

Shape optimisation of stirring rods for mixing binary fluids

MAXIMILIAN F. EGGL

Department of Mathematics, Imperial College London, London SW7 2AZ, U.K.

*Corresponding author: maximilian.eggl11@imperial.ac.uk

PETER J. SCHMID

Department of Mathematics, Imperial College London, London SW7 2AZ, U.K.

[Received on 18 November 2019]

Mixing is an omnipresent process in a wide-range of industrial applications, which supports scientific efforts to devise techniques for optimising mixing processes under time and energy constraints. In this endeavor, we present a computational framework based on nonlinear direct-adjoint looping for the enhancement of mixing efficiency in a binary fluid system. The governing equations consist of the nonlinear Navier-Stokes equations, complemented by an evolution equation for a passive scalar. Immersed and moving stirrers are treated by a Brinkman-penalisation technique, and the full system of equations is solved using a Fourier-based pseudospectral approach. The adjoint equations provide gradient and sensitivity information which is in turn used to improve an initial mixing strategy, based on shape, rotational and path modifications. We utilise a Fourier-based approach for parameterising and optimising the embedded stirrers and consider a variety of geometries to achieve enhanced mixing efficiency. We consider a restricted optimisation space by limiting the time for mixing and the rotational velocities of all stirrers. In all cases, non-intuitive shapes are found which produce significantly enhanced mixing efficiency.

Keywords: Mixing, Optimization, Penalization, Adjoint method, Fourier-based Shape Parameterisation

2000 Math Subject Classification: 34K30, 35K57, 35Q80, 92D25

1. Introduction

The mixing of binary fluids is a problem of fundamental concern in fluid dynamics, as its underlying mechanisms play an important role in a wide variety of industrial fields and in many fluid processes encountered in daily life. This provides ample motivation to study the processes and characteristics that drive this natural phenomenon.

The study of mixing can broadly be split into two areas, theoretical and application-driven. The theoretical aspects of mixing have been investigated in great detail, and there exists a vast body of literature pertaining to many aspects of mixing for binary fluids. Issues that have been studied include the formulation of emulsions and the stirring mechanics used to create them [Taylor, 1934], mixing in stratified flows [Long, 1978, Peltier & Caulfield, 2003] as well as the optimisation of mixing [Marcotte & Caulfield, 2018], among many other efforts. For example, Stokes mixing, i.e., the mixing of highly viscous fluids, has drawn great interest in the research community, exemplified by investigations [Spencer & Wiley, 1951] that identified stretching, cutting and folding as the primary mechanisms of mixing. Chien et al. [1986], in a related effort, found that chaotic mixing can be induced from laminar mixing by varying the oscillations of the walls, a configuration that was further studied by Ottino et al. [1988] and Mohr et al. [1957] who argued for a simplified analysis of viscous mixing by ignoring interfacial tension. Their simplification formed the basis for many subsequent developments in the mixing of viscous fluids. Concurrently, at the other end of

the parametric range, turbulent mixing, i.e., the mixing of binary fluids mostly by inertial effects, has been studied extensively, starting with the early articles by Corrsin [1957], Corrsin & Karweit [1969] and Batchelor [1952]. More recently, mixing has been linked to chaotic processes and chaotic advection, yielding the accepted consensus that turbulence naturally leads to ‘better’ mixing [Dimotakis & Catrakis, 1999]. This fundamental finding was subsequently the topic of many more studies and has been verified experimentally, e.g. in Dimotakis [2000]. As a consequence, it has also spawned a trend where efficient mixing is achieved by techniques that inject energy into the system so as to generate turbulence, which subsequently produces better mixing [Aamo & Krstic, 2013, Jalali et al., 2015]. Turbulence is also a key component in studies of mixing in stratified fluids, e.g., Linden [1979]; these investigations suggest that turbulence exists for all values of the ratio of buoyancy to flow shear and argue for the existence of an overall stratification where mixing reaches an optimum. This general configuration has also been studied by Tang et al. [2009], who put a value on this optimal stratification and confirmed Linden’s supposition.

The above selection of articles, while certainly incomplete and not exhaustive, is intended to demonstrate the breadth and depth of scientific efforts that have been devoted to mixing. For readers who seek a more in-depth review of the different aspects of mixing theory, the monographs by Uhl [2012], Paul et al. [2004] and Ottino [1989] provide a more complete background on the rich aspects of fluid mixing.

The above monographs and investigations focus on the more theoretical aspects of fluid mixing, but are less representative of studies concerned with the applications of mixing in natural or technological settings. Samples from the latter category examine the spreading and mixing of pollutants in the ocean [Lekien et al., 2005], the ventilation of buildings [Hunt & Linden, 1999, Linden, 1999], the mixing of air and fuel for subsequent combustion [Annaswamy & Ghoniem, 1995] or the mixing in microfluidic devices [Hessel et al., 2005, Nguyen & Wu, 2005]. Additional fields of application include the food processing, pharmaceutical and consumer-product industry. Interest in these industrial application fields is fueled by the fact that even modest improvements of mixing efficiency would translate into immediate and significant profits as well as into a more consistent quality of the respective final product.

Mixing in industrial applications is often accomplished by stirrers, i.e., rotating bodies of a given shape embedded in a mixing vessel, whose task it is to produce long filaments [Aref, 1984] which subsequently diffuse. This two-step shearing-diffusion process, which is principally at the core of binary fluid mixing by stirrers, has been recognised as a fundamental mechanism in inertia-driven mixing [Spencer & Wiley, 1951] and has been studied extensively to gain further insight and to guide control strategies. The present article follows in this path and considers stirrer-induced inertia-driven mixing coupled with a shape optimisation of the stirrers.

From a computational point of view, several difficulties arise from the treatment of the embedded stirrers, in particular the treatment of the fluid-structure interaction between these moving bodies and the binary fluid. Various options exist and have been pursued by previous studies, such as body-fitted meshes [Mattiussi, 2000] or the method of fictitious domains [Glowinski et al., 1996]. These computational techniques are, however, often restricted to simple configurations or are excessively costly due to the need to remesh at each time step (after the solid body has moved). An attractive alternative to these methods is the Brinkman-penalisation approach which will be adopted in this paper and will provide the necessary flexibility and efficiency to treat more complex geometries and flow configurations. Introduced by Arquis & Caltagirone [1984], solutions obtained by the penalisation method have been rigorously shown to converge to the corresponding solutions of the Navier-Stokes equations [Angot et al., 1999] for the respective complex domain. Supporting studies, including an asymptotic analysis of this convergence process, are summarised in Liu & Vasilyev [2007], and applications to high Mach-number flows, e.g. Boiron et al. [2009], and turbulent flow past cylinders, e.g. Kevlahan & Ghidaglia [2001], among many other examples, have demonstrated the effectiveness and adaptability of this approach. Brinkman-penalisation has also been employed by Chanta-

lat et al. [2009], coupled with a level-set technique to express the geometry, and applied by Bruneau et al. [2013] to optimise the shape of stirrers. An additional appeal of this method lies in its simple derivation and straightforward numerical implementation.

A second difficulty of the mixing problem lies in the nonlinear nature of its underlying governing equations. Optimising stirring strategies using a gradient-based approach will have to cope with the solution of nonlinear equations and the checkpointing problem for the dual/adjoint problem (see details below). The complexity of the flow does not furnish equilibrium points about which to linearise, rather the full nonlinear problem has to be tackled, and nonlinear adjoint looping techniques (Juniper [2010]) have to be employed for the stirrer geometry and/or stirrer paths.

A key challenge arises from the parameterisation for the stirrers' geometry. Even though the pointwise definition of the penalisation method would allow for an extensive and diverse number of solid shapes [Hejlesen et al., 2015], the resulting excessive dimensionality of the control variable does not lend itself to realistic and applicable optimisation outcomes. On the other hand, any choice of reducing the dimensions of the design space will influence the final optimisation results, as more exotic, yet potentially optimal, shapes are excluded. Consequently, care must be taken in selecting a suitable parameterisation to represent the embedded stirrers. Examples of shape parameterisations include Bézier curves [Han et al., 2009] or Hermite splines [De Boor et al., 1987]. While these methods allow for an acceptable range of possible complex shapes, they quickly approach a break-even point between flexibility and computational cost. Instead, we turn to a Fourier-based shape parameterisation, also referred to as Fourier descriptors [Lin & Chellappa, 1987]. Relying on Fourier series, we are thus able to select the dimensionality of the control space – without substantially compromising on the design space. The effectiveness of this approach will be demonstrated throughout this paper.

With these challenges identified, and respective algorithms chosen, we can then attempt to enhance mixing via stirrer-based strategies using a nonlinear direct-adjoint optimisation framework, coupled to a shape-optimisation approach. While various partial aspects of this undertaking have been attempted earlier, the synthesis of all components constitutes a novel pursuit.

Previously, optimal control of mixing via entropy maximisation of a flow governed by two orthogonal shear flows has been attempted [D'Alessandro et al., 1999], as well as the optimal mixing of binary fluids by optimising the mix-norm [Lin et al., 2011, Mathew et al., 2007, Vermach & Caulfield, 2018]; a concise review is presented in Ottino [1990] and Thiffeault [2012]. These studies, however, do not address shape optimisation in the context of mixing.

Relevant studies that do attempt shape optimisation fall into two general categories: (i) Brinkman-penalisation-based shape optimisation of mixing coupled to algorithms that are not gradient-based or nonlinear, and (ii) gradient-based methods for shape optimisation that do not focus specifically on mixing. Among some samples from the former category, we mention the mixing optimisation of CO₂ ejectors using genetic algorithms [Palacz et al., 2016], which achieved modest gains in mixing efficiency, the optimisation of diffuser shapes by a derivative-free optimisation approach [Madsen et al., 2000], and the shape optimisation of a heating element for desired temperature uniformity [Smolka, 2013], again using genetic meta-heuristic algorithms. In the latter category, we mention the work by Alexandersen et al. [2014], Pinggen et al. [2007], Saglietti et al. [2017; 2018] which employ adjoint (gradient-based) technology as well as Brinkman-penalisation, but consider natural convection rather than inertia-driven mixing of binary fluids.

The work presented in this article builds on the nonlinear direct-adjoint framework presented in Eggel & Schmid [2018], where a first attempt at shape optimisation, computing the optimal parameters of elliptical cross-sections, has been made. Building on this starting point, we generalise the stirrer geometry, while maintaining the circular shape of the mixing container (a commonplace configuration in industrial applications [Hemrajani & Tattersson, 2004]). More specifically, we consider three test cases of increasing geometric

complexity, which allows the optimisation framework to fully explore the design space and enhance mixing efficiency through the use of optimal shapes.

