt number  $Sc$ . As was observed in the base geometry, when plotting washout against  $Sc$  across the set of Reynolds numbers, numerous crossings of the lines are observed. This further underscores the complex relation between washout time, flow speed, vortices, and diffusivity. Of course, from a device design point of view, the key question is *what shape should the inflow and outflow channels be?*. Here, in lieu of a single answer, our analysis provides several points of insight: we have shown that the optimal shape is somewhat dependent on the Reynolds number, and the optimal Reynolds number is dependent on the diffusivity. The diffusivity may not be well known, but we see the importance of a good estimate. In terms of Reynolds number, in practice this can be controlled by changing the upstream driving pressure or fluid properties, in which case Figure 5 (or its equivalent in a different geometry) enables to determine the best combination of shape and Reynolds number. A natural extension would be to optimise for both the Reynolds number and the shape simultaneously. However, in addition to non-trivial extensions of the Fireshape library being necessary for such a combined optimisation, we also expect more difficulties navigating the multiple solution branches if the Reynolds number changes during the optimisation. For these reasons, we do not explore this combined optimisation at this stage.

Even in the absence of a global optimal shape, Figure 5 shows a clear qualitative trend that curving the jet of fluid via bending the inlet channel leads to a decrease in the size of vortical structures, and we anticipate that this qualitative finding could carry over to more complex (and possibly 3D) domain geometries. To test this idea, in Figure 7 we have run the optimization algorithm on a set of three distinct geometries: in Figure 7(a) and (d) we replace the square domain with a smooth curved boundary; in (b) and (e) we replace the back edge with an inwardly-curved boundary; and (c) and

(e) break the symmetry of the domain through the addition of a circular bulge in the corner. Here, the top row shows the flow pattern and computed  $T_{90}$  for the straight inlet channel, while the bottom row shows the optimised shape and corresponding  $T_{90}$ . For each domain, the total cavity area was preserved, allowing direct comparison of the washout times. Unsurprisingly, the smooth domain in (a) and (d) exhibits a much faster washout than the domain in (c) and (e), in which debris remains stuck in the small vortex in the additional ‘mini-cavity’. Nevertheless, in all cases we observe a significant reduction in washout time. Perhaps more importantly, the optimised shape for each of the domains is qualitatively very similar, following the same trend of curving the tip of the inlet channel as was found in Figure 5. While this is clearly not exhaustive in terms of possible geometric features, it does demonstrate a robustness to the framework and reinforces a design principle that curving the fluid can minimize the size of vortical structures and thus reduce the recirculation of debris.

An obvious extension to the study reported in this manuscript is to consider three-dimensional – rather than two-dimensional – geometries representative of real medical devices. However, a salient factor to consider is computational cost: solving shape optimisation problems involves many ( $\sim 200$ ) state and adjoint equation solves. In two dimensions this is possible on a standard workstation, but the computational cost is several orders of magnitude larger in three dimensions. For this reason we focused on a 2D case in this study, but are cognizant of the significance of the extension to 3D. We note that the choice of vortex metric is more open-ended in 3D, and a thorough comparison of results using different available metrics remains necessary future work (Jeong & Hussain 1995; Haller 2005).

Finally, we note that in the optimisation scheme, the choice of which boundaries are fixed and which are free is important. In our set-up, we fixed the walls at the front of the scope, thereby preventing the inflow channel from flaring or tapering as it enters the cavity. Relaxing this restriction potentially introduces competing effects: a flared tip geometry may reduce the size of vortices, but it also reduces the flow speed of the open streamlines, or ‘racetrack’, as described in Section 3.1. For high  $Sc$  (tracer with low diffusivity), debris can get stuck in areas of slow velocity, inducing longer washout times, even though the presence of vortices has been reduced. In this instance, an alternative objective that promotes fast velocity in non-vortex regions – the racetrack – might be suitable, e.g. one could subtract a term from the objective of the form  $\int_{\Omega_c} |\mathbf{u}| \mathcal{H}(-\det \nabla \mathbf{u}) d\Omega_c$ . However, numerical complexities arise with determining a smoothed version of the indicator function  $\mathcal{H}$ , and further work is required to implement such a modified metric in practice.

#### *Code availability*

For reproducibility we have archived the code and the exact software versions used to produce the results in this paper on Zenodo (Zenodo 2020).

This publication is based on work supported (or partially supported) by the EPSRC Doctoral Prize Scheme (EP/R513295/1). The authors would like to acknowledge the support of Boston Scientific Corporation. The authors are grateful for discussions with Patrick Farrell and Simon Tavener on using defcon to compute the bifurcation diagrams. Declaration of Interests. The authors report no conflict of interest.

