A space–time integral minimisation method for the reconstruction of velocity fields from measured scalar fields

Jurriaan J. J. Gillissen^{1,†}, Alexandre Vilquin², Hamid Kellay²,
Roland Bouffanais³ and Dick K. P. Yue⁴

¹Center for Environmental Sensing and Modeling (CENSAM) IRG Singapore-MIT
Alliance for Research and Technology (SMART) Centre, Science Drive 2, 117543 Singapore

²Laboratoire Ondes et Matière d’Aquitaine (UMR CNRS 5798), Université de Bordeaux,
351 cours de la Libération, 33405 Talence, France

³Singapore University of Technology and Design, 8 Somapah Road, Singapore 487372, Singapore

⁴Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, Massachusetts 02139,
USA

Scalar image velocimetry (SIV) is the technique to extract velocity vectors from scalar field measurements. The usual technique involves minimising a cost functional, that penalises the deviation from the scalar conservation equation. This approach requires the measured scalar field to be sufficiently resolved and relatively noise free, such that space and time derivatives of the measured scalar field can be accurately evaluated. We quantify these requirements for a synthetic two-dimensional (2-D) turbulent flow field by evaluating the velocity reconstruction accuracy as a function of the temporal and spatial resolution and the noise level. We propose an improved SIV scheme, that reconstructs not only the velocity field but also the scalar field, which does not require approximating the space and time derivatives of the measured scalar field. Improved velocity reconstruction is demonstrated for the 2-D synthetic field. We furthermore apply the scheme to interferograms of the thickness field of a falling soap film, where 2-D turbulence is generated by an array of cylindrical obstacles. The statistics of the reconstructed velocity field are within 10 % of laser Doppler velocimetry measurements.

Key words: mathematical foundations, variational methods

1. Introduction

Scalar image velocimetry (SIV) is the reconstruction of the fluid velocity field $\mathbf{u}$ from measurements of a scalar field ψ , that is advected by $\mathbf{u}$. The technique is applied in weather forecasting models using e.g. satellite images of clouds or ocean temperature (see e.g. Kalnay 2003). SIV also finds applications in medical flow imaging and in experimental fluid mechanics. For instance in the laser induced

[†] Email address for correspondence: jurriaangillissen@gmail.com

fluorescence (LIF) technique a fluid is seeded with fluorescent molecules and laser light is focused into a thin sheet, where it is absorbed by the fluorescent molecules followed by spontaneous emission of light which is recorded by a camera (see e.g. Su & Dahm 1996). Assuming that the recorded light intensity is proportional to the fluorescence concentration ψ , the velocity field $\mathbf{u}$ can be reconstructed by invoking the scalar transport equation: $\partial_t \psi + \nabla \cdot (\mathbf{u}\psi) - \lambda \nabla^2 \psi = 0$, where λ is the scalar diffusivity. A direct inversion of the scalar transport equation only provides the component of $\mathbf{u}$ that is normal to ψ isolines: $\mathbf{u}^\perp = (-\partial_t \psi + \lambda \nabla^2 \psi) \nabla \psi / |\nabla \psi|^2$. Finding all components of $\mathbf{u}$ requires additional constraints, e.g. conservation of hydrodynamic variables. These constraints are naturally implemented using a variational technique, involving a minimisation of a cost functional $\mathcal{J}$. Previous works have used the following cost functional, which penalises the deviation from the scalar transport equation (see e.g. Su & Dahm 1996; Liu & Shen 2008; Papadakis & Mmin 2008; Corpetti *et al.* 2009):

$$\mathcal{J} = \frac{1}{2} \|\partial_t \psi + \nabla \cdot (\mathbf{u}\psi) - \lambda \nabla^2 \psi\|^2. \quad (1.1)$$

Throughout this work we assume that the norm $\|\cdot\|$ is based on the inner product $\langle \cdot, \cdot \rangle$, which, when applied to two vector (or scalar) fields $\mathbf{a}$ and $\mathbf{b}$ reads $\langle \mathbf{a}, \mathbf{b} \rangle = \int dV \mathbf{a} \cdot \mathbf{b}$. Reconstructing $\mathbf{u}$ using (1.1) is limited to cases where ψ is sufficiently smooth and well resolved in space and time, such that the space and time derivatives of the measured scalar field ψ can be accurately approximated. In the present work we alleviate these restrictions by formulating a variational scheme, that reconstructs not only the velocity field $\mathbf{u}$ but also the scalar field ϕ , which is not to be confused with the measured scalar field ψ . This scheme minimises the difference between ϕ and ψ :

$$\mathcal{J} = \frac{1}{2} \|\psi - \phi\|^2, \quad (1.2)$$

under the constraint of conserving fluid momentum, fluid mass and scalar field. As opposed to (1.1), equation (1.2) does not require approximating space and time derivatives of ψ . Therefore (1.2) is expected to perform better than (1.1), when the spatial or temporal resolution is low or the noise is high.

In this paper, we verify this hypothesis by comparing the velocity reconstruction based on (1.1) (Method 1) to that based on (1.2) (Method 2). In § 2 we formulate Method 1 and Method 2 and in § 3 we compare the accuracy of both Methods for a two-dimensional (2-D) synthetic, decaying, turbulent flow field. In § 4 we apply Method 2 to experimental data of the thickness field of a turbulent soap film. The statistics of the reconstructed velocity field are compared to laser Doppler velocimetry measurements. Conclusions are summarised in § 5.

2. Method derivation

2.1. Introductory remarks

In this work we formulate a new SIV method (Method 2) to reconstruct the space and time dependent velocity field $\mathbf{u}$, pressure field p and scalar field ϕ from noisy scalar field measurements ψ at discrete times $t = i\tau$, where i is an integer and τ is the time interval between measurements, which is referred to as the sampling time. The reconstruction method divides the time domain into segments, where the start $t_0 = (i-1)\tau$ and the end $t_1 = t_0 + \tau$ of the i th segment coincide with the i th and $(i+1)$ th scalar measurements. The reconstruction scheme solves a sequence of optimisation problems for the unknown state variable $\mathbf{w} = (\mathbf{u}, p, \phi)$ at the start of each segment,

i.e. at $t = t_0$. We use a subscript on a field variable to indicate a time instance, e.g. $\mathbf{w}_0 = \mathbf{w}(t_0)$. Finding $\mathbf{w}_0$ in each segment involves an iterative scheme, and the initial guess for the iteration is taken from the reconstructed field $\mathbf{w}_1$ at $t = t_1$ obtained in the preceding segment. It is noted that the segment time is arbitrarily chosen to be equal to the sampling time. Although beyond the scope of this work, using segment times larger than τ might improve the performance of the method.

2.2. Method 1

Before we derive Method 2, we first derive Method 1, which is a previously proposed SIV scheme (Papadakis & Mémin 2008; Corpetti *et al.* 2009). Method 1 reconstructs velocity and pressure, which corresponds to the state variable: $\mathbf{w} = (\mathbf{u}, p)$. In each time segment, Method 1 finds the initial conditions for the state variable $\mathbf{w}_0$, by minimising the deviation from the scalar conservation equation at the start t_0 as well as at the end t_1 of the segment. The corresponding cost functional for Method 1 reads:

$$\mathcal{J} = \frac{1}{2} \left(\sum_{i=0,1} \|\partial_t \psi_i + \nabla \cdot (\mathbf{u}_i \psi_i) - \lambda \nabla^2 \psi_i\|^2 + \kappa \|\nabla^2 \mathbf{u}_0\|^2 \right). \quad (2.1)$$

In (2.1) κ is a regularisation constant and the regularisation term $\kappa \|\nabla^2 \mathbf{u}_0\|^2$ is added to suppress the large wavenumber content of $\mathbf{u}$, which is important when the reconstruction method is confronted with noisy scalar measurement data and is prone to instability. Including regularisation terms is standard practice in inverse problems such as SIV (see e.g. Corpetti *et al.* 2006), and the mathematical theory of regularisation can be found in e.g. Tikhonov & Arsenin (1977). It is noted, that there is no need to add a pressure regularisation term, e.g. $\kappa \|\nabla^2 p_0\|^2$, since in the present methodology, p_0 is regularised due to its coupling to the regularised $\mathbf{u}_0$, via the equation of state, which is given by the incompressibility condition in (2.2) below.