The remainder of this article is organised as follows. Section §2 introduces the governing equations and the mathematical background related to the optimisation of mixing. This is followed by a brief discussion of Brinkman-style penalisation, applied to the governing equations; discretisation details are given as well. We then turn to presenting the direct-adjoint optimisation procedure, together with the Fourier-shape parameterisation, in section §3. Next, we present the results of the optimisation studies in section §4; enhanced mixing is reported in all cases. Conclusions are offered in section §5.

2. Mathematical background

2.1 Governing equations

We will consider the active mixing of a binary fluid by stirring devices, for a parameter regime where inertial effects cannot be dismissed, whilst turbulent fluid motion is not yet present. The material behavior of the two fluid components is aptly described as a Newtonian fluid. This particular choice of regime and characteristics reflects the fact that these properties represent values commonly encountered in industrial settings; moreover, they allow a rich and interesting suite of physical mechanisms that can be exploited for the design of optimised stirring protocols.

The underlying governing equations that describe the motion of our fluid are the incompressible Navier-Stokes equations, which in primitive form read

$$\partial_t \mathbf{u} + \mathbf{u} \cdot \nabla \mathbf{u} + \nabla p - Re^{-1} \nabla^2 \mathbf{u} = 0 \quad (2.1)$$

$$\nabla \cdot \mathbf{u} = 0 \quad (2.2)$$

with \mathbf{u} as the velocity vector and p denoting the pressure field. We have stated the equations in non-dimensional form, using an appropriate characteristic length (L_0) and velocity (u_0), based upon which the dependent and independent variables are rendered dimensionless. This latter process introduces the non-dimensional fluid viscosity as the inverse Reynolds number with $Re = \frac{u_0 L_0}{\nu}$ and ν is the kinematic viscosity.

To fully describe a mixing process, an additional equation is introduced that tracks the two fluids (and their mixing) in the form of a passive scalar θ . This passive scalar, which takes on the value of $\theta = 0$ in fluid one, $\theta = 1$ in fluid two, and values in-between for a mixture of the two fluids, is governed by an advection-diffusion equation of the form

$$\partial_t \theta + \mathbf{u} \cdot \nabla \theta - Pe^{-1} \nabla^2 \theta = 0. \quad (2.3)$$

Another non-dimensional parameter, the Péclet number Pe , is defined as $Pe = \frac{u_0 L_0}{\kappa}$ and represents the dimensionless form of the diffusivity κ of the mixing process. The passive scalar will be used to track the composition of the binary mixture, to measure mixedness of our fluid and to subsequently design active stirring strategies to optimise this mixedness.

2.2 Measuring mixedness

Numerous studies have considered effective measures for quantifying mixedness, among them Foures et al. [2014], Liu [2008], Thiffeault [2012], where the latter performed a comparative study of energy norm, variance and mix-norm. While all measures have been successful in generating efficient mixing strategies within an optimisation framework, the most effective has been found to be the mix-norm defined by

$$\|\theta\|_{\text{mix}} = \frac{1}{V_\Omega} \int_\Omega \|\nabla^k \theta(\mathbf{x}, t)\| d\Omega. \quad (2.4)$$




FIG. 1: Representative geometric setup of the computational mixing examples.

We note that $k = 0$ is equivalent to the variance of the passive scalar. The mix-norm, which is a Sobolev norm of fractional and/or negative index k , was introduced by Mathew et al. [2007] and has its mathematical origins in measure theory. The mix-norm, which requires k to be negative [Thiffeault, 2012], focuses specifically on the amount of small scales in the passive scalar and the transfer between large and small scales during mixing. There are two mechanisms which drive mixing, advection and diffusion, which are active and passive mechanisms, respectively. As we are only able to directly control advective processes, the seemingly most efficient way to exploit the passive diffusion process is by creating small scales and elongated filaments, using advective processes. Once these filaments have been created, diffusion will be far more effective and rapid. This emphasis in scale-transfer is best represented by the mix-norm – making the mix-norm a crucial ingredient of our optimisation attempt. We wish to stress, however, that the choice of measure does not affect the optimisation framework: no conceptual (but certainly quantitative) changes would occur if, for example, the variance were chosen instead.

2.3 Fluid-solid interactions

We accomplish our mixing with the help of embedded stirrers; this setup defines our attempt as a fluid-solid interaction problem. A representative geometry of our setup is displayed in figure 1. We will use a penalisation method [Arquis & Caltagirone, 1984] to reformulate the problem. In essence, this method recasts the governing equations by including additional penalisation terms, which model the solid stirrers as Brinkman-style porous media with a vanishing permeability C_η . As the permeability tends to zero, it has been shown that the modified equations converge towards solutions of the Navier-Stokes equations for the complex geometry [Angot et al., 1999]. The attractiveness of this method lies in its simple implementation, its flexibility in expressing boundary conditions and its numerical efficiency. Solids are simply implemented using an indicator function, χ , which is also referred to as the mask function. As we will treat a number of stirrers, we will introduce sub-masks which describe each individual solid

$$\chi_i = \begin{cases} 1, & \text{if } \mathbf{x} \in \Omega_{s_i} \\ 0, & \text{else} \end{cases} . \quad (2.5)$$

Since we impose a non-overlapping geometry we can combine these sub-masks into one general mask χ according to

$$\chi = \sum_{i=1}^n \chi_i \quad (2.6)$$

which allows us to distinguish between the solid and the fluid parts of the computational domain. An additional benefit of this method is the rather straightforward treatment of time-dependent (moving) solids. The movement of solid bodies through the fluid is achieved by remapping the mask function according to the updated locations of the bodies, without any need for a complex regridding operation. The mask function allows us to augment the Navier-Stokes equations (2.2) which then model our fluid-solid interactions as

$$\partial_t \mathbf{u} + \mathbf{u} \cdot \nabla \mathbf{u} + \frac{\chi}{C_\eta} \mathbf{u} - \frac{\chi_i}{C_\eta} \mathbf{u}_{s,i} + \nabla p - Re^{-1} \nabla^2 \mathbf{u} = 0, \quad (2.7)$$

$$\nabla \cdot \mathbf{u} = 0. \quad (2.8)$$

We note that for $\chi = 0$, we describe the fluid domain and recover the standard Navier-Stokes equations. Previous studies have shown that the optimal value of C_η , that provides the most accurate numerical result, is proportional to $(\Delta x)^2$; see Kadoch et al. [2012]. This penalisation approach will analogously be applied to the passive scalar equation (2.3), yielding the penalised version [Engels et al., 2015]

$$\partial_t \theta + (1 - \chi) \mathbf{u} \cdot \nabla \theta + \chi_i (\mathbf{u}_{s,i} \cdot \theta) - \nabla \cdot \left(\left[Pe^{-1} (1 - \chi) + \frac{\chi}{C_\eta} \right] \nabla \theta \right) = 0 \quad (2.9)$$

The terms $(1 - \chi) \mathbf{u} \cdot \nabla \theta$ and $\chi_i (\mathbf{u}_{s,i} \cdot \theta) - \nabla \cdot \left(\left[Pe^{-1} (1 - \chi) + \chi / C_\eta \right] \nabla \theta \right)$ prevent the scalar field θ from advecting or diffusing, respectively, into any of the solid bodies, while the term $\chi_i (\mathbf{u}_{s,i} \cdot \theta)$ imposes the transport of the passive scalar with the velocity of the solid.

2.4 Numerical method

As details of the numerical method can be found in Eggl & Schmid [2018], we will only briefly touch upon implementation specifics; for more technicalities the interested reader is referred to the original paper. Starting point for the numerical implementation of the penalised governing equations is the open-source software FluSI [Engels et al., 2015]. This Fortran 90 Fourier-based, pseudo-spectral code solves the nonlinear, three-dimensional, incompressible Navier-Stokes equations on an equidistant grid and incorporates adaptive time-stepping as well as periodically uncoupled domains. A formulation in spectral space and a pressure-projection approach is used. The fluid-solid interactions are treated by the Brinkmann-style penalisation method discussed above, and specific grid layers are used to impose outflow and open boundaries.

Following the spectral discretisation strategy implemented in FluSI, we will replace the spatial derivatives with a multiplication of the discretised velocity, pressure and scalar fields by a Fourier discretisation matrix. To this end, we introduce the discrete differentiation analog

$$\frac{\partial}{\partial x_i} \rightarrow A_i \quad (2.10)$$

where A_i is a $n \times n$ matrix with n as the number of grid points in a single coordinate direction. To represent the associated gradient vector ∇ , we introduce the tensor \mathbf{A} , with A_i as its entries. State variables are similarly discretised, i.e. \mathbf{u} and $\mathbf{u}_{s,i}$ become $3 \times n$ vectors containing the fluid and solid velocity components, respectively. We note that $(\mathbf{u}_{s,i})_j$ is a $1 \times n$ vector and refers to the j^{th} velocity component of the i^{th} solid. The variables χ , θ and p turn into $1 \times n$ vectors that represent the mask, passive scalar field and pressure, respectively. With these new variables, and assuming Einstein summation convention, we can state the

spatially discretised form of the governing equations as follows

$$\partial_t \mathbf{u} + \mathbf{u}_j \circ [\mathbf{A}_j \mathbf{u}] + \frac{\chi}{C_\eta} \circ \mathbf{u} - \frac{\chi_i}{C_\eta} \circ \mathbf{u}_{s,i} + \mathbf{A}p - Re^{-1} \mathbf{A}_i \mathbf{A}_i \mathbf{u} = 0, \quad (2.11)$$

$$\mathbf{A}_i \mathbf{u}_i = 0, \quad (2.12)$$

$$\partial_t \theta - \mathbf{A}_i ([Pe^{-1}(\mathbf{1} - \chi) + \kappa \chi] \circ \mathbf{A}_i \theta) + (\mathbf{1} - \chi) \circ \mathbf{u}_j \circ [\mathbf{A}_j \theta] + \chi_i \circ (\mathbf{u}_{s,i})_j \circ [\mathbf{A}_j \theta] = 0 \quad (2.13)$$

where \circ denotes the Hadamard (elementwise) product [Horn & Johnson, 2012]. To enforce the continuity equation, we utilise a operator-splitting approach, producing a pressure Poisson equation of the form

$$\mathbf{A}_j \mathbf{A}_j p + \mathbf{A}_i (\mathbf{u}_j \circ [\mathbf{A}_j \mathbf{u}_i]) + \mathbf{A}_i \left[\frac{\chi}{C_\eta} \circ \mathbf{u} - \frac{\chi_i}{C_\eta} \circ \mathbf{u}_{s,i} \right] = 0 \quad (2.14)$$

which completes the discretised system of equations describing the motion of the fluid and passive scalar.