#### REFERENCES

- ABERGEL, F. & TEMAM, R. 1990 On some control problems in fluid mechanics. *Theoret. Comput. Fluid Dynamics* **1** (6), 303–325.
- ALNÆS, M. S., LOGG, A., ØLGAARD, K. B., ROGNES, M. E. & WELLS, G. N. 2014 Unified Form Language: A domain-specific language for weak formulations of partial differential equations. *ACM Transactions on Mathematical Software* **40** (2), 9:1–9:37.
- AMESTOY, P.R., DUFF, I.S., L’EXCELLENT, J.-Y. & KOSTER, J. 2000 MUMPS: a general purpose distributed memory sparse solver. In *International Workshop on Applied Parallel Computing*, pp. 121–130. Springer.
- BALAY, S., ABHYANKAR, S., ADAMS, M.F., BROWN, J., BRUNE, P., BUSCHELMAN, K., DALCIN, L., DENER, A., EIJKHOUT, V., GROPP, W.D., KAUSHIK, D., KNEPLEY, M.G., MAY, D.A., MCINNES, L.C., MILLS, R.T., MUNSON, T., RUPP, K., SANAN, P., SMITH, B.F., ZAMPINI, S., ZHANG, H. & ZHANG, H. 2018 PETSc users manual. *Tech. Rep.* ANL-95/11 - Revision 3.10. Argonne National Laboratory.
- BALAY, S., ABHYANKAR, S., ADAMS, M.F., BROWN, J., BRUNE, P., BUSCHELMAN, K., DALCIN, L., EIJKHOUT, V., GROPP, W.D., KAUSHIK, D., KNEPLEY, M.G., MCINNES, L.C., RUPP, K., SMITH, B.F., ZAMPINI, S., ZHANG, H. & ZHANG, H. 2016 PETSc Web page.
- BALAY, S., GROPP, W.D., MCINNES, L.C. & SMITH, B.F. 1997 Efficient management of parallelism in object oriented numerical software libraries. In *Modern Software Tools in Scientific Computing*, pp. 163–202. Birkhäuser Press.
- BRENNER, S.C. & SUNG, L.-Y. 2005  $C^0$  interior penalty methods for fourth order elliptic boundary value problems on polygonal domains. *J. Sci. Comput.* **22–23** (1-3), 83–118.
- CLAVICA, F., ZHAO, X., ELMAHDY, M., DRAKE, M.J., ZHANG, X. & CARUGO, D. 2014 Investigating the flow dynamics in the obstructed and stented ureter by means of a biomimetic artificial model. *PLoS ONE* **9** (2), e87433.
- COPPOLA, G. & CARO, C. 2009 Arterial geometry, flow pattern, wall shear and mass transport: potential physiological significance. *J. Roy. Soc. Int.* **6**, 519–528.
- DALCIN, L.D., PAZ, R.R., KLER, P.A. & A.COSIMO 2011 Parallel distributed computing using Python. *Adv. Water Resour.* **34** (9), 1124–1139, new Computational Methods and Software Tools.
- DELFOUR, M.C. & ZOLÉSIO, J.-P. 2011 *Shapes and Geometries: Metrics, Analysis, Differential Calculus, and Optimization*. Society for Industrial and Applied Mathematics.
- FARRELL, P.E., BIRKISSON, Á. & FUNKE, S.W. 2015 Deflation techniques for finding distinct solutions of nonlinear partial differential equations. *SIAM J. Sci. Comput.* **37** (4), A2026–A2045.
- FRANCA, LEOPOLDO P., FREY, SERGIO L. & HUGHES, THOMAS J.R. 1992 Stabilized finite element methods: I. Application to the advective-diffusive model. *Comput. Methods in Appl. Mech. Eng.* **95** (2), 253–276.
- GEUZAIN, C. & REMACLE, J.-F. 2009 Gmsh: A 3-D finite element mesh generator with built-in pre- and post-processing facilities. *Int. J. Numer. Meth. Eng.* **79** (11), 1309–1331.
- GUZMN, JOHNNY & SCOTT, L. RIDGWAY 2018 The Scott-Vogelius finite elements revisited. *Mathematics of Computation* **88** (316), 515–529.
- HALLER, G. 2005 An objective definition of a vortex. *J. Fluid Mech.* **525**, 1–26.
- HAM, D.A., MITCHELL, L., PAGANINI, A. & WECHSUNG, F. 2019 Automated shape differentiation in the Unified Form Language. *Structural and Multidisciplinary Optimization* Accepted.
- HOMOLYA, M., KIRBY, R.C. & HAM, D.A. 2017a Exposing and exploiting structure: optimal code generation for high-order finite element methods.
- HOMOLYA, M., MITCHELL, L., LUPORINI, F. & HAM, D.A. 2017b TSFC: a structure-preserving form compiler.
- JEONG, J. & HUSSAIN, F. 1995 On the identification of a vortex. *J. Fluid Mech.* **285**, 69–94.
- JIMÉNEZ, J.M. & DAVIES, P.F. 2009 Hemodynamically driven stent strut design. *Ann. Biomed. Eng.* **37** (8), 1–24, arXiv: NIHMS150003.
- JOHN, V., LINKE, A., MERDON, C., NEILAN, M. & REBHOLZ, L. G. 2017 On the divergence constraint in mixed finite element methods for incompressible flows. *SIAM Review* **59** (3), 492–544.
- KASUMBA, H. & KUNISCH, K. 2010 Shape design optimization for viscous flows in a channel