It is noted, that in the hypothetical case of reconstructing flows with very large Reynolds numbers, one could resort to the ‘large eddy simulation’ technique, where the effect of the unresolved scales is modelled by adding dissipative terms to the equations of motion (equation (2.2) below). It would be interesting to investigate to what extent the regularisation terms in (2.1), could serve this purpose, and to relate the corresponding regularisation strength κ to the grid size. This however is beyond the present scope, which is restricted to the ‘direct numerical simulation’ technique, where the reconstructed hydrodynamic and scalar fields are fully resolved on the computational grid.

Equation (2.1) is minimised under the constraint, that the velocity field $\mathbf{u}$ and the pressure field p are governed by the Navier–Stokes equations and the continuity equation:

$$\mathbf{R}(\mathbf{w}) = \begin{pmatrix} \partial_t \mathbf{u} + \mathbf{u} \cdot \nabla \mathbf{u} + \nabla p - \nu \nabla^2 \mathbf{u} \\ \nabla \cdot \mathbf{u} \end{pmatrix} = \mathbf{0}. \quad (2.2)$$

Here ν is the fluid kinematic viscosity. Adding constraint (2.2) to (2.1) using the Lagrange multiplier $\hat{\mathbf{w}} = (\hat{\mathbf{u}}, \hat{p})$ results in the following constrained cost functional, which is referred to as the Lagrangian $\mathcal{L}$:

$$\mathcal{L} = \frac{1}{2} \left(\sum_{i=0,1} \|\partial_t \psi_i + \nabla \cdot (\mathbf{u}_i \psi_i) - \lambda \nabla^2 \psi_i\|^2 + \kappa \|\nabla^2 \mathbf{u}_0\|^2 \right) + \int_{t_0}^{t_1} \langle \hat{\mathbf{w}}, \mathbf{R}(\mathbf{w}) \rangle dt. \quad (2.3)$$

Reconstructing velocity fields from scalar fields

As ψ is discretely sampled in time, the time derivative $\partial_t \psi_i$ in (2.3) is approximate. In § 3.2 we study the effect of the sampling time on the reconstruction accuracy, where we use the second-, third- and fourth-order finite difference approximations for the time derivative of the measured scalar field:

$$\partial_t \psi_i = \tau^{-1} \begin{cases} \frac{1}{2} \psi_{i+1} - \frac{1}{2} \psi_{i-1}: & \text{second order} \\ -\frac{1}{6} \psi_{i+2} + \psi_{i+1} - \frac{1}{2} \psi_i - \frac{1}{3} \psi_{i-1}: & \text{third order} \\ -\frac{1}{12} \psi_{i+2} + \frac{2}{3} \psi_{i+1} - \frac{2}{3} \psi_{i-1} + \frac{1}{12} \psi_{i-2}: & \text{fourth order} \end{cases} \quad (2.4)$$

Spatial derivatives are computed using the Fourier collocation method, and in § 3.3, we study the effect of the spatial resolution, by artificially removing large wavenumbers from the measured scalar field.

Minimising $\mathcal{L}$ (equation (2.3)) with respect to $\mathbf{w}_0$ involves computing the gradient of $\mathcal{L}$ with respect to $\mathbf{w}_0$, i.e. $\delta \mathcal{L} / \delta \mathbf{w}_0$. To derive an expression for $\delta \mathcal{L} / \delta \mathbf{w}_0$, we start by writing the variation of the Lagrangian $\delta \mathcal{L}$ due to an infinitesimal variation in the state variable $\delta \mathbf{w} = (\delta \mathbf{u}, \delta p)$:

$$\begin{aligned} \delta \mathcal{L} = & \sum_{i=0,1} \langle \partial_t \psi_i + \nabla \cdot (\mathbf{u}_i \psi_i) - \lambda \nabla^2 \psi_i, \delta[\nabla \cdot (\mathbf{u}_i \psi_i)] \rangle + \kappa \langle \nabla^2 \mathbf{u}_0, \delta[\nabla^2 \mathbf{u}_0] \rangle \\ & + \int_{t_0}^{t_1} \langle \hat{\mathbf{w}}, \delta \mathbf{R}(\mathbf{w}) \rangle. \end{aligned} \quad (2.5)$$

We rewrite the various terms in (2.5) using integration by parts; see e.g. Gunzburger (2003):

$$\begin{aligned} \delta \mathcal{L} = & \langle -\psi_0 \nabla[\partial_t \psi_0 + \nabla \cdot (\mathbf{u}_0 \psi_0) - \lambda \nabla^2 \psi_0] + \kappa \nabla^4 \mathbf{u}_0 - \hat{\mathbf{u}}_0, \delta \mathbf{u}_0 \rangle \\ & + \langle -\psi_1 \nabla[\partial_t \psi_1 + \nabla \cdot (\mathbf{u}_1 \psi_1) - \lambda \nabla^2 \psi_1] + \hat{\mathbf{u}}_1, \delta \mathbf{u}_1 \rangle \\ & + \int_{t_0}^{t_1} \langle \hat{\mathbf{R}}(\mathbf{w}, \hat{\mathbf{w}}), \delta \mathbf{w} \rangle, \end{aligned} \quad (2.6)$$

where $\hat{\mathbf{R}}$ is the adjoint operator of the linearised version of $\mathbf{R}$ (equation (2.2)):

$$\hat{\mathbf{R}}(\mathbf{w}, \hat{\mathbf{w}}) = \begin{pmatrix} -\partial_t \hat{\mathbf{u}} - \mathbf{u} \cdot (\nabla \hat{\mathbf{u}} + \nabla \hat{\mathbf{u}}^T) - \nabla \hat{p} - \nu \nabla^2 \hat{\mathbf{u}} \\ -\nabla \cdot \hat{\mathbf{u}} \end{pmatrix}, \quad (2.7)$$

and where $\nabla \hat{\mathbf{u}}^T$ is the transposed of $\nabla \hat{\mathbf{u}}$. In (2.6) the $\hat{\mathbf{u}}_0$ and $\hat{\mathbf{u}}_1$ are time-boundary terms, that are obtained by integrating by parts the time derivative term in $\langle \hat{\mathbf{w}}, \delta \mathbf{R}(\mathbf{w}) \rangle$. For simplicity, we restrict ourselves to periodic domains, which negates the space-boundary terms, that are produced by integrating by parts the space derivative terms in $\langle \hat{\mathbf{w}}, \delta \mathbf{R}(\mathbf{w}) \rangle$.

From (2.6) we find, that the functional derivative of $\mathcal{L}$ (equation (2.3)) with respect to $\mathbf{w}_0 = (\mathbf{u}_0, p_0)$ equals:

$$\frac{\delta \mathcal{L}}{\delta \mathbf{w}_0} = \begin{pmatrix} \frac{\delta \mathcal{L}}{\delta \mathbf{u}_0} \\ \frac{\delta \mathcal{L}}{\delta p_0} \end{pmatrix} = \begin{pmatrix} -\psi_0 \nabla[\partial_t \psi_0 + \nabla \cdot (\mathbf{u}_0 \psi_0) - \lambda \nabla^2 \psi_0] + \kappa \nabla^4 \mathbf{u}_0 - \hat{\mathbf{u}}_0 \\ 0 \end{pmatrix}. \quad (2.8)$$

This expression for $\delta\mathcal{L}/\delta\mathbf{w}_0$ contains the Lagrange multiplier $\hat{\mathbf{w}} = (\hat{\mathbf{u}}, \hat{p})$. The evolution equation of this quantity as well as its initial and final conditions are found by demanding that $\delta\mathcal{L}$ (equation (2.6)) equals zero under arbitrary, infinitesimal $\delta\mathbf{w}$:

$$\hat{\mathbf{R}}(\mathbf{w}, \hat{\mathbf{w}}) = \mathbf{0}, \quad (2.9a)$$

$$-\psi_1 \nabla[\partial_t \psi_1 + \nabla \cdot (\mathbf{u}_1 \psi_1) - \lambda \nabla^2 \psi_1] + \hat{\mathbf{u}}_1 = \mathbf{0}, \quad (2.9b)$$

$$-\psi_0 \nabla[\partial_t \psi_0 + \nabla \cdot (\mathbf{u}_0 \psi_0) - \lambda \nabla^2 \psi_0] + \kappa \nabla^4 \mathbf{u}_0 - \hat{\mathbf{u}}_0 = \mathbf{0}. \quad (2.9c)$$