3. Optimisation framework

3.1 The cost functional

In the previous study by Eggl & Schmid [2018], the above system of equations has been used to optimise the velocities of rotating circular stirrers or the eccentricities of rotating elliptical stirrers. The primary conclusion from this study was that (i) increasing the rotational velocity generally enhanced mixing efficiency, and (ii) the shape of the stirrers changed significantly, once vortex shedding from distorted geometries could be exploited. The fusing of these two effects leaves unclear the role of pure shape optimisation. For this reason, we will concentrate in this study only on topology changes and assess their isolated effect on the overall mixing efficiency. More specifically, the control variables of our optimisation are shape-describing parameters. As a consequence of this setup, there is no need for an otherwise required L_2 -penalisation of \mathbf{u}_s . We thus can state our cost functional solely in terms of the mixing measure as

$$\mathcal{J} = \frac{1}{V_\Omega} \int_\Omega \|\nabla^k \theta\| \, d\Omega \Big|_T. \quad (3.1)$$

For our study, we select the mix-norm with an exponent k of $-2/3$, slightly deviating from the previous choices of $k = -1$ [Lin et al., 2011] or $k = -1/2$ [Foures et al., 2014], but falling within this range. In line with the discretisation of the governing equations, we present the semi-discretised form of the cost functional \mathcal{J} as

$$\mathcal{J} = \frac{\sqrt{[\mathbf{A}^{-2/3} \theta]^H \mathbf{M} [\mathbf{A}^{-2/3} \theta]}}{V_\Omega} \Big|_T. \quad (3.2)$$

where \mathbf{M} is a symmetric, positive weight matrix, taking into account the grid spacing and spatial weightings, and the superscript H refers to the conjugate transpose of the particular matrix or vector. As we are confined to an equidistant grid of n^2 mesh points, the weight matrix in our case is simply a diagonal matrix of the form,

$$\mathbf{M} = \frac{1}{n^2} \mathbf{I}, \quad (3.3)$$

with I as the identity matrix. We restate that this type of norm encourages the formation of filamented structures in the passive scalar field, which are subsequently subjected to diffusion, once the optimisation window has passed.

So far, we have not imposed any constraints on the control variables. However, as the shapes need to remain physically realisable, as well as non-overlapping, some restrictions need to be enforced. However, these constraints are not included directly in the cost functional, but are instead applied during the update step of the control variables, once the direct-adjoint optimisation is completed.

3.2 The adjoint system

The cost functional \mathcal{J} depends explicitly only on θ . The dependence on our control variables as well as other state variables is implicit via their effect on θ . The minimisation of our cost functional \mathcal{J} has to observe the full governing equations, which is accomplished by embedding the governing equations as PDE-constraints into a single augmented cost functional \mathcal{L} . This embedding is achieved using Lagrange multipliers or adjoint variables. The augmented cost functional then takes the following form

$$\begin{aligned} \mathcal{L} = \mathcal{J} - \int_0^T (\mathbf{u}^\dagger)_k^H M \left\{ \partial_t \mathbf{u} + \mathbf{u}_j \circ [\mathbf{A}_j \mathbf{u}] + \frac{\chi}{C_\eta} \circ \mathbf{u} - \frac{\chi_i}{C_\eta} \circ \mathbf{u}_{s,i} + \mathbf{A}p - Re^{-1} [\mathbf{A}_i \mathbf{A}_i \mathbf{u}] \right\}_k \\ + p^\dagger \cdot^H M \left\{ [\mathbf{A}_i \mathbf{A}_i] p + \mathbf{A}_i (u_j \circ [\mathbf{A}_j u_i]) + \mathbf{A}_i \left[\frac{\chi}{C_\eta} - \frac{\chi_i}{C_\eta} \circ \mathbf{u}_{s,i} \right] \right\} \\ + \theta^\dagger \cdot^H M \left\{ \partial_t \theta + (1 - \chi) \circ \mathbf{u}_j \circ [\mathbf{A}_j \theta] - \chi \circ (\mathbf{u}_{s,i})_j \circ [\mathbf{A}_j \theta] - \mathbf{A}_i ([Pe^{-1}(1 - \chi) + \kappa \chi] \circ \mathbf{A}_i \theta) \right\} \\ + \chi_i^\dagger \cdot^H M [\chi_i - g_i(\mathbf{x}, t)] dt. \end{aligned} \quad (3.4)$$

This expression demonstrates that all governing equations and boundary conditions have been enforced using Lagrange multipliers, denoted by $(\cdot)^\dagger$; these multipliers are also referred to as adjoint variables and will subsequently supply our optimisation algorithm with sensitivity or gradient information. The key to deriving the optimality conditions is the minimisation of the augmented cost functional \mathcal{L} with respect to all dependent variables, direct or adjoint. This is accomplished by taking the first variation of (3.4) and evaluating it to zero, i.e., enforcing $\delta \mathcal{L} = 0$, breaking it into its various subcomponents. We note that the variation with respect to the adjoint variables recovers the governing equations (momentum, continuity and passive scalar equations). The first variation with respect to the state variables involves more algebraic effort, but leads to an evolution equation/algebraic equation for the adjoint variables. The explicit derivation will not be shown here; details of it can be found in the appendix of Eggl & Schmid [2018]. The final adjoint equations, governing \mathbf{u}^\dagger , p^\dagger and θ^\dagger reads as follows

$$\partial_t \mathbf{u}_i^\dagger - \Pi_k^\dagger \circ [\mathbf{A}_i \mathbf{u}_k] - \mathbf{A}_j^H [\mathbf{u}_j \circ \Pi_i^\dagger] - \frac{\chi}{C_\eta} \circ \Pi_i^\dagger + Re^{-1} \mathbf{A}_j^H \mathbf{A}_j^H \mathbf{u}_i^\dagger - (\mathbf{1} - \chi) \circ \theta^\dagger \circ [\mathbf{A}_i \theta] = 0, \quad (3.5)$$

$$\mathbf{A}_j^H \Pi_j^\dagger = 0, \quad (3.6)$$

$$\partial_t \theta^\dagger - \mathbf{A}_j^H [(1 - \chi) \circ \mathbf{u}_j \circ \theta^\dagger] + \mathbf{A}_i^H ([Pe^{-1}(1 - \chi) + \kappa \chi] \circ \mathbf{A}_i^H \theta^\dagger) - \mathbf{A}_j^H [\chi_i \circ (\mathbf{u}_{s,i})_j \circ \theta^\dagger] = 0, \quad (3.7)$$

with initial conditions

$$\mathbf{u}^\dagger(\mathbf{x}, T) = 0, \quad \theta^\dagger(\mathbf{x}, T) = \frac{2}{V_\Omega} (\mathbf{A}^{-2/3})^H M (\mathbf{A}^{-2/3} \theta). \quad (3.8)$$

The optimality condition arises from the first variation with respect to the control variable, i.e., the mask χ . We obtain

$$\chi_i^\dagger = \left[\theta^\dagger \circ [A_j \theta] - \frac{\Pi_j^\dagger}{C_\eta} \right] \circ (\mathbf{u}_j - (\mathbf{u}_{s,i})_j) + (\kappa - Pe^{-1} A_j^H \theta^\dagger \circ A_j \theta) \quad (3.9)$$

where $\Pi_i^\dagger = \mathbf{u}_i^\dagger + A_i^H p^\dagger$. We recognise that the above system constitutes a set of simultaneous equations that has to be solved for the optimal solution. In practice, this solution is obtained progressively, by solving the direct and adjoint equations exactly and by iterating the optimality condition until a user-specified criterion is satisfied.

Due to the fact that our governing equations are nonlinear and, as a consequence, the optimisation problem is non-convex, the iterative scheme outlined above can only yield a local minimum. A globally optimal solution cannot be guaranteed. For this reason, care must be exercised when choosing an initial condition, as this choice may have a noticeable influence on the overall result.

The optimality equation (3.9), based on the mask χ_i , allows for a maximum of flexibility in choosing the cross-sectional shape of the stirrers. To avoid discontinuous geometries and to reduce the dimensionality of the control space, we will make further modifications to the geometric representation of the stirrers' shape (beyond the mask function) to make the optimisation problem computationally more tractable.

3.3 Complex shape parameterisation

Previous studies on stirrer shape optimisation dealt with a very simplified geometry, the morphing from circular to elliptical cross-sections, intended to validate the computational framework [Eggl & Schmid, 2018] rather than accommodate complex shapes. However, with a view towards industrially applicable configurations that cannot be approached by intuition alone, a geometric parameterisation with a far higher degree of flexibility – while still conforming to the optimisation framework of the previous section – has to be adopted. In our choice, we are motivated by our spatial discretisation and hence introduce a Fourier-based shape parameterisation [Crimmins, 1982] shown to allow for an wide array of logically circular shapes. We have

$$f_x(\alpha) = \frac{a_0}{2} + \sum_{k=1}^n a_k \cos(k\alpha) - b_k \sin(k\alpha), \quad (3.10a)$$

$$f_y(\alpha) = \frac{c_0}{2} + \sum_{k=1}^n c_k \cos(k\alpha) - d_k \sin(k\alpha), \quad (3.10b)$$

where α is a vector representing the discretised interval $[0, 2\pi]$. A parametric curve in Cartesian coordinates is then given by

$$x = f_x(\alpha), \quad (3.11a)$$

$$y = f_y(\alpha), \quad (3.11b)$$

or, in a more compact form,

$$f(a_i, b_i, c_i, d_i) = [f_x, f_y], \quad (3.12)$$

which forms the basis for the optimisation of the i^{th} stirrer shape. The parametric representation of the stirrer geometries introduces additional computational issues that need to be addressed. The relation of

the parametric shape back to the mask χ requires an indicator function which specifies whether a point falls inside or outside a chosen geometry. For the previous cases of circular or elliptical cross-sections this condition could be formulated in a trivial manner. As our new parameterisation allows for shapes of significant concavity, we have to introduce a more sophisticated scheme to distinguish the interior from the exterior of our shape. Furthermore, additional complications arise from highly distorted geometric shapes whose barycentre falls outside the solid structure. To resolve these issues, we introduce the winding number $w(\mathbf{x})$. This number is defined (within a discrete setting) as

$$w(\mathbf{x}_k) = \sum_{i=1}^n \phi_{i,i-1}, \quad (3.13)$$

where $\phi_{i,j}$ denotes the angle of the arc, which is centred at the test point \mathbf{x}_k and runs along the parametric curve from \mathbf{x}_i to \mathbf{x}_j . As the points $\mathbf{x}_i, \mathbf{x}_j$ and the angle ϕ sweep the full perimeter of the parametric curve (from index i to n), the total angle travelling around the curve will be 2π if the point \mathbf{x}_k is inside the curve, or zero otherwise.