- with a bump and an obstacle. *2010 15th International Conference on Methods and Models in Automation and Robotics, MMAR 2010* pp. 284–289.
- KASUMBA, H. & KUNISCH, K. 2012 Vortex control in channel flows using translational invariant cost functionals. *Comput. Optim. Appl.* **52**, 691–717.
- LUPORINI, F., HAM, D.A. & KELLY, P.H.J. 2017 An algorithm for the optimization of finite element integration loops. *ACM T. Math. Software* **44**, 3:1–3:26.
- MICHELETTI, A.M. 1972 Metrica per famiglie di domini limitati a proprietà generiche degli autovalori. *Annali della Scuola Normale Superiore di Pisa* **26** (3), 683–694.
- MIN, D., FISCHER, P.F. & PEARLSTEIN, A.J. 2020 High schmidt number ”washout” of a soluble contaminant downstream of a backward facing step. *Int. J. Heat Mass Trans.* **159**, 119740.
- MURAT, F. & SIMON, J. 1976 Etude de problèmes d’optimal design. In *Optimization Techniques Modeling and Optimization in the Service of Man Part 2*, pp. 54–62. Springer Berlin Heidelberg.
- NOCEDAL, J. & WRIGHT, S.J. 2006 *Numerical Optimization*. Berlin: Springer.
- PAGANINI, A. & WECHSUNG, F. 2020 Fireshape: a shape optimization toolbox for Firedrake. *In preparation* .
- PASTER, A., SHLAIN, S., BARGHOUTHY, Y., LIBERZON, A., AVIRAM, G. & SOFER, M. 2019 Assessing the Influence of Irrigation Flows on Clearance of Calculi Fragments During Percutaneous Nephrolithotomy: A Numerical and Physical Model Study. *Urology* **124**, 46–51.
- QIN, J. 1994 On the convergence of some low order mixed finite elements for incompressible fluids. PhD thesis, Pennsylvania State University.
- RATHGEER, F., HAM, D.A., MITCHELL, L., LANGE, M., LUPORINI, F., MCRAE, A.T.T., BERCEA, G.-T., MARKALL, G.R. & KELLY, P.H.J. 2016 Firedrake: automating the finite element method by composing abstractions. *ACM Trans. Math. Softw.* **43** (3), 24:1–24:27.
- RHINES, B. & YOUNG, W.R. 1983 How rapidly is a passive scalar mixed within closed streamlines? *J. Fluid Mech.* **133**, 133–145.
- RIDZAL, D. & KOURI, D. 2014 Rapid Optimization Library.
- SILVESTER, D., ELMAN, H. & WATHEN, A. 2014 *Finite Elements and Fast Iterative Solvers: with applications in incompressible fluid dynamics*. Oxford University Press.
- SIMON, J. 1980 Differentiation with respect to the domain in boundary value problems. *Numer. Func. Anal. Opt.* **2** (7-8), 649–687.
- WECHSUNG, F. 2019 Shape optimisation and robust solvers for incompressible flow. PhD thesis, University of Oxford.
- WELLS, G.N., KUHLE, E. & GARikipati, K. 2006 A discontinuous Galerkin method for the Cahn–Hilliard equation. *J. Comput. Phys.* **218** (2), 860–877.
- WILLIAMS, J.G., CASTREJON-PITA, A.A., TURNER, B.W., FARRELL, P.E., TAVENER, S.J., MOULTON, D.E. & WATERS, S.L. 2020 Cavity flow characteristics and applications to kidney stone removal. *J. Fluid Mech.* **902**, A16.
- XHANG, P.H., TKATCH, C., NEWMAN, R., GRIMME, W., VAINCHTEIN, D. & KRESH, J.Y. 2019 The mechanics of spiral flow: Enhanced washout and transport. *Art. Organs.* **43**, 1144–1153.
- Zenodo 2020 Software used in ‘Shape optimisation for faster washout in recirculating flows’, <https://doi.org/10.5281/zenodo.4075330>.