Equations (2.7) and (2.9a) govern the evolution of the Lagrange multiplier $\hat{\mathbf{w}}$, showing that $\hat{\mathbf{u}}$ is incompressible and it is advected by $\mathbf{u}$, and it is subjected to diffusion. Remarkably, the diffusion coefficient $-\nu$ of this transport equation is negative, and therefore this equation is integrated backward in time from $t = t_1$ to $t = t_0$. The ‘starting’ conditions $\hat{\mathbf{w}}_1$ at $t = t_1$ are given by (2.9b), while the ‘final’ conditions $\hat{\mathbf{w}}_0$ at $t = t_0$ define the optimisation update direction of $\mathbf{w}_0$ via (2.8). This direction approaches zero, when $\mathbf{w}_0$ reaches an extremum of $\mathcal{L}$, which corresponds to (2.9c).

Equation (2.9c) also shows that the regularisation term in the cost functional (equation (2.1)) effectively adds a pulse of hyper viscosity to the transport equations at $t = t_0$. The order of the hyper viscosity, which in this case equals four, depends on the exponent on the velocity gradient in the κ -term in (2.1), which in this case equals two. For a unit exponent we would have recovered normal, second-order viscosity (see e.g. Corpetti *et al.* 2006). Fourth-order viscosity is chosen above second-order viscosity however, since the former affects more selectively the large wavenumbers, while leaving small wavenumbers intact.

2.3. Method 2

Next we derive Method 2. In addition to velocity and pressure, which are reconstructed by Method 1, Method 2 also reconstructs the scalar field. The corresponding state variable reads: $\mathbf{w} = (\mathbf{u}, p, \phi)$. Consequently, Method 2 requires some 30% more computational effort than Method 1. In each time segment, Method 2 finds the initial conditions for the state variable $\mathbf{w}_0$, by minimising the deviation between the reconstructed scalar field ϕ and the measured scalar field ψ at the start t_0 as well as at the end t_1 of the segment. The corresponding cost functional for Method 2 reads:

$$\mathcal{J} = \frac{1}{2} \left(\sum_{i=0,1} \|\phi_i - \psi_i\|^2 + \kappa \|\nabla^2 \phi_0\|^2 + \kappa \|\nabla^2 \mathbf{u}_0\|^2 \right). \quad (2.10)$$

It is noted that the regularisation constant κ for the velocity field is chosen to be equal to that for the scalar field. This choice requires that $\nabla^2 \mathbf{u}_0$ is of the same order of magnitude as $\nabla^2 \phi_0$. When variables are non-dimensionalised, such that $\mathbf{u}_0 \sim \phi_0$, this condition is met when the smallest length scale of $\mathbf{u}_0$ is of the same order of magnitude as that of ϕ_0 . This requirement restricts the applicability of (2.10) to cases where the Schmidt number $Sc = \nu/\lambda \sim 1$. Extending the method to arbitrary Schmidt number requires having two distinct regularisation parameters.

Equation (2.10) is minimised under the constraint that $\mathbf{w}_1$ is related to $\mathbf{w}_0$ via the conservation equations of fluid momentum, fluid mass and scalar:

$$\mathbf{R}(\mathbf{w}) = \begin{pmatrix} \partial_t \mathbf{u} + \mathbf{u} \cdot \nabla \mathbf{u} + \nabla p - \nu \nabla^2 \mathbf{u} \\ \nabla \cdot \mathbf{u} \\ \partial_t \phi + \mathbf{u} \cdot \nabla \phi - \lambda \nabla^2 \phi \end{pmatrix} = \mathbf{0}. \quad (2.11)$$

Reconstructing velocity fields from scalar fields

Adding constraint (2.11) to (2.10) using the Lagrange multiplier $\hat{\mathbf{w}} = (\hat{\mathbf{u}}, \hat{p}, \hat{\phi})$ results in the following constrained cost functional, which is again referred to as the Lagrangian $\mathcal{L}$:

$$\mathcal{L} = \frac{1}{2} \left(\sum_{i=0,1} \|\phi_i - \psi_i\|^2 + \kappa \|\nabla^2 \phi_0\|^2 + \kappa \|\nabla^2 \mathbf{u}_0\|^2 \right) + \int_{t_0}^{t_1} \langle \hat{\mathbf{w}}, \mathbf{R}(\mathbf{w}) \rangle dt. \quad (2.12)$$

Minimising $\mathcal{L}$ with respect to $\mathbf{w}_0$ involves computing the gradient of $\mathcal{L}$ with respect to $\mathbf{w}_0$, i.e. $\delta\mathcal{L}/\delta\mathbf{w}_0$. To derive an expression for $\delta\mathcal{L}/\delta\mathbf{w}_0$ we start by writing the variation of the Lagrangian $\delta\mathcal{L}$ due to an infinitesimal variation of the state variable $\delta\mathbf{w} = (\delta\mathbf{u}, \delta p, \delta\phi)$:

$$\delta\mathcal{L} = \sum_{i=0,1} \langle \phi_i - \psi_i, \delta\phi_i \rangle + \kappa \langle \nabla^2 \mathbf{u}_0, \delta(\nabla^2 \mathbf{u}_0) \rangle + \kappa \langle \nabla^2 \phi_0, \delta(\nabla^2 \phi_0) \rangle + \int_{t_0}^{t_1} \langle \hat{\mathbf{w}}, \delta\mathbf{R}(\mathbf{w}) \rangle. \quad (2.13)$$

We rewrite the regularisation terms and the Lagrange multiplier term in (2.13) using integration by parts (see e.g. Gunzburger 2003):

$$\begin{aligned} \delta\mathcal{L} = & \langle \phi_1 - \psi_1 + \hat{\phi}_1, \delta\phi_1 \rangle + \langle \kappa \nabla^4 \phi_0 + \phi_0 - \psi_0 - \hat{\phi}_0, \delta\phi_0 \rangle \\ & + \langle \hat{\mathbf{u}}_1, \delta\mathbf{u}_1 \rangle + \langle \kappa \nabla^4 \mathbf{u}_0 - \hat{\mathbf{u}}_0, \delta\mathbf{u}_0 \rangle + \int_{t_0}^{t_1} \langle \hat{\mathbf{R}}(\mathbf{w}, \hat{\mathbf{w}}), \delta\mathbf{w} \rangle, \end{aligned} \quad (2.14)$$

where $\hat{\mathbf{R}}$ is the adjoint of the linearised version of $\mathbf{R}$ (equation (2.11)):

$$\hat{\mathbf{R}}(\mathbf{w}, \hat{\mathbf{w}}) = \begin{pmatrix} -\partial_t \hat{\mathbf{u}} - \mathbf{u} \cdot (\nabla \hat{\mathbf{u}} + \nabla \hat{\mathbf{u}}^T) - \nabla \hat{p} - \nu \nabla^2 \hat{\mathbf{u}} - \phi \nabla \hat{\phi} \\ -\nabla \cdot \hat{\mathbf{u}} \\ -\partial_t \hat{\phi} - \mathbf{u} \cdot \nabla \hat{\phi} - \lambda \nabla^2 \hat{\phi} \end{pmatrix}. \quad (2.15)$$

In (2.14) the $\hat{\mathbf{u}}_0$, $\hat{\mathbf{u}}_1$, $\hat{\phi}_0$ and $\hat{\phi}_1$ are time-boundary terms, that are produced by integrating by parts the time derivative terms in $\langle \hat{\mathbf{w}}, \delta\mathbf{R}(\mathbf{w}) \rangle$. It is again noted that we restrict ourselves to periodic domains, which negates the space-boundary terms, that are produced by integrating by parts the space derivative terms in $\langle \hat{\mathbf{w}}, \delta\mathbf{R}(\mathbf{w}) \rangle$. It is shown in §4, that, when confronted with non-periodic, scalar measurement data, the periodic reconstruction method performs surprisingly well, except for a small region near the boundaries.