Combining this concept with our previous definition of the mask χ , we obtain

$$\chi_i(\mathbf{x}, t) = \begin{cases} 1, & w_i(\mathbf{x}) = 2\pi, \\ 0, & w(\mathbf{x}) = 0, \end{cases} \quad (3.14)$$

We note that this formulation is discontinuous and thus would generate spurious numerical oscillations as the stirrers move across the Cartesian background grid. For this reason, we introduce a smoothing layer to avoid these artefacts. The design of this layer is further complicated, as the smoothing layer cannot simply be defined as a distance from the centre (as in Engels et al. [2015] or Eggl & Schmid [2018]); instead we must rely on the minimum distance from our chosen point to our curve, i.e.,

$$d = \min_{\alpha} \sqrt{(x - f_x)^2 + (y - f_y)^2}. \quad (3.15)$$

The angle of α that assumes this minimum value will be defined as $\tilde{\alpha}$ and can be uniquely determined for each point (x, y) inside the shape. We then augment our mask functions χ_i with d to obtain the smoothed mask

$$\chi_i(x, t) = \begin{cases} 1, & w = 2\pi \text{ and } d > h, \\ \frac{1}{2} \left(1 + \cos \left(\frac{\pi(h-d)}{h} \right) \right), & w = 2\pi \text{ and } d < h, \\ 0, & w = 0. \end{cases} \quad (3.16)$$

With this formulation we ensure a continuous and differentiable mask-based geometry.

A benefit of using the parameterisation above is its adaptability to the adjoint system. Without loss of generality, we will focus on optimising the Lagrangian \mathcal{L} with respect to one representative coefficient a_k , with the remaining Fourier coefficients following the same procedure. Since \mathcal{L} does not explicitly depend on a_k , but is linked implicitly through the geometric mask, we have to express the derivative with respect to a_k in terms of its partial components. We have

$$\frac{\partial \mathcal{L}}{\partial a_k} = \frac{\partial \mathcal{L}}{\partial \chi_i} \frac{\partial \chi_i}{\partial d} \frac{\partial d}{\partial a_k}. \quad (3.17)$$

The first term, the gradient $\frac{\partial \mathcal{L}}{\partial \chi_i}$, follows from equation (3.4), while the second term has been reported (in analogous form) in Eggl & Schmid [2018]. The last term is novel and will be further discussed below. While

equation (3.15) does include a minimisation, the derivative term in equation (3.17) is independent of this minimum and does not affect the subsequent derivation. Given the form of the parameterisation (equation (3.10)), the derivative is in essence isolating the wavenumber within the minimum distance function, i.e.,

$$\frac{\partial d}{\partial a_k} = \frac{(f_x - x) \cos(k\alpha)}{d} \Big|_{\bar{\alpha}}. \quad (3.18)$$

Combining this expression with the terms above, we arrive at the complete formulation of the optimality condition with respect to the chosen Fourier-coefficient a_k . The composite nature of equation (3.17) allows for an efficient implementation as the first two terms are independent of our choice of a_k and can be determined ahead of the direct-adjoint iterations. Furthermore, the final term only changes minimally, either by including a $(f_y - y)$ term or by replacing a sine- with a cosine-term, depending on our choice of coefficient. In short, the calculation of the optimality condition is efficient and readily scalable, as the majority of the required terms can be calculated *a priori*.

One further issue that needs attention is area conservation. The optimisation routine at this point is not constrained by the cross-sectional area of the stirrers and therefore can and will modify the stirrers accordingly. This may encourage shapes that achieve better mixing simply by changing the size of the stirrers. To counter this tendency, we need to define and conserve the area of our parametric curves. We start by deriving the area as follows

$$A = - \int_{\alpha=0}^{\alpha=2\pi} f_y df_x \quad (3.19)$$

$$= - \int_0^{2\pi} f_y \frac{\partial f_x}{\partial \alpha} d\alpha \quad (3.20)$$

$$= - \int_0^{2\pi} \left[\frac{c_0}{2} + \sum_{k=1}^n c_k \cos(k\alpha) - d_k \sin(k\alpha) \right] \left[\sum_{l=1}^n l(-a_l \sin(l\alpha) - b_l \cos(l\alpha)) \right] d\alpha. \quad (3.21)$$

We will absorb the factor $\frac{1}{2}$ into c_0 to simplify the notation

$$A = \sum_{k=0}^n \sum_{l=1}^n l \left(\int_0^{2\pi} [c_k \cos(k\alpha) - d_k \sin(k\alpha)] [a_l \sin(l\alpha) + b_l \cos(l\alpha)] d\alpha \right) \quad (3.22)$$

$$= \sum_{k=0}^n \sum_{l=1}^n l \left(\int_0^{2\pi} c_k b_l \cos(k\alpha) \cos(l\alpha) + c_k a_l \cos(k\alpha) \sin(l\alpha) \right. \\ \left. - d_k a_l \sin(k\alpha) \sin(l\alpha) - d_k b_l \sin(k\alpha) \cos(l\alpha) d\alpha \right). \quad (3.23)$$

Observing orthogonality between the trigonometric functions over the interval $(0, 2\pi)$, the only non-zero

contributions are given by $k = l$ for the $\cos \times \cos$ and $\sin \times \sin$ terms. This results in

$$\sum_{k=0}^n \sum_{l=1}^n \int_0^{2\pi} c_k b_l \cos(k\alpha) \cos(l\alpha) d\alpha = \sum_{k=0}^n \pi c_k b_k, \quad (3.24)$$

$$\sum_{k=0}^n \sum_{l=1}^n \int_0^{2\pi} c_k a_l \cos(k\alpha) \sin(l\alpha) d\alpha = 0, \quad (3.25)$$

$$\sum_{k=1}^n \sum_{l=1}^n \int_0^{2\pi} d_k b_l \sin(k\alpha) \cos(l\alpha) d\alpha = 0, \quad (3.26)$$

$$\sum_{k=1}^n \sum_{l=1}^n \int_0^{2\pi} d_k a_l \sin(k\alpha) \sin(l\alpha) d\alpha = \sum_{k=1}^n d_k a_k \pi. \quad (3.27)$$

Using the above and substituting into equation (3.23) we obtain

$$A = \sum_{k=1}^n \pi k (c_k b_k - d_k a_k). \quad (3.28)$$

The special cases for the area of a circle and ellipse can be recovered easily. We will use the general expression to normalise the coefficients and thus preserve the initial area. More specifically, for an initial area

$$\tilde{A} = \sum_{k=1}^n \pi k (\tilde{c}_k \tilde{b}_k - \tilde{d}_k \tilde{a}_k), \quad (3.29)$$

we multiply all Fourier coefficients for the new geometric shape (with area A) by $\sqrt{\frac{\tilde{A}}{A}}$ to obtain

$$\sum_{k=1}^n \pi k \left(\sqrt{\frac{\tilde{A}}{A}} c_k \sqrt{\frac{\tilde{A}}{A}} b_k - \sqrt{\frac{\tilde{A}}{A}} d_k \sqrt{\frac{\tilde{A}}{A}} a_k \right) = \frac{\tilde{A}}{A} \sum_{k=1}^n \pi k (c_k b_k - d_k a_k) \quad (3.30)$$

$$= \frac{\tilde{A}}{A} A \quad (3.31)$$

$$= \tilde{A}, \quad (3.32)$$

i.e., an area-normalised new shape.

We now have in place the key ingredients to successfully optimise any number of stirrers of arbitrary Fourier-based shapes. This geometric parameterisation has been found to be particularly simple to implement and robust to optimise.

3.4 Shape constraints

The possibilities within the Fourier-based approach require additional constraints that go beyond area conservation. These additional constraints arise from the desire to generate physically feasible shapes that are implementable in an industrial setting.

When generating geometric shapes with a multitude of Fourier coefficients, it is rather likely that the parametric curve self-intersects, an example of which can be seen in the left shape of figure 2. In essence, the




FIG. 2: Unconstrained shape optimisation of two rotating stirrers. We can see that the left shape has resulted in two pinched off pieces due to the optimisation.

direct-adjoint optimisation overrules the imposed number of prescribed stirrers by generating new pinched-off geometric components that subsequently enhance mixing in later optimisation steps, albeit with the same dynamic constraints. While this option may be beneficial to the mix-norm, the geometric shape is, of course, unrealisable in a realistic setting. It may appear sufficient to simply expand the distance between the points on the curve and therefore avoid such a pinching off. However, the full path along the parametric curve needs to be considered as well, as any pure thickening of the curve will inevitably leave some twisting due to our use of the winding number w . What is instead needed is an untwisting strategy.


We consider two line segments \overline{AB} and \overline{CD} of arbitrary length and evaluate a criterion for their mutual intersection. If this criterion is met, without loss of generality, we exchange B and D and mirror the parameterisation of the part of the parametric curve from A to D to generate an untwisted, and simply connected, shape. An example of such an untwisting can be seen in figure 3a, which contains the previously shown shape that had two pinched-off areas. After this untwisting procedure, new Fourier coefficients can be calculated, supplied to the area calculation and passed on to further optimisation steps.

The above untwisting would be sufficient for a continuous setting. However, on a finite grid we can arrive at pinched-off shapes even after untwisting has been performed. This is due to small geometric features that fall below the resolution of the Cartesian background grid despite the fact that the shape is simply connected. To avoid this complication, we introduce a displacement strategy for points that fall below a chosen distance, so that their modified distances remain above the minimum grid spacing. This displacement is based on the following steps. Consider two points which are a distance r apart. If this distance falls below the minimum distance constraint, i.e., $r < r_{min}$, we displace the two points along their respective outward pointing normals to the parametric curve. The amount of displacement, ζ , is given by a Buckingham potential [Buckingham, 1938]


$$\zeta = r_{min} \left[e^{-\frac{r}{r_{min}}} \right]. \quad (3.33)$$

A demonstration of this procedure is displayed in figure 3b.

This strategy, in combination with the untwisting method, ensures a simply connected shape that is properly represented on our underlying grid.



(a) Untwisting algorithm applied to the left shape from figure 2. The new curve is displayed in red.



(b) Application of thickening routine to untwisted shape.

4. Optimisation results

4.1 Geometric configuration

Our previous investigations of shape-optimisation by direct-adjoint looping [Eggl & Schmid, 2018] using a circular initial geometry for the stirrers revealed a severely impeded convergence towards an optimal design as the initial solid-body rotation did not induce vortex shedding and hence resulted in negligible gradient information for enhanced mixing. As a consequence, a substantial amount of computational effort was expended, and ultimately squandered, over many iterations, before vortex shedding set in and further modifications to the stirrers' cross-section yielded improved mixedness. For this reason a non-circular initial geometry will be chosen, in an attempt to introduce vortex shedding already in the first iteration and on which the direct-adjoint optimisation will further improve. A more expedited convergence is thus expected. More particularly, we initialise the stirrers as four-pointed hypocycloids (astroids), examples of which can be seen in figure 4. While vortex shedding from the cusps will commence immediately, initially the vortices will be trapped in the astroids' cavities and show little interaction. This situation is then the starting point for our direct-adjoint shape-optimisation algorithm.

The geometric arrangements and rotational speed ($\omega = \frac{2\pi}{8}$) of the star-shaped stirrers have been chosen identical to our previous study [Eggl & Schmid, 2018] to allow a qualitative comparison. The Reynolds and Péclet number have also been taken from the same reference. In addition, the control space for the stirrer shapes has been limited to the five largest wave numbers, i.e.,

$$\{a_i, b_i, c_i, d_i\} \text{ for } i = 1, \dots, 5. \quad (4.1)$$

We curtail the optimisation horizon to one single rotation, i.e., to $T = 8$. This choice focuses the optimisation effort on short-term features, rather than on residual diffusive phenomena. In effect, we force the optimisation to achieve efficient mixing with severe limitations in time.

4.2 Case 1: Shape optimisation of one stirrer

We begin by considering a single, centrally placed stirrer. The effect on the scalar field by the action of the unoptimised (astroidal) stirrer and the fully optimised stirrer is displayed in figure 6. We observe that the shape generated by the optimisation is drastically more asymmetric than the initial (symmetric) configuration. This complex shape induces the shedding of vortices and their mutual interaction. In the absence of a second stirrer, a fully cooperative mixing strategy is not possible; rather, the single stirrer must achieve a self-collaborative strategy among its shed vortices. One strategy is to elongate the stirrer's shape: with a constant rotational speed, a longer shape causes a higher velocity at the trailing edge and a more vigorous shedding of vortices. This principal tendency is further augmented by an accompanying alignment




FIG. 4: Sketch of the initial configurations for the three shape optimisation cases. Case 1 (top left): one centred stirrer rotating about its centre. Case 2 (top right): two stirrers aligned on the horizontal axis, rotating in the clockwise directions. Case 3 (bottom centre): five rotating stirrers, with three cylinders aligned along the horizontal axis and two cylinders vertically displaced.

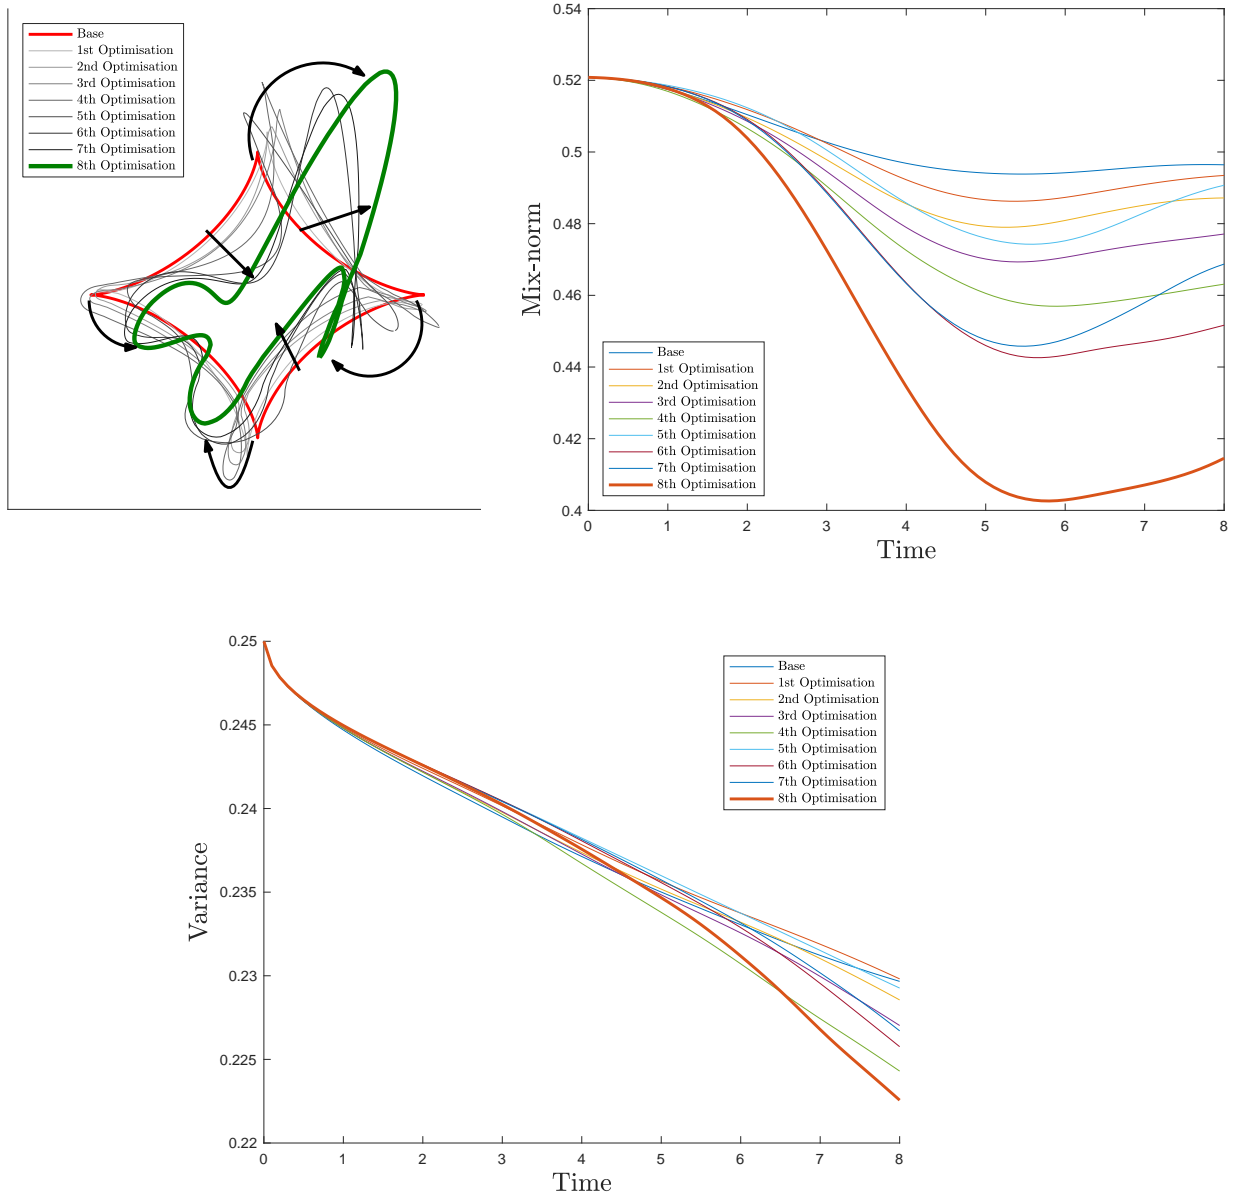


FIG. 5: Case 1: mixing optimisation using one fixed rotating stirrer. No untwisting or thickening was applied. (a) Evolution of the shape over the optimisation cycle. The red line is the initial configuration, the green line the final shape. Arrows have been added in black to illustrate the changes in shape. (b) Mix-norm of the passive scalar versus time $t \in [0, T]$. (c) Variance of the passive scalar versus time $t \in [0, T]$.

of the three remaining cusps along the axis of the principal cusp (see figure 5 for the gradual evolution into a final shape, with black arrows guiding the eye). We recall the conservation of the initial cross-sectional area of the stirrer, causing the observed stretching and thinning of the solid body.

The converged shape, after our final iteration, leads to the formation of a trailing-line vortex street which can be seen at $t = 6$ in the right column of figure 6. Furthermore, one of the protruding cusps near the centre of the stirrer generates a patch of high circulation which entrains unmixed passive scalar and, later in the rotation, launches vorticity towards the trailing-line vortex street ($t = 6$ to $t = 8$). We urge the reader to consult the animations in the supplemental material to fully appreciate this dynamic process.

One item of particular interest in this optimisation is the capability of the direct-adjoint looping, coupled with shape optimisation, to explore competing minima beyond a current local solution. This manifests itself in figure 5b where the final values of the mix-norm after the fifth and seventh iteration are significantly higher, indicating worse mixing, than during the previous iterations. Further analysis of this observation reveals that the shapes (shown in figure 5a) go through drastic changes, i.e., there is a determined shift from an initial, presumably optimal configuration to a more efficient one. Throughout the initial stages of the optimisation, the first attempts at a better shape reinforce a particular configuration. During this reinforcement, and owing to the global nature of the Fourier-based shape representation, the optimisation scheme becomes aware of a more optimal shape at a different point in design space. Though in the pursuit of this new configuration, the optimisation must pass through a regime of sub-optimality. Once this traversal is completed, we ultimately obtain a stirrer geometry with significantly enhanced mixing efficiency. Aiding in this behaviour is a strong and global link between all points of the Fourier-represented curves, and thus changes at a particular point are communicated effectively to the remaining curve. This strong global link allows the optimisation to probe multiple local minima. However, it must be called to mind that, due to the non-convexity of the problem, convergence to a global minimum cannot be guaranteed.

After eight optimisations we terminate the direct-adjoint loop, even though better mixing efficiency could be achieved in subsequent iterations, as the resulting shapes increasingly become structurally delicate and physically unfeasible. During these eight optimisations, no post-processing such as thickening or untwisting has become necessary; the resulting optimal shapes can thus conceivably be implemented in an industrial setting.

As a final observation, we note the rise of the mix-norm towards the end of the simulation. Similar behaviour has been observed in Vermach & Caulfield [2018]. This may seem counter-intuitive and raise the question why the optimisation did not terminate when the mix-norm is lowest. We remind the reader that our cost-functional focuses on the end-time evaluation of the mix-norm and, therefore, any influence prior to the final time horizon does not directly enter the optimisation. Parenthetically, the variance continues to decrease beyond the optimisation horizon, which implies that the fluid continues to mix.

4.3 Case 2: Shape optimisation of two stirrers

We complicate the geometry and optimisation scheme by introducing a second stirrer which we also place at the interface between the two fluids (see figure 4). Both stirrers are initially taken as a four-pointed hypocycloid and are placed sufficiently far from each other such that communication between them is weak. The interest in this configuration is two-fold: on one hand, we intend to assess the proximity-effects of the outer wall, on the other hand, we wish to investigate a possible collaborative mixing strategy between the two stirrers.