From (2.14) we find, that the functional derivative of $\mathcal{L}$ (equation (2.12)) with respect to $\mathbf{w}_0 = (\mathbf{u}_0, p_0, \phi_0)$ equals:

$$\frac{\delta\mathcal{L}}{\delta\mathbf{w}_0} = \begin{pmatrix} \frac{\delta\mathcal{L}}{\delta\mathbf{u}_0} \\ \frac{\delta\mathcal{L}}{\delta p_0} \\ \frac{\delta\mathcal{L}}{\delta\phi_0} \end{pmatrix} = \begin{pmatrix} \kappa \nabla^4 \mathbf{u}_0 - \hat{\mathbf{u}}_0 \\ 0 \\ \kappa \nabla^4 \phi_0 + \phi_0 - \psi_0 - \hat{\phi}_0 \end{pmatrix}. \quad (2.16)$$

This expression for $\delta\mathcal{L}/\delta\mathbf{w}_0$ contains the Lagrange multiplier $\hat{\mathbf{w}}$. The evolution equation of this quantity as well as its initial and final conditions are found by demanding that $\delta\mathcal{L}$ (equation (2.14)) equals zero under arbitrary, infinitesimal $\delta\mathbf{w}$:

$$\hat{\mathbf{R}}(\mathbf{w}, \hat{\mathbf{w}}) = \mathbf{0}, \quad (2.17a)$$

$$\hat{\mathbf{u}}_1 = \mathbf{0}, \quad \phi_1 - \psi_1 + \hat{\phi}_1 = 0, \quad (2.17b)$$

$$\kappa \nabla^4 \mathbf{u}_0 - \hat{\mathbf{u}}_0 = \mathbf{0}, \quad \kappa \nabla^4 \phi_0 + \phi_0 - \psi_0 - \hat{\phi}_0 = 0. \quad (2.17c)$$

Equations (2.15) and (2.17a) govern the evolution of the Lagrange multiplier $\hat{\mathbf{w}} = (\hat{\mathbf{u}}, \hat{p}, \hat{\phi})$, showing that $\hat{\mathbf{u}}$ is incompressible and is forced along gradients of $\hat{\phi}$, and that both $\hat{\mathbf{u}}$ and $\hat{\phi}$ are advected by $\mathbf{u}$ and are subjected to diffusion. The diffusion coefficients $-\nu$ and $-\lambda$ of these transport equations are negative, and therefore these equations are integrated backward in time from $t=t_1$ to $t=t_0$. The ‘starting’ conditions $\hat{\mathbf{w}}_1$ at $t=t_1$ are given by (2.17b), while the ‘final’ conditions $\hat{\mathbf{w}}_0$ at $t=t_0$ define the optimisation update direction of $\mathbf{w}_0$ via (2.16). This direction approaches zero, when $\mathbf{w}_0$ reaches an extremum of $\mathcal{L}$, which corresponds to (2.17c).

2.4. Final remarks

The SIV schemes used in this work (Method 1 and Method 2) find $\mathbf{w}_0$ using the Polak–Rebiere variant of the conjugate gradient method (Polak 1971), which updates $\mathbf{w}_0$ along a search direction $\mathbf{h}$ related to $\delta\mathcal{L}/\delta\mathbf{w}_0$. The initial guess for $\mathbf{w}_0$ is $\mathbf{w}_1$ from the previous time segment, and the step length along $\mathbf{h}$ is varied using Brent’s line minimisation algorithm (Brent 2013), until the minimum of $\mathcal{J}$ (equation (2.1) for Method 1 or (2.10) for Method 2) in this direction is found. The conjugate gradient algorithm is continued until the relative change in $\mathcal{J}$ between two consecutive iterations drops below 0.01. Both Method 1 and Method 2 typically require $\sim 10^2$ conjugate gradient steps and ~ 10 Brent minimisation steps per conjugate gradient step. Therefore the computational effort of both methods is equivalent to that of $\sim 10^3$ computational fluid dynamics simulations.

3. Methods comparison

3.1. Set-up

In this section we compare the performance of Method 1 (equation (2.1)) and Method 2 (equation (2.10)) in reconstructing the velocity field $\mathbf{u}$ from a series of synthetic scalar fields ψ_i . In these tests, we know the exact equations that govern the measured scalar field, which in the ‘reconstruction literature’ is referred to as the ‘perfect model’ assumption. In §4 below, we analyse experimental data of soap film thickness fluctuations. In those tests, the governing equations are not entirely known, which adds an additional layer of uncertainty to the problem. In the present section we focus on the effects of the temporal resolution, i.e. the sampling time τ , the spatial resolution and the measurement noise. In Method 1 we use the second-, third- and fourth-order finite difference approximations for the time derivative of the measured scalar field $\partial_t \psi$ (equation (2.4)).

To compare their respective performances we apply both methods to a synthetic scalar field ψ in a two-dimensional (2-D) incompressible, decaying turbulent flow on a square, bi-periodic domain of size $L = 2\pi$. We generate a reference velocity field $\mathbf{v}$ and a reference scalar field ψ using numerical simulation of the hydrodynamic equations and the scalar transport equation in the absence of forcing mechanisms (equation (2.11)). Spatial derivatives in these equations are computed using the Fourier basis functions. Time integration is performed using the second-order explicit Adams–Bashforth scheme for the advection terms and the second-order implicit

FIGURE 1. Autocorrelation $\mathcal{C}$ of the synthetic fluid velocity and scalar fields as functions of the separation distance at $t = 1$.

Crank–Nicolson scheme for the diffusion terms. The number of grid points is $N^2 = 128^2$ and the numerical integration time step is $\Delta t = 1 \times 10^{-3}$. The reconstruction is applied to the time interval $0 < t < 4$. The initial reference velocity and scalar fields are constructed by assigning random numbers to the Fourier components of the wave vectors with absolute value $|\mathbf{k}| \leq 8$, while the remaining Fourier components are assumed zero. The initial velocity and scalar fields are normalised, such that $\mathcal{U} = \|\mathbf{v}\| = 1$ and $\|\psi\| = 1$ at $t = 0$. The diffusivity is $\lambda = 2.2 \times 10^{-3}$ and the viscosity is $\nu = 7.3 \times 10^{-4}$, which corresponds to a Reynolds number of $Re = \mathcal{U}L/\nu = 8.6 \times 10^3$ based on and the initial velocity scale $\mathcal{U}$ and a Schmidt number of $Sc = \nu/\lambda = 1/3$, indicating that the smallest velocity length scale is of the same order of magnitude as the smallest scalar length scale.

To quantify the size of the turbulent structures we compute the autocorrelation function $\mathcal{C}$, which is plotted for both $\mathbf{v}$ and ψ at $t = 1$ in figure 1. Defining the correlation length scale $\mathcal{L}$ as the distance where $\mathcal{C} = 1/2$, we observe from figure 1 that for both $\mathbf{v}$ as well as for ψ , $\mathcal{L} \sim 10^{-1}$ at $t = 1$. Given the velocity scale of $\mathcal{U} \approx 1$ we estimate the correlation time as $\mathcal{T} = \mathcal{L}/\mathcal{U} \sim 10^{-1}$. It is noted that in decaying turbulence $\mathcal{L}$ and $\mathcal{T}$ grow in time (not shown) and the above numbers are order of magnitude estimates. With time, the decaying flow field loses structure, and the spatial and temporal derivatives weaken. Therefore to stringently test both SIV methods, we restrict the reconstruction to $0 < t < 4$.

In § 3.2 we study the effect of the temporal resolution on the velocity reconstruction, by using varying sampling times τ . To illustrate the magnitude of a typical value for τ , we show in figure 2(a,b) two consecutive scalar field samples, using $\tau = 0.4$.

In § 3.3 we study the effect of the spatial resolution, by subjecting the synthetic scalar field to a low pass filter, with a varying cutoff wavenumber. Figure 2(c) shows a low pass filtered scalar field, where the cutoff wavenumber equals $k_c = |\mathbf{k}_c| = 48\pi/L$.

In § 3.4 we study the effect of scalar noise, by adding at each time step an independent random number to each Fourier component of the synthetic scalar field. The random numbers are drawn from a uniform distribution with zero mean and a standard deviation, that is η times the absolute value of the corresponding Fourier coefficient. Parameter η is referred to as the noise level. To illustrate the effect of noise, we show in figure 2(d) a synthetic scalar field, that is subjected

FIGURE 2. (a,b) Two consecutive samples at $t = 1$ (a) and $t = 1.4$ (b) of the synthetic scalar field, using a sampling time of $\tau = 0.4$. (c) Low pass filtered, synthetic scalar field, using a cutoff wavenumber of $|k_c|L/(2\pi) = 24$. (d) Noisy, synthetic scalar field, using a noise level of $\eta = 0.5$.

FIGURE 3. (a) Reconstruction error ϵ (equation (3.1)) versus time t , using noise level $\eta = 0$ and various sampling times τ , for Method 1 (filled markers) using regularisation constant $\kappa = 0$, and for Method 2 (open markers) using $\kappa = 0$ for $\tau < 0.1$ and $\kappa = 6 \times 10^{-7}$ for $\tau \geq 0.1$. Note, that the correlation time is estimated in §3.1 as: $\mathcal{T} \approx 0.1$. (b) ϵ versus τ at $t = 4$ using $\eta = 0$ for Method 1 using $\kappa = 0$, and for Method 2 using $\kappa = 0$ for $\tau < 0.1$ and $\kappa = 6 \times 10^{-7}$ for $\tau \geq 0.1$. The time derivative of ψ in the cost functional of Method 1 (equation (2.1)) is approximated using second-, third- and fourth-order finite difference schemes (equation (2.4)). The lines are drawn to guide the eye. (c) ϵ versus κ at $t = 4$ using Method 2, $\eta = 0$ and $\tau = 0.4$. These data show, that for $\tau = 0.4$, the unregularised Method 2 is unstable, and that the optimum regularisation strength equals $\kappa \sim 6 \times 10^{-7}$.

to a noise level of: $\eta \approx 0.5$ (50 %). It is shown in §3.4, that for very large η , i.e. $\eta \gtrsim 10$, the finite difference approximations in Method 1 become meaningless, and the reconstruction deteriorates. It is furthermore shown, that for $\eta \sim 1$, which is still a considerable amount of noise, Method 1 may produce accurate results. This behaviour demonstrates the merits of performing the space-time integral minimisation.

3.2. Effect of temporal resolution

We start by studying the accuracy of Method 1 in the absence of noise, i.e. $\eta = 0$, as a function of the sampling time τ , which has been varied between 0.025 and 3.2. Unless stated otherwise we use the second-order finite difference approximation (equation (2.4)) of $\partial_t \psi$ in (2.1). All the corresponding reconstruction runs were stable with zero regularisation strength $\kappa = 0$. Figure 3(a) shows for various values of τ the reconstruction error ϵ as a function of time.

$$\epsilon = \frac{\|\mathbf{u} - \mathbf{v}\|^2}{\|\mathbf{v}\|^2}, \quad (3.1)$$

where it is recalled, that $\mathbf{u}$ is the reconstructed velocity and $\mathbf{v}$ is the reference velocity field. It is noted, that for clarity, figure 3(a) only shows results up to $\tau = 0.4$, since results for larger τ fall on top of each other. It is seen in figure 3(a), that with time, the error first decreases and after a few time units reaches a steady value. Recalling the time segmentation described in § 2, this time dependent behaviour indicates, that the reconstruction depends on the quality of the initial guess for the initial condition at the start of each segment. In the first segment the initial guess is zero, while in each consecutive segment the initial guess gets closer to the ‘exact’ solution, explaining the observed decrease in ϵ with time.

Figure 3(b) shows the final and steady error ϵ , which is taken at $t=4$, as a function of the sampling time, in the range: $0.025 \leq \tau \leq 3.2$. The reconstruction error ϵ is seen to increase with τ . There are two causes for this behaviour. The first cause is related to the ill posedness of the initial value problem for chaotic systems, which corresponds to the cost functional developing multiple minima, when τ exceeds a threshold (see e.g. Pires, Vautard & Talagrand 1996). Figure 3(b) shows that in the present system this threshold is of the order of 0.1, which is in the same range as the estimated correlation time $\mathcal{T}$; see § 3.1. The second cause for the increase in ϵ with τ (figure 3b) is the error in the discretisation of $\partial_t \psi$ in (2.1). To study this effect we have repeated the reconstruction, using third- and fourth-order finite difference schemes (equation (2.4)). The results in figure 3(b) show only marginal improvements, by using higher-order time discretisation schemes. Therefore we use the second-order scheme in the remainder of this work.

Next, we evaluate the performance of Method 2 in reconstructing the velocity field from noise-free scalar fields ($\eta = 0$). We still focus on the effect of the sampling time τ , which is varied between 0.025 and 3.2. We found that in the current set-up, Method 2 is unstable for $\tau \gtrsim 0.1$. Since we employ the ‘direct numerical simulation’ technique, this instability is not triggered by a lack of energy dissipation due to unresolved scales. Instead, it is believed, that the instability is related to the ill posedness of the problem, as discussed above. In order to stabilise these cases we use regularisation κ -terms in (2.10). The optimal κ is determined by trial and error. Figure 3(c) shows the reconstruction error ϵ (equation (3.1)) for various values of the regularisation strength κ using $\tau = 0.4$. It is seen, that the optimum value is $\kappa \sim 6 \times 10^{-7}$, while for smaller values the method is unstable, and for larger values the error increases.

Figure 3(a,b) shows the reconstruction error ϵ of Method 2, alongside with that of Method 1. Figure 3(a) shows ϵ as a function of time t for various values of the sampling time τ , while figure 3(b) shows ϵ as a function of τ for a fixed value of $t = 4$. For Method 2 we have used $\kappa = 0$ for $\tau < 0.1$ and $\kappa = 6 \times 10^{-7}$ for $\tau \geq 0.1$. The comparison demonstrates the improved performance of Method 2 in the range $0.05 < \tau < 1$. It is moreover observed, that with increasing τ , the error of Method 2 continues to grow, and the corresponding reconstructed velocity becomes increasingly noisy and finally unstable. This behaviour is in contrast to that of Method 1, whose error saturates, and whose solution remains stable for large τ . We elaborate on the stability issue in § 3.4. It is noted, that for $\tau \gtrsim 1$, the error of Method 1: $\epsilon \sim 1$, which is smaller than $\epsilon \gtrsim 1$ for Method 2, but both reconstructions are too inaccurate to be useful.

3.3. Effect of spatial resolution

Next we consider the effect of the spatial measurement resolution, on the performance of Method 1 and Method 2. Since our numerical method operates in Fourier space,

FIGURE 4. Reconstruction error ϵ (equation (3.1)) versus low pass, cutoff wavenumber k_c , using sampling time $\tau = 0.025$, noise level $\eta = 0$ and regularisation strength $\kappa = 0$, for Method 1 and Method 2.

we vary the spatial resolution, by subjecting the synthetic scalar fields, to a low pass filter, with a variable cutoff wavenumber $k_c = |k_c|$. In physical space, this operation is equivalent to coarse graining, with a grain size of $\sim k_c^{-1}$. Since the unfiltered reference field has a spatial resolution of 128×128 , and the domain size is $L = 2\pi$, the maximum cutoff wavenumber is 64, which corresponds to no filtering. In the following, we use Method 1 and Method 2 to reconstruct the velocity field on the 128×128 grid, based on low pass filtered scalar fields, where k_c is varied between 8 and 64. We use a sampling time of $\tau = 0.025$, a noise level of $\eta = 0$ and a regularisation strength of $\kappa = 0$. Under these conditions, both methods provide equal accuracy, with $\epsilon \approx 10^{-2}$, when applied to the fully resolved ($k_c = 64$) synthetic scalar field; see figure 3(b).