As can be seen from the left column of figure 8, showing the unoptimised state, there is no visible interaction between the stirrers and the outer wall; the shed vortices appear unaware of any other solid in




FIG. 6: Case 1: mixing optimisation using one rotating stirrer. Left column: unoptimised configuration, with snapshots at $t = 2, 4, 6, 8$ (top to bottom). Right column: after eight direct-adjoint optimisations, with snapshots at $t = 2, 4, 6, 8$ (top to bottom). For videos of these scenarios please refer to `Shape1NoOpt.mp4` and `Shape1Opt.mp4` for the left and right column, respectively.

their vicinity. In contrast, the optimised results illustrate an influence of the outer wall, in particular for the left stirrer, while the right stirrer tends to concentrate its mixing efforts on the local vortex dynamics. This break in symmetry can be attributed to the identical rotational direction of the stirrers. Firstly, we observe an elongation in the horizontal axis for both stirrers. This echoes the behaviour of the single cylinder optimisation, where elongation has been used to generate a vortex street. However, one fundamental difference prevails: whereas the only elongation for the single-stirrer optimisation was tilted at an angle to the interface to take advantage of a maximum plunging action, the elongation for the left stirrer in the present case is in the horizontal direction. The principal reason can be found in the presence of the wall, which enables and encourages effective vortex-wall collisions. The right stirrer, sensing far less influence of the wall, engages more in a plunging action and therefore develops a secondary elongation at an angle to the interface, similar to the one-stirrer case.

In general, we observe geometric features similar to the one-stirrer case, i.e., elongation in a principal direction and formation of concavities and protuberances, which generate and induce interacting vortices. While the left stirrer takes advantage of the outer wall in its vortex dynamics, the right stirrer follows the previous strategy of more local vortex interplays.

Overall, only a weak level of collaboration between the two stirrers can be discerned. Minor and final changes in the right stirrer geometry can nonetheless be attributed to an exchange of information via the velocity field; an example of this can be seen in form of the centre vortex at $t = 8$, which forms the start of a joint and collaborative mixing strategy.

Lastly, constraints on the geometry (untwisting and thickening) had to be employed to the left stirrer, ensuring a singly connected stirrer cross-section. The enforcement of these constraints mildly compromised the optimality suggested by the adjoint system. In other words, we traded a minor loss in optimality for physical feasibility. Despite this loss, a significant enhancement in mixing efficiency could be achieved again (see figure 7c).

4.4 Case 3: Shape optimisation of five stirrers

Encouraged by the advancement of a collaborative mixing strategy for multiple stirrers, we further facilitate this process by embedding additional stirrers into the fluid. Three stirrers will be placed on the interface, with two additional stirrers at an off-set to the interface; see figure 4 for a sketch of the configuration. The close proximity of the three stirrers is expected to result in a coordinated strategy between them and, possibly, an interaction with the off-set elements. Intuitively we may identify the off-set cylinders as non-critical to the overall mixing process. With no imposed energy constraints, however, this may not necessarily be the case for the optimisation algorithm.

Considering the unoptimised case (see the left column of figure 11), we note that, despite the close proximity of the central stirrers, little to no interaction is observed between the generated vortices. After only four iterations (right column), an entirely different picture arises. Stirrers one and two have been modified regarding their relative angles between the cusps and by the introduction of areas of high concavity (see figures 10a,b; the black arrows illustrate the evolution of the stirrer shapes). These modifications, when combined with the changes that the middle stirrer (five) has undergone, produce an interlocking gear mechanism. In the narrow gaps between the central stirrers, a great many filaments in the passive scalar are produced and thoroughly mixed. Less emphasis is placed on the interaction with the outer wall, and more focus is directed towards the collaboration between the stirrers, resulting in a highly distorted interface and the generation of small-scale structures. Vortices generated by one stirrer collide with a neighbouring

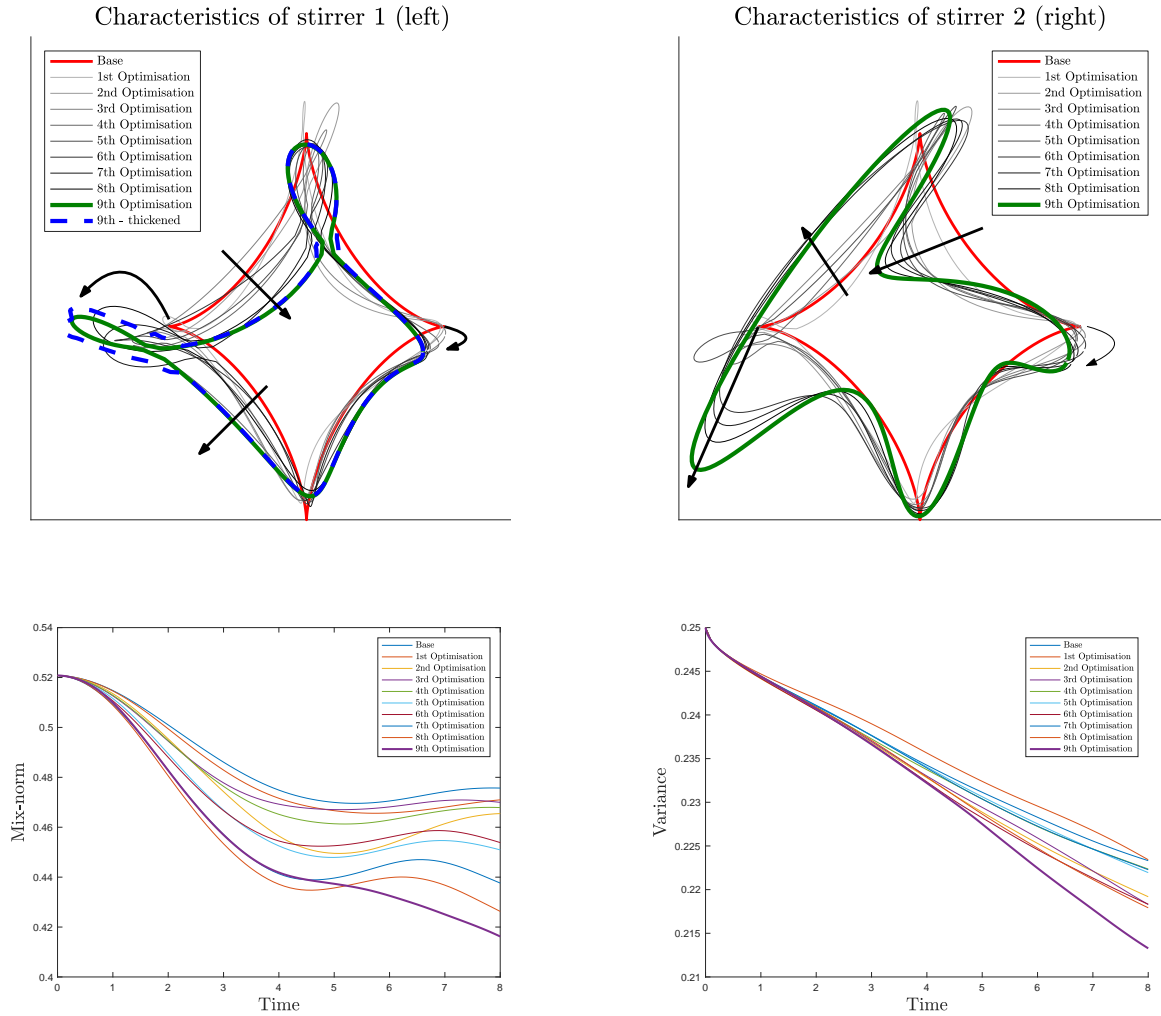


FIG. 7: Case 2: mixing optimisation using two fixed rotating stirrers. (a) Evolution of the shape of the left stirrer throughout the optimisation steps. The red line is the initial configuration, green the final shape, and the blue dashed line refers to the thickening routine applied to the final optimised geometry. Arrows have been added in black to illustrate the changes in shape. (b) Evolution of the shape of the right stirrer throughout the optimisation steps. (c) Mix-norm of the passive scalar versus time $t \in [0, T]$. (d) Variance of the passive scalar versus time $t \in [0, T]$.