In figure 4, we compare the reconstruction error of both methods, as a function of the spatial resolution of the measured scalar field, i.e. as a function of k_c . As expected the error of both methods increases with decreasing k_c , i.e. when removing flow features with increasing length scales. Perhaps unexpected, at low spatial resolution, Method 2 somewhat underperforms compared to Method 1. This result reflects that both methods are based on the same governing equations, and therefore they reconstruct the missing information in a similar fashion. It furthermore shows, that the Fourier approximation of the derivatives of resolved modes, is unaffected by the removal of unresolved modes. When the (resolved) modes are contaminated by noise however, then the discretisation errors do become important, and this is discussed in §3.4.

3.4. Effect of noise

We now turn our focus to the effect of the scalar noise. Figure 5(b) visualises the vorticity of the reconstructed flow field $\nabla \times \mathbf{u}$ using Method 1, $\tau = 0.1$, $\eta = 0.29$ and $\kappa = 0$. For comparison we show the reference solution in figure 5(a). It is noted that the vorticity fields in figure 5 vary in magnitude and to allow a proper visualisation we have used different mappings between the vorticity and the grey scale values in each of the panels.

FIGURE 5. Vorticity visualisation at $t = 4$ of the reference simulation (a) and of the reconstruction using $\tau = 0.1$ and $\eta = 0.29$ for Method 1 using $\kappa = 0$ (b), $\kappa = 1.9 \times 10^{-4}$ (c), $\kappa = 6.0 \times 10^{-3}$ (d), and for Method 2 using $\kappa = 6.0 \times 10^{-7}$ (e).

As expected figure 5(b) shows, that scalar noise results in velocity reconstruction noise. For the case of figure 5(b) the noise is concentrated at the small length scales. As an effect the large-scale flow structures, which are reasonably accurately reconstructed, are hardly visible in figure 5(b). To substantiate this point we plot in figure 6(a) the relative error spectrum ϵ_k :

$$\epsilon_k = \frac{\left| \int [\mathbf{u} - \mathbf{v}] \exp(-i\mathbf{k} \cdot \mathbf{x}) dV \right|^2}{\left| \int \mathbf{v} \exp(-i\mathbf{k} \cdot \mathbf{x}) dV \right|^2}. \quad (3.2)$$

This quantity measures the relative difference between the reconstruction and the reference velocity field, at wave vectors with an absolute value of k . The filled squares in figure 6(a) correspond to the case, that is visualised in figure 5(b). It is seen that the large scales are captured relatively well with: $\epsilon_k \sim 10^{-1}$. The small scales on the other hand are heavily contaminated by noise with: $\epsilon_k \sim 10^2$.

In order to suppress the large wavenumber noise in the reconstruction, a regularisation term is added to the cost functional (equation (2.1)). The value for the regularisation constant κ must be chosen with care as to suppress the small-scale noise, while leaving the large-scale flow structures intact. Figure 5(c,d) shows that with increasing κ , the small-scale noise diminishes but also the large-scale structures become less accurate. This is confirmed by the relative error spectra in figure 6(a), which shows that for intermediate $\kappa = 1.9 \times 10^{-4}$ (filled diamonds) the error is reduced over the entire spectrum. For large $\kappa = 6.0 \times 10^{-3}$ (filled triangles) on the other hand, although slightly better at the small scales, the error increases significantly at the large scales.

Figure 6(b) shows the corresponding reconstruction error ϵ (equation (3.1)) as a function of κ . The figure demonstrates that for $\kappa \sim 6 \times 10^{-4}$ the reconstruction error is minimum. For this optimum regularisation strength there is however still small-scale noise observed in the visualisation (figure 5c). Despite this the large-scale structures are distinguishable and resemble those of the reference field (figure 5a). Increasing κ beyond the optimum value results in a smooth velocity reconstruction, but also in a large ϵ . The corresponding large-scale structures (figure 5d) no longer resemble the reference field (figure 5a). It is noted that for severe noise levels, up to $\eta = 16$, Method 1 produced stable results, without regularisation, i.e. $\kappa = 0$.

The stability behaviour is different for Method 2, where instability is triggered by a relatively small amount of noise $\eta \gtrsim 0.1$. Method 2 can efficiently be stabilised however using the regularisation terms in (2.10). Again the regularisation constant κ

FIGURE 6. (a) Relative error spectrum ϵ_k (equation (3.2)) versus wavenumber k using sampling time $\tau = 0.1$ and noise level $\eta = 0.29$ for Method 1 (filled markers) using various regularisation constants κ , and for Method 2 (open markers) using $\kappa = 6.0 \times 10^{-7}$. (b) Reconstruction error ϵ (equation (3.1)) versus κ using Method 1, $\tau = 0.1$ and $\eta = 0.29$. The dashed line indicates the ϵ -value for $\kappa = 0$. (c) ϵ versus κ using Method 2, $\tau = 0.1$ and $\eta = 0.29$. (d) ϵ versus η using $\tau = 0.2$. Comparison between Method 1 using $\kappa = 6.0 \times 10^{-4}$ and Method 2 using $\kappa = 6.0 \times 10^{-7}$.

was found by trial and error. Figure 6(c) shows the corresponding ϵ as a function of κ using $\tau = 0.1$ and $\eta = 0.29$. It is observed that under these conditions, Method 2 is unstable for $\kappa \lesssim 10^{-7}$, while for $\kappa \gtrsim 10^{-7}$ the error is a non-monotonic function of κ with an optimum value at around $\kappa \sim 6 \times 10^{-7}$. It is noteworthy that this value is roughly the same as in §3.2, where the instability set in due to a relatively large sampling time τ . The corresponding vorticity visualisation (figure 5e) is free of noise and resembles the reference field (figure 5a) well.

The relative error spectrum ϵ_k in figure 6(a) shows the improved accuracy of Method 2 over Method 1, when confronted with measurement noise, and when using the optimum regularisation parameters. In this regard, it is noted that, as defined in (3.2), ϵ_k clearly visualises the k -dependence of the reconstruction error, which is hidden when directly comparing the energy spectra E_k (not shown):

$$E_k = \left| \int \mathbf{v} \exp(-i\mathbf{k} \cdot \mathbf{x}) d\mathbf{V} \right|^2. \quad (3.3)$$

It may seem peculiar, that the optimum κ for Method 1 (figure 6b) is three orders of magnitude larger than that for Method 2 (figure 6c). This difference can be explained

however by an order of magnitude analysis of the various terms in the cost functionals for Method 1 (equation (2.1)) and for Method 2 (equation (2.10)). As discussed in § 2 the regularisation terms $\kappa \|\nabla^2 \mathbf{u}\|^2$ and $\kappa \|\nabla^2 \phi\|^2$ are of the same order of magnitude provided that variables are properly scaled and the Schmidt number is $Sc = \nu/\lambda \sim 1$. To estimate the appropriate value for κ we balance the magnitude of the regularisation terms to that of the remaining term, which is $\|\partial_t \psi + \nabla \cdot (\mathbf{u}\psi) - \lambda \nabla^2 \psi\|^2$ in Method 1 and $\|\phi - \psi\|^2$ in Method 2. The ratio of κ in Method 1: κ_1 to that of Method 2: κ_2 is therefore proportional to the ratio of these remaining terms, i.e. $\kappa_1/\kappa_2 \sim \|\partial_t \psi + \nabla \cdot (\mathbf{u}\psi) - \lambda \nabla^2 \psi\|^2 / \|\phi - \psi\|^2$. Assuming that $\mathbf{u} \sim \psi \sim \phi \sim 1$ and that $\nabla \sim k_{max} = N/2$ we estimate $\kappa_1/\kappa_2 \sim (N/2)^2 \approx 4 \times 10^3$, where $N = 128$ is the number of grid points per dimension. This explains that Method 1 requires a regularisation strength that is three orders of magnitude larger than that of Method 2.

In figure 6(d) we compare the reconstruction error ϵ for Method 1 and Method 2 as a function of the noise level η , that has been varied between $3.1 \times 10^{-2} \leq \eta \leq 1.6 \times 10^1$. Here we use a sampling time of $\tau = 0.2$ and we have kept the regularisation constant fixed to $\kappa = 6 \times 10^{-4}$ for Method 1 and to $\kappa = 6 \times 10^{-7}$ for Method 2. It is noted that these κ values were optimised for $\tau = 0.1$ and $\eta = 0.29$ (see figure 6b,c), which are different from the parameters used in figure 6(d). It would be possible to derive optimum κ values for each data point in figure 6(d), but this would only marginally alter the results and has therefore been omitted. Figure 6(d) shows that for $\eta \lesssim 0.1$ the error in both methods is insensitive to η . For $\eta \gtrsim 0.1$ the error in Method 1 grows gradually and reaches a stable value for $\eta \gtrsim 10$. The error in Method 2, on the other hand, grows rapidly for $\eta \gtrsim 0.1$, and for this particular κ the solution becomes unstable for $\eta \gtrsim 1$. Although Method 1 is stable for large noise levels $\eta \gtrsim 10$, the reconstruction error of Method 1 equals: $\epsilon \sim 1$, under these conditions.

4. Application to soap film turbulence

Finally, we apply SIV to experimental data of the thickness field of a turbulent soap film. The experimental set-up is sketched in figure 7(a) and a photograph of the set-up is provided in figure 7(b). This set-up has previously been used to study two dimensional turbulence; see e.g. Greffier, Amarouchene & Kellay (2002). In the experiment, a 2 % detergent (Fairy Dreft) in water solution is injected at the top of a two-dimensional channel, that consists of two nylon wires, with a thickness of 0.5 mm. The channel has a straight section with a length of $l_2 = 1$ m and a width of $D = 0.06$ m. The wires are configured to form a diverging channel entrance and a converging channel exit, as to allow fluid injection and collection, respectively. As sketched in figure 7(a), three cylinders with a radius of 4 mm are placed in the film. The distance between the cylinder centres is 12 mm. The cylinders induce flow disturbances that evolve into 2-D turbulence; see figure 8(a). The downstream distance to the cylinders is denoted x and the distance to the centreline is denoted y . Film thickness and velocity measurements are taken between $x = 0.06$ and 0.45 m. In this region the flow speed is relatively constant and around $U \approx 2$ m s⁻¹. The average film thickness is estimated as $h_0 = Q/DU = 3$ μ m, where the volumetric flow rate is $Q = 3 \times 10^{-7}$ m³ s⁻¹.

Velocity fluctuations induce fluctuations in the film thickness $h(x, y, t)$. Experiments under similar conditions have shown, that these fluctuations are around $h_{rms} \approx 100$ nm (Greffier *et al.* 2002), which are small compared to the average thickness of $h_0 \approx 3$ μ m. It has been argued theoretically (Bruinsma 1995), as well as experimentally (Wu *et al.* 1995; Greffier *et al.* 2002), that the film behaves nearly as a 2-D

FIGURE 7. (a) Gravity g drives a $3\text{ }\mu\text{m}$ thick soap film with a velocity of around $U = 2\text{ m s}^{-1}$ between two wires, that form a vertical channel, with a diverging width, a constant width $D = 0.06\text{ m}$ and a converging width. The lengths of the respective sections are $l_3 = 0.7\text{ m}$, $l_2 = 1\text{ m}$ and $l_1 = 0.45\text{ m}$. Soap is collected at the channel exit, pumped upwards and injected at the channel entrance. Three cylinders are placed in the film to generate 2-D turbulence. The downstream distance to the cylinders and the wall normal distance to the channel centreline are denoted x and y , respectively. (b) Photograph of the set-up, showing the cylindrical obstacles, the LDV probe and the soap film.

incompressible fluid, and that the thickness field is a passive scalar, that is being advected by the flow, but with an unusual dissipative term:

$$\partial_t h + \mathbf{u} \cdot \nabla h = -\sigma \nabla^4 h. \quad (4.1)$$

Here σ is referred to as the thickness hyper diffusivity, which is conceived to be a function of the fluid viscosity and the fluid surface tension (Bruinsma 1995), and has an estimated value of $\sigma \sim 10^{-18}\text{ m}^4\text{ s}^{-1}$.

In the experiment, variations in h are visualised by illuminating the film with eight mercury lamps (Proxistar 85w 5500k) behind a diffuser and recording the reflected light using a high speed colour camera (Phantom v641), with an exposure time of $300\text{ }\mu\text{s}$, and a frame rate of 2000 Hz . Reflections from both sides of the film produce an interference pattern in the recorded image $\psi(x, y, t)$, as shown in figure 8(a,b). The pattern is sharpened by isolating the green channel from the recordings. The interference pattern is periodic, when h varies over half of the wavelength of the recorded light. However, since variations in h are fewfold smaller than the wavelength, this periodicity may be ignored. Therefore, ψ is assumed to be proportional to h , and to behave as a passive scalar.

Since the exact form of (4.1) is uncertain (Chomaz 2001; Auliel *et al.* 2015), and since Method 2 requires diffusive regularisation (κ -terms in (2.10)), we do not attempt to accurately reconstruct the dissipative range of the scalar energy spectrum. As an effect, we do not incorporate the fourth-order diffusion of (4.1) into Method 2, but instead, we replace $-\sigma \nabla^4 h$ in (4.1), with the usual, second-order diffusion $\lambda \nabla^2 h$, and apply Method 2 without further modification. The substitute diffusivity (λ in (2.11)) is chosen to be equal to the kinematic viscosity, which is assumed to be that of water $\lambda = \nu = 10^{-6}\text{ m}^2\text{ s}^{-1}$. Furthermore, to stabilise the method, we use a regularisation

FIGURE 8. (a) The soap film thickness field between $0.06 < x < 0.43$ m and $-0.027 < y < 0.027$ m is visualised by the interference pattern of reflected light. (b) The interference pattern in an interrogation window between $0.201 < x < 0.254$ m and $-0.027 < y < 0.027$ m. (c) The y component of the reconstructed velocity field in the interrogation window, that corresponds to the interference pattern, shown in figure 8(b).

strength of $\kappa = 1.2 \times 10^{-11} \text{ m}^4$. Despite these unphysical assumptions at the dissipative scales, it is shown below, that Method 2 correctly reproduces the statistics of the energy containing velocity fluctuations.

It is noted, that for this case, Method 1 (equation (2.1)) does not produce a reasonable velocity reconstruction (not shown). This is most likely due to inaccurate approximations of the spatial derivatives of the noisy, high frequency content of the interference patterns (figure 8b). To improve on these aspects, we have developed Method 2, which does not rely on computing spatial nor temporal derivatives of the measured scalar field. Therefore, the subsequent discussion focusses on Method 2.

We apply Method 2 to a square domain (interrogation window) with a size of 0.053 m. A snapshot of the recorded light intensity ψ in the window is shown in figure 8(b). The velocity and the scalar fields are reconstructed on a computational grid, where the number of grid points $N^2 = 180^2$ is one quarter of the number of pixels in the interrogation window, which corresponds to a grid spacing of $\Delta x = 2.94 \times 10^{-4}$ m. The computational time step is $\Delta t = 2 \times 10^{-5}$ s, which is one twenty-fifth of the sampling time $\tau = 5 \times 10^{-4}$ s.

Method 2 assumes periodic boundary conditions, which are not satisfied experimentally. In order to mitigate the detrimental effects of these unphysical conditions, the interrogation window moves downwards with a velocity of 1.88 m s^{-1} . In the moving window the turbulent eddies evolve without being advected by the mean flow, and can therefore be observed for a long period of time. Furthermore, the moving window minimises the influx of unknown information at the cross-stream boundaries.

FIGURE 9. (a) The standard deviation of the velocity components $u_{x,rms}$ and $u_{y,rms}$ for $y=0$ as functions of x . Comparison between SIV reconstruction (lines) and LDV measurement (markers). (b) The mean streamwise velocity component $\bar{u}_x$ for $y=0$ as a function of x . Comparison between SIV reconstruction (line) and LDV measurement (markers). (c) The reconstructed, mean streamwise velocity component $\bar{u}_x$ as a function of x and y . The cylinders induce a cross-stream variation of $\bar{u}_x$ with a local minimum on the centreline.

Figure 8(c) shows a snapshot of the y -component of the velocity reconstruction. It is seen that effects, due to the assumed periodicity of the experimental data, are confined to a thin (few pixels) region at the cross-stream boundaries. Apparently these effects propagate from the boundaries inward with a speed, that is small compared to the speed of the reconstruction.