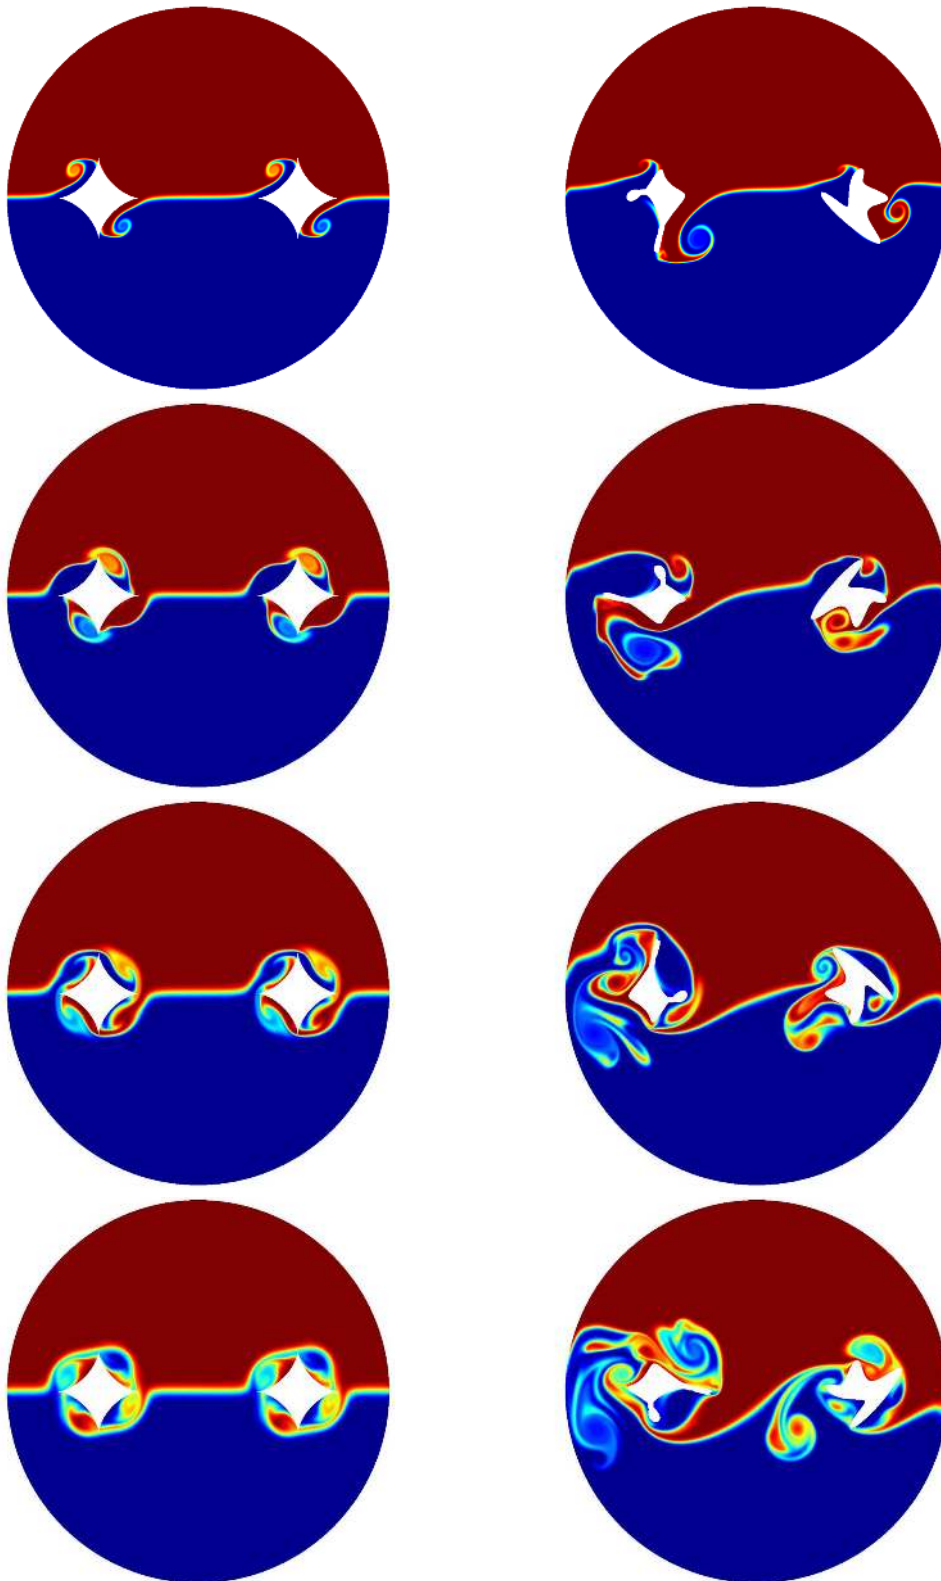


FIG. 8: Case 2: mixing optimisation using two rotating stirrers. Left column: unoptimised configuration, with snapshots at $t = 2, 4, 6, 8$ (top to bottom). Right column: after nine direct-adjoint optimisations, with snapshots at $t = 2, 4, 6, 8$ (top to bottom). For videos of these scenarios, please refer to `Shape2NoOpt .mp4` and `Shape2Opt .mp4` for the left and right column, respectively.




FIG. 9: Case 3: mixing optimisation using five stationary rotating stirrers. (a) Mix-norm of the passive scalar versus time $t \in [0, T]$. (b) Variance of the passive scalar versus time $t \in [0, T]$.

stirrer. The middle stirrer acts as a mediator between the two outer and more pointed stirrers, enabling a full and effective communication between the three central cylinders. We stress that the interlocking geometry resulting in the gear-like motion of the rotating stirrers was produced by the adjoint optimisation without explicitly enforcing or encouraging it.

The outer off-set cylinders undergo far less change, but nevertheless contribute in a minor fashion to the overall mixing process. In contrast to the analogous case studied in Eggl & Schmid [2018], where energy considerations (L^2 -norm of \mathbf{u}_s) played an important role to de-emphasise the outer cylinders, in the present case (with no energy constraints) the off-set stirrers eventually participate in decreasing the mix-norm. Specifically the upper off-set stirrer has been altered, so that two of its cusps have been extended to reach the mixed fluid. This can be observed at $t = 8$ in figure 11, where minute filaments of fluid are transported into the upper half of the vessel as a consequence.

One of the more remarkable aspects of this result consists of the fact that it was achieved in only four iterations. The proximity of the stirrers, and the information they communicate to each other, is an effective incentive for the optimisation algorithm and ensures rapid convergence to a collaborative strategy. We emphasise once again that the initial rotational velocity of each stirrer has been kept constant throughout, and solely shape optimisation has been performed. In view of these restrictions, and despite a rather limited time horizon of one rotation, a significant decrease in the mix-norm has been achieved.




FIG. 10: Case 3: mixing optimisation using five stationary rotating stirrers. (a) Evolution of the shape of stirrer 1 throughout the optimisation steps. The red line is the initial configuration, green the final one and the blue dashed line refers to the thickening routine applied to the final optimised shape. Arrows have been added in black to illustrate the changes in the shape. (b) Evolution of the shape of stirrer 2 throughout the optimisation steps. (c) Evolution of the shape of stirrer 5 throughout the optimisation steps.




FIG. 11: Case 3: mixing optimisation using five rotating stirrers. Left column: unoptimised configuration, with snapshots at $t = 2, 4, 6, 8$ (top to bottom). Right column: after four direct-adjoint optimisations, with snapshots at $t = 2, 4, 6, 8$ (top to bottom). For videos of these scenarios please refer to `Shape5NoOpt .mp4` and `Shape5Opt .mp4` for the left and right column, respectively.

5. Summary and conclusions

We have introduced a mathematical and computational framework, based on direct-adjoint iterative optimisation, for treating the mixing enhancement in binary fluids by topological shape modification of the embedded stirrers, which are represented by a Fourier-based parametric curve. The derivation and application of this shape parameterisation is surprisingly simple, and the resulting design space, while rather low-dimensional, allows for a large number of possible stirrer configurations.

Three test cases of increasing complexity have been presented to demonstrate the ease of use and power of the computational methodology: (i) the optimisation of one central stirrer, (ii) a double configuration along the initial interface, and (iii) a layout with five stirrers. In all cases significant improvement in mixing efficiency could be accomplished, and the optimisation algorithm showed notable robustness in finding a more optimal solution. These optimal solutions often include non-intuitive shapes that, when implemented in an industrial setting, could lead to substantial savings in expended effort or to higher-quality products.

A key observation of the obtained solutions is the generation of vortex filaments by the moving stirrers. Their mutual interactions dominates the mixing process. In contrast to the spatially fixed stirrers, these vortical elements are able to move freely within the mixing vessel and therefore can affect areas of the fluid that are otherwise inaccessible to the stirrers but that are highly beneficial to mixing enhancement. This effect could be observed by the convergence of the geometries towards protuberances and concavities for each of the stirrers. These two geometric components, combined with the rotation of the stirrers, generated areas of fluid recirculation and lead to the creation and shedding of vortices. The vortices then interacted and collided with previously and subsequently generated vortices, further distorting the interface and creating small scales for their later diffusion.

A further observation that lead to significant mixing enhancement (also reported in Eggl & Schmid [2018]) was the evolution of collaborative mixing strategies when multiple stirrers are optimised. The interlocking gear design of the optimised shapes of the five stirrer configuration is an example of this cooperation. In this case, the cusp of a rotating stirrer is placed directly in the path of a vortex being shed by another solid. This joint strategy requires little extra expended kinetic energy, but can yield a substantial increase in mixing efficiency. Further mixing studies may take advantage of a multi-stirrer configuration and exploit this collaboration of stirrers to additionally enhance mixing.

Despite these promising results, a few challenges remain. As is the case for any gradient-based optimisation applied to a non-convex problem, only a local minimum can be guaranteed by our algorithm. The Fourier-based shape parameterisation somewhat alleviates this problem, as the global nature of the representation allows the exploration of multiple local minima. Nonetheless, a global minimum is unattainable with our approach. A possible solution could lie in coupling our framework with an annealing-type stochastic algorithm Kirkpatrick et al. [1983], but the excessive computational cost associated with these methods prohibits its use in our case. Despite this issue, local optima should not be dismissed *ab initio*, as even small improvements in efficiency can equate into substantial gains and profits at an industrial scale.

A further challenge arises from the application of the Fourier-based shape parameterisation. The thickening and untwisting procedures have been embedded via a projection approach, after the optimisation step has taken place. While this approach achieved a significant increase in mixing efficiency, the resulting shapes were not strictly optimal. Mathematically, the geometric constraints should be incorporated into the cost functional; the details of this operation, however, remain an open challenge.

While keeping these challenges in mind, the investigations presented in this article point towards several areas where our optimisation framework can excel and achieve enhanced mixing efficiency.

A natural extension consists in a cooperative combination of the Fourier-based shape optimisation and the optimisation of time-dependent quantities, such as the velocity of stirrers along a given path. Exploring this *combined* design space will also give insight into the underlying mechanisms that drive mixing, e.g.,

flapping, plunging, heaving or rotation.

Extending the Fourier representation of the shape, the equivalent mathematical representation based on parametric curves can also be applied to optimise the trajectory of moving stirrers. The Fourier-based parameterisation for the stirrers' paths opens new areas of the design space. For example, proposing parameterised Lissajous curves [Cundy & Rollett, 1961] for each stirrer could yield pertinent and interesting results, with potentially novel and non-intuitive mixing strategies. This approach has the added benefit that for path optimisation, the constraints of untwisting and thickening of the geometry can be safely ignored, since any continuous, and even self-intersecting, trajectory represents a viable mixing strategy. However, care must be taken to avoid collisions among stirrers and/or the wall.

All the above extensions follow the two-dimensional nature of the test cases studied in this article and ignore three-dimensional effects. The computational framework is highly parallelised and scales well to many cores on parallel computer architectures; the Fourier-based parameterisation can also naturally be extended to three dimensions. Scalable performance for large-scale shape optimisations of binary mixing problems is thus expected.

In view of the wide scope of possibilities offered by the direct-adjoint optimisation approach and its robustness in finding enhanced mixing solutions, promising results can be expected in future studies.