The total recorded time is one second, which corresponds to 2000 frames. The interrogation window takes 0.17 s to move from the top to the bottom of the recorded image. This translation corresponds to a sequence of 340 images. From the 2000 recorded images, we extract a total of 30 such sequences, by varying the time, at which the window starts moving from the top to the bottom. This starting time is increased by 0.028 s, for each consecutive sequence, which corresponds to the translation of the window over its own size.

To validate Method 2 we compare the reconstructed velocity statistics to laser Doppler velocimetry (LDV) measurements of the streamwise x and wall normal y velocity components on the centreline $y=0$ as a function of the downstream position x . To this end, the fluid is dispersed with $1.1 \mu\text{m}$ polystyrene beads (SIGMA). Figure 9(a) shows the measured profiles of the standard deviations, i.e. turbulent intensities, of the streamwise velocity $u_{x,rms}$ (triangles) and wall normal velocity $u_{y,rms}$ (squares). The data show that $u_{y,rms}$ is 5% smaller than $u_{x,rms}$ and both quantities decrease as functions of x . In the same figure, we show the reconstructed $u_{x,rms}$ (solid line) and $u_{y,rms}$ (dashed line). These quantities are zero at the start of the reconstruction, which is at $x=0.06$ m, and increase until $x=0.2$ m, after which they decrease. The initial growth of the turbulent intensity reflects, that the reconstruction requires time to develop. As discussed in § 3.2, this development is related to the time segmentation, where the initial guess for the reconstruction of each segment is taken from the reconstruction of the previous segment. Therefore the reconstructed turbulent intensity is zero at $x=0.06$ m, increases with x , and approaches the converged profile at $x=0.2$ m. For $x > 0.2$ m the reconstructed $u_{y,rms}$ is 5% smaller than $u_{x,rms}$, and the profiles are within 5% of the experimental data.

In figure 9(b) we show the profile of the mean streamwise velocity component $\bar{u}_x$, obtained by LDV (squares). It is observed, that the mean flow accelerates from $\bar{u}_x = 1.6 \text{ m s}^{-1}$ at $x = 0.06$ m to $\bar{u}_x = 2.4 \text{ m s}^{-1}$ at $x = 0.43$ m. The corresponding, reconstructed velocity profile is also plotted (line) in figure 9(b). Since the interrogation window moves with a velocity of 1.88 m s^{-1} , the reconstructed mean

flow starts at $\bar{u}_x = 1.88 \text{ m s}^{-1}$ at $x = 0.06 \text{ m}$. With increasing x , the reconstructed $\bar{u}_x$ decreases and approaches the converged profile at $x = 0.15 \text{ m}$. After $x = 0.15 \text{ m}$, the reconstructed $\bar{u}_x$ follows the upward trend of the LDV data. There is a 10 % discrepancy between the reconstructed $\bar{u}_x$ and the LDV data. This may be attributed to steep variations of the mean flow in the y -direction (see figure 9c), in combination with non-converged Reynolds averages, slight off-centreline positioning of the LDV probe, as well as small differences in the positioning of the cylinders, and the flow speed during the video imaging and the LDV data acquisition.

5. Conclusions

We have analysed the performance of two scalar image velocimetry (SIV) methods, being an earlier proposed method (Method 1) and a new method (Method 2). Method 1 requires approximating the space and time derivatives of the measured scalar field ψ . Consequently when the scalar measurement resolution is relatively low, or when the scalar fields are relatively noisy, discretisation errors deteriorate the velocity reconstruction. Method 2 improves on these aspect by reconstructing not only the velocity field but also the scalar field, without approximating the space and time derivatives of ψ .

We have demonstrated the improvement of Method 2 over Method 1 for a synthetic, 2-D decaying turbulent flow field. The results show that Method 2 produces smaller reconstruction errors, especially when sampling rates are relatively low, and when there is considerable noise in the measurements. Despite being more accurate Method 2 is observed to be less stable than Method 1. Stability is however efficiently restored by adding regularisation terms to the cost functional, which effectively add pulses of hyper viscosity and hyper diffusivity to the conservation equations of momentum and scalar. Although Method 1 is stable without regularisation, a similar regularisation term is shown to improve the accuracy of Method 1, when confronted with noisy measurement data.

To further demonstrate Method 2, we analysed thickness field measurements in a gravity driven, turbulent soap film. The mean and standard deviations of the reconstructed velocity field are within 10 % of laser Doppler velocimetry measurements. The method therefore provides an alternative to particle image velocimetry, to experimentally obtain instantaneous velocity fields in turbulent soap films.

It is finally noted, that, even though we have focussed on 2-D test problems, the proposed Method 2 generalises to three spatial dimensions in a straightforward way. In fact, benefits of Method 2 over Method 1, are expected to be more pronounced in three than in two dimensions. This expectation is based on the notion that Method 2 requires smaller sampling rates than Method 1, and that volumetric sampling rates tend to be smaller than planar sampling rates. For instance acquiring 3-D laser induced fluorescence data requires time consuming sweeps of the laser sheet through the measurement volume (see e.g. Su & Dahm 1996) or 3-D tomographic imaging, which typically operates at a lower sampling rate than planar imaging.

Acknowledgements

This research is supported by the National Research Foundation Singapore under its Campus for Research Excellence and Technological Enterprise programme. The Center for Environmental Sensing and Modeling is an interdisciplinary research group of the Singapore MIT Alliance for Research and Technology. We also wish to

acknowledge the financial support from the Department of Mathematics of University College London.

REFERENCES

- AULIEL, M. I., CASTRO, F., SOSA, R. & ARTANA, G. 2015 Gravity-driven soap film dynamics in subcritical regimes. *Phys. Rev. E* **92** (4), 043009.
- BRENT, R. P. 2013 *Algorithms for Minimization Without Derivatives*. Courier Corporation.
- BRUINSMA, R. 1995 Theory of hydrodynamic convection in soap films. *Physica A* **216** (1–2), 59–76.
- CHOMAZ, J.-M. 2001 The dynamics of a viscous soap film with soluble surfactant. *J. Fluid Mech.* **442**, 387–409.
- CORPETTI, T., HÉAS, P., MÉMIN, E. & PAPADAKIS, N. 2009 Pressure image assimilation for atmospheric motion estimation. *Tellus A* **61** (1), 160–178.
- CORPETTI, T., HEITZ, D., ARROYO, G., MEMIN, E. & SANTA-CRUZ, A. 2006 Fluid experimental flow estimation based on an optical-flow scheme. *Exp. Fluids* **40** (1), 80–97.
- GREFFIER, O., AMAROUCHENE, Y. & KELLAY, H. 2002 Thickness fluctuations in turbulent soap films. *Phys. Rev. Lett.* **88** (19), 194101.
- GUNZBURGER, M. D. 2003 *Perspectives in Flow Control and Optimization*. SIAM.
- KALNAY, E. 2003 *Atmospheric Modeling, Data Assimilation and Predictability*. Cambridge University Press.
- LIU, T. & SHEN, L. 2008 Fluid flow and optical flow. *J. Fluid Mech.* **614**, 253–291.
- PAPADAKIS, N. & MÉMIN, É. 2008 Variational assimilation of fluid motion from image sequence. *SIAM J. Imaging Sci.* **1** (4), 343–363.
- PIRES, C., VAUTARD, R. & TALAGRAND, O. 1996 On extending the limits of variational assimilation in nonlinear chaotic systems. *Tellus A* **48** (1), 96–121.
- POLAK, E. 1971 *Computational Methods in Optimization: A Unified Approach*. Academic.
- SU, L. K. & DAHM, W. J. 1996 Scalar imaging velocimetry measurements of the velocity gradient tensor field in turbulent flows. I. Assessment of errors. *Phys. Fluids* **8** (7), 1869–1882.
- TIKHONOV, A. N. & ARSENIN, V. Y. 1977 *Solutions of Ill-Posed Problems*. Winston.
- WU, X. L., MARTIN, B., KELLAY, H. & GOLDBURG, W. I. 1995 Hydrodynamic convection in a two-dimensional Couette cell. *Phys. Rev. Lett.* **75** (2), 236–239.