REFERENCES

- Aamo, O. M. & Krstic, M. (2013) *Flow control by feedback: stabilization and mixing*. Springer Science & Business Media.
- Alexandersen, J., Aage, N., Andreasen, C. S. & Sigmund, O. (2014) Topology optimisation for natural convection problems. *Int. J. Num. Meths. Fluids*, **76**(10), 699–721.
- Angot, P., Bruneau, C.-H. & Fabrie, P. (1999) A penalization method to take into account obstacles in incompressible viscous flows. *Num. Math.*, **81**(4), 497–520.
- Annaswamy, A. & Ghoniem, A. (1995) Active control in combustion systems. *IEEE Contr. Sys.*, **15**(6), 49–63.
- Aref, H. (1984) Stirring by chaotic advection. *J. Fluid Mech.*, **143**(1), 1–21.
- Arquis, E. & Caltagirone, J. (1984) Sur les conditions hydrodynamiques au voisinage d'une interface milieu fluide-milieu poreux: application a la convection naturelle. *C.R. Acad. Sci. Paris II*, **299**, 1–4.
- Batchelor, G. K. (1952) The effect of homogeneous turbulence on material lines and surfaces. *Proc. Roy. Soc. Lond. A*, **213**(1114), 349–366.
- Boiron, O., Chiavassa, G. & Donat, R. (2009) A high-resolution penalization method for large Mach number flows in the presence of obstacles. *Comp. & Fluids*, **38**(3), 703–714.
- Bruneau, C.-H., Chantalat, F., Iollo, A., Jordi, B. & Mortazavi, I. (2013) Modelling and shape optimization of an actuator. *Struct. Multidisc. Optim.*, **48**(6), 1143–1151.
- Buckingham, R. A. (1938) The Classical Equation of State of Gaseous Helium, Neon and Argon. *Proc. Roy. Soc. Lond. A*, **168**, 264–283.
- Chantalat, F., Bruneau, C.-H., Galusinski, C. & Iollo, A. (2009) Level-set, penalization and Cartesian meshes: A paradigm for inverse problems and optimal design. *J. Comp. Phys.*, **228**(17), 6291–6315.
- Chien, W.-L., Rising, H. & Ottino, J. (1986) Laminar mixing and chaotic mixing in several cavity flows. *J. Fluid Mech.*, **170**, 355–377.
- Corrsin, S. (1957) Simple theory of an idealized turbulent mixer. *AIChE J.*, **3**(3), 329–330.
- Corrsin, S. & Karweit, M. (1969) Fluid line growth in grid-generated isotropic turbulence. *J. Fluid Mech.*, **39**(1), 87–96.
- Crimmins, T. R. (1982) A Complete Set of Fourier Descriptors for Two-Dimensional Shapes. *IEEE Trans. Sys., Man, and Cybern.*, **12**(6), 848–855.
- Cundy, H. M. & Rollett, A. P. (1961) *Mathematical models*. Clarendon Press Oxford.
- D'Alessandro, D., Dahleh, M. & Mezić, I. (1999) Control of mixing in fluid flow: a maximum entropy approach. *IEEE Trans. Autom. Contr.*, **44**(10), 1852–1863.
- De Boor, C., Höllig, K. & Sabin, M. (1987) High accuracy geometric Hermite interpolation. *Comp. Aided Geom.*

- Design*, **4**(4), 269–278.
- Dimotakis, P. E. (2000) The mixing transition in turbulent flows. *J. Fluid Mech.*, **409**, 69–98.
- Dimotakis, P. E. & Catrakis, H. J. (1999) Turbulence, fractals, and mixing. In *Mixing*, pages 59–143. Springer.
- Eggl, M. & Schmid, P. (2018) A gradient-based framework for maximizing mixing in binary fluids. *J. Comp. Phys.*, **368**, 131 – 153.
- Engels, T., Kolomenskiy, D., Schneider, K. & Sesterhenn, J. (2015) FluSI: A novel parallel simulation tool for flapping insect flight using a Fourier method with volume penalization. *SIAM J. Sci. Comp.*, pages 1–21.
- Foures, D., Caulfield, C. & Schmid, P. (2014) Optimal mixing in two-dimensional plane Poiseuille flow at finite Péclet number. *J. Fluid Mech.*, **748**, 241–277.
- Glowinski, M., Pan, T., Wells Jr., R. & Zhou, X. (1996) Wavelet and finite element solutions for the Neumann problem using fictitious domains. *J. Comp. Phys.*, **126**(1), 40–51.
- Han, X.-A., Ma, Y. & Huang, X. (2009) The cubic trigonometric Bézier curve with two shape parameters. *Appl. Math. Letts.*, **22**(2), 226–231.
- Hejlesen, M. M., Koumoutsakos, P., Leonard, A. & Walther, J. H. (2015) Iterative Brinkman penalization for remeshed vortex methods. *J. Comp. Phys.*, **280**, 547–562.
- Hemrajani, R. & Tattersson, G. (2004) *Mechanically Stirred Vessels*, pages 345–390. John Wiley & Sons, Inc.
- Hessel, V., Löwe, H. & Schönfeld, F. (2005) Micromixers – A review on passive and active mixing principles. In *Chemical Engineering Sciences*, volume 60, pages 2479–2501.
- Horn, R. & Johnson, C. (2012) *Matrix Analysis*. Cambridge University Press, New York, NY, USA, 2nd edition.
- Hunt, G. & Linden, P. (1999) The fluid mechanics of natural ventilation – displacement ventilation by buoyancy-driven flows assisted by wind. *Buildg. and Environm.*, **34**(6), 707–720.
- Jalali, M. A., Khoshnood, A. & Alam, M.-R. (2015) Microswimmer-induced chaotic mixing. *J. Fluid Mech.*, **779**, 669–683.
- Juniper, M. (2010) Optimization with nonlinear adjoint looping. In *Int. Workshop on Nonnormal and Nonlinear Effects in Aero- and Thermoacoustics*, Munich, Germany.
- Kadoch, B., Kolomenskiy, D., Angot, P. & Schneider, K. (2012) A volume penalization method for incompressible flows and scalar advection-diffusion with moving obstacles. *J. Comp. Phys.*, **231**(12), 4365–4383.
- Kevlahan, N. & Ghidaglia, J.-M. (2001) Computation of turbulent flow past an array of cylinders using a spectral method with Brinkman penalization. *Eur. J. Mech. B/Fluids*, **20**(3), 333–350.
- Kirkpatrick, S., Gelatt, C. D. & Vecchi, M. P. (1983) Optimization by simulated annealing. *Science*, **220**(4598), 671–680.
- Lekien, F., Coulliette, C., Mariano, A., Ryan, E., Shay, L., Haller, G. & Marsden, J. (2005) Pollution release tied to invariant manifolds: A case study for the coast of Florida. *Physica D*, **210**(1-2), 1–20.
- Lin, C. & Chellappa, R. (1987) Classification of partial 2-D shapes using Fourier descriptors. *IEEE Trans. Pat. Anal. and Mach. Intel.*, pages 686–690.
- Lin, Z., Thiffeault, J.-L. & Doering, C. (2011) Optimal stirring strategies for passive scalar mixing. *J. Fluid Mech.*, **675**, 465–476.
- Linden, P. (1979) Mixing in stratified fluids. *Geophys. & Astrophys. Fluid Dyn.*, **13**(1), 3–23.
- Linden, P. (1999) The fluid mechanics of natural ventilation. *Annu. Rev. Fluid Mech.*, **31**(1), 201–238.
- Liu, Q. & Vasilyev, O. (2007) A Brinkman penalization method for compressible flows in complex geometries. *J. Comp. Phys.*, **227**(2), 946–966.
- Liu, W. (2008) Mixing enhancement by optimal flow advection. *SIAM J. Contr. Opt.*, **47**(2), 624–638.
- Long, R. R. (1978) A theory of mixing in a stably stratified fluid. *J. Fluid Mech.*, **84**(1), 113124.
- Madsen, J. I., Shyy, W. & Haftka, R. T. (2000) Response surface techniques for diffuser shape optimization. *AIAA J.*, **38**(9), 1512–1518.
- Marcotte, F. & Caulfield, C. (2018) Optimal mixing in two-dimensional stratified plane Poiseuille flow at finite Péclet and Richardson numbers. *J. Fluid Mech.*, **853**, 359385.
- Mathew, G., Mezic, I., Grivopoulos, S., Vaidya, U. & Petzold, L. (2007) Optimal control of mixing in Stokes fluid flows. *J. Fluid Mech.*, **580**, 261–281.
- Mattiusi, C. (2000) The finite volume, finite element, and finite difference methods as numerical methods for physical

- field problems. *Adv. Imag. and Elec. Phys.*, **113**, 1–146.
- Mohr, W., Saxton, R. & Jepson, C. (1957) Mixing in laminar-flow systems. *Ind. & Eng. Chemistry*, **49**(11), 1855–1856.
- Nguyen, N.-T. & Wu, Z. (2005) Micromixers: a review. *J. Micromech. Microeng.*, **15**(2), R1–R16.
- Ottino, J. (1990) Mixing, chaotic advection, and turbulence. *Annu. Rev. Fluid Mech.*, **22**(1), 207–254.
- Ottino, J., Leong, C., Rising, H. & Swanson, P. (1988) Morphological structures produced by mixing in chaotic flows. *Nature*, **333**(6172), 419.
- Ottino, J. M. (1989) *The kinematics of mixing: stretching, chaos, and transport*, volume 3. Cambridge University Press.
- Palacz, M., Smolka, J., Kus, W., Fic, A., Bulinski, Z., Nowak, A. J., Banasiak, K. & Hafner, A. (2016) CFD-based shape optimisation of a CO₂ two-phase ejector mixing section. *Appl. Thermal Eng.*, **95**, 62–69.
- Paul, E. L., Atiemo-Obeng, V. A. & Kresta, S. M. (2004) *Handbook of industrial mixing: science and practice*. John Wiley & Sons.
- Peltier, W. & Caulfield, C. (2003) Mixing efficiency in stratified shear flows. *Annu. Rev. Fluid Mech.*, **35**(1), 135–167.
- Pingen, G., Evgrafov, A. & Maute, K. (2007) Topology optimization of flow domains using the lattice Boltzmann method. *Struct. Multidisc. Optim.*, **34**(6), 507–524.
- Saglietti, C., Schlatter, P., Monokrousos, A. & Henningson, D. S. (2017) Adjoint optimization of natural convection problems: differentially heated cavity. *Theor. Comput. Fluid Dyn.*, **31**(5-6), 537–553.
- Saglietti, C., Schlatter, P., Wadbro, E., Berggren, M. & Henningson, D. S. (2018) Topology optimization of heat sinks in a square differentially heated cavity. *Int. J. Heat & Fluid Flow*, **74**, 36–52.
- Smolka, J. (2013) Genetic algorithm shape optimisation of a natural air circulation heating oven based on an experimentally validated 3-D CFD model. *Int. J. Thermal Sci.*, **71**, 128–139.
- Spencer, R. & Wiley, R. (1951) The mixing of very viscous liquids. *J. Coll. Sci.*, **6**(2), 133–145.
- Tang, W., Caulfield, C. & Kerswell, R. (2009) A prediction for the optimal stratification for turbulent mixing. *J. Fluid Mech.*, **634**, 487–497.
- Taylor, G. I. (1934) The formation of emulsions in definable fields of flow. *Proc. Roy. Soc. Lond. A*, **146**(858), 501–523.
- Thiffeault, J.-L. (2012) Using multiscale norms to quantify mixing and transport. *Nonlinearity*, **25**(2), R1.
- Uhl, V. (2012) *Mixing VI: Theory And Practice*. Elsevier.
- Vermach, L. & Caulfield, C. (2018) Optimal mixing in three-dimensional plane Poiseuille flow at high Péclet number. *J. Fluid Mech.*, **850**, 875–923.