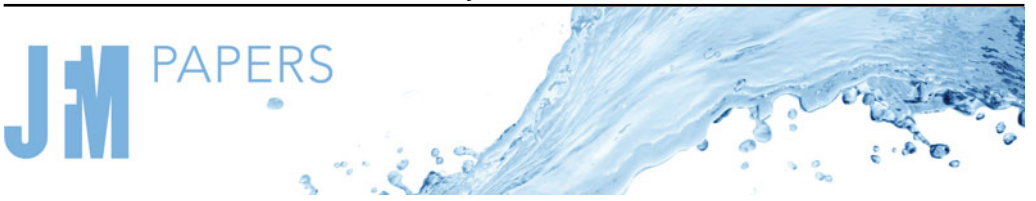
J. Fluid Mech. (2026), *vol.* 1038, A12, doi:10.1017/jfm.2026.11722



Searching for invariant solutions to wall-bounded flows using resolvent-based optimisation

Thomas Burton[1], **Sean Symon**[1] **and Davide Lasagna**[1]

[1]Faculty of Engineering and Physical Sciences, University of Southampton, University Road, Southampton SO17 1BJ, UK

Corresponding author: Thomas Burton, t.p.burton@soton.ac.uk

(Received 3 November 2025 UTC; revised 6 April 2026 UTC; accepted 18 May 2026 UTC)

We present a robust optimisation framework for computing invariant solutions of wall-bounded flows by recasting the Navier–Stokes equations as a variational problem as established in Ashtari & Schneider (*J. Fluid Mech.*, vol. 977, 2023, A7). The approach minimises the residual of the governing equations over a finite-time horizon, seeking periodic or equilibrium solutions. A novel contribution is made by including a Galerkin projection onto a basis of divergence-free modes that satisfy the no-slip boundary conditions. This projection not only makes the variational framework applicable to wall-bounded flows but it also yields a low-order representation of the dynamics. The basis is derived from resolvent analysis, which provides an orthonormal set. We demonstrate the method on a streamwise invariant formulation of rotating plane Couette flow, obtaining exact equilibrium and periodic solutions consistent with direct numerical simulations. The conditioning of the optimisation problem is analysed in detail, showing that convergence rates depend on the stability properties of the targeted solutions. Finally, we highlight a direct link between the conditioning of the optimisation and the structure of the resolvent operator, suggesting a unifying perspective on both the efficiency of the optimisation and the dynamical significance of resolvent modes.

Key words: nonlinear dynamical systems, variational methods

1. Introduction

Turbulence is often treated as a stochastic process. Progress in its study has largely come from statistical descriptions that exploit the self-similarity across different scales, yielding a universal picture of small-scale dynamics. Such approaches, however, offer little insight

© The Author(s), 2026. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

into the geometry-dependent large-scale structures or the physical mechanisms that govern turbulence more generally. An alternative view of turbulence, motivated by the theory of chaos for low-dimensional dissipative systems, was proposed by Hopf (Hopf 1942, 1948). In this view, the state of a fluid is asymptotically confined to a finite-dimensional invariant subspace embedded within the infinite-dimensional state space of the flow. The dimension of this invariant subspace can be much smaller than the infinite-dimensional nature of the state would suggest, owing to the strong dissipative power of viscosity. As the Reynolds number increases, the dimension of the invariant subspace usually increases allowing for more complex temporal evolution of the flow and requiring more degrees of freedom to accurately describe the resulting smaller-scale motions observed. However, understanding this invariant subspace is quite difficult as a result of its complex fractal structure indicative of strange attractors (Ruelle & Takens 1971; Aizawa 1982).

For low-dimensional chaotic systems, the complex structure of strange attractors has been shown in Auerbach *et al.* (1987) to be determined by the set of unstable periodic orbits (UPOs) embedded within the attractor, providing a 'skeleton' for the dynamics. It was shown in Cvitanović (1988) and Cvitanović (1995) that building these UPOs up in the form of a weighted sum can yield the ergodic properties of the dynamics on the strange attractor (Artuso, Aurell & Cvitanovic 1990). This method is called cycle expansion theory. Analogously, exact nonlinear solutions to the Navier–Stokes equations have been numerically shown to exist, dubbed exact coherent structures (ECS) due to their close resemblance to coherent structures observed in experiment and direct numerical simulation (DNS) studies. The exact nature of these ECS is not yet fully understood, however, there is evidence that turbulent flows repeatedly shadow ECS for a finite amount of time before moving away along one of the unstable manifolds of the solution towards a different ECS (Suri *et al.* 2020; Krygier, Pughe-Sanford & Grigoriev 2021; Crowley *et al.* 2022). It was also demonstrated (Yalnız *et al.* 2021; Page *et al.* 2024) that using ECS in a cycle expansion approach yields accurate statistics of the turbulent flow. The results do however generally rely on a tuning of the relative contribution of each ECS as opposed to deriving them exactly according to cycle expansion theory (Wang *et al.* 2025). The numerical computation of ECS was first shown in Nagata (1990) and Nagata (1997) where equilibrium and travelling wave solutions to the plane Couette flow were computed. Since then there has been regular work to expand the sets of solutions for various flows (Waleffe 1998, 2001; Faisst & Eckhardt 2003; Wedin & Kerswell 2004; Viswanath 2007; Pringle & Kerswell 2007; Gibson, Halcrow & Cvitanovic 2008; Waleffe 2009; Itano & Generalis 2009; Wedin, Bottaro & Nagata 2009; Okino *et al.* 2010; Uhlmann, Kawahara & Pinelli 2010; Willis, Cvitanović & Avila 2013; Nagata, Song & Wall 2021). A general review of the relevance of invariant solutions to fluid turbulence can be found in Kawahara *et al.* (2012).

The numerical methods employed to find ECS are generally divided into local and global approaches. Local methods, primarily shooting algorithms, are initialised with a flow snapshot that is evolved forward over the current period, with the mismatch between initial and final states used to update both the initial condition and the period. Examples include Christiansen, Cvitanovic & Putkaradze (1997), Tomoaki & Sadayoshi (2001), Kawahara & Kida (2001), Sánchez *et al.* (2004), van Veen, Kida & Kawahara (2006), Cvitanović & Gibson (2010) and Chandler & Kerswell (2013). Shooting algorithms are typically combined with a Newton–Krylov iteration, where the Jacobian system is solved by the generalised minimal residual (GMRES) method (Saad & Schultz 1986) to avoid explicit formation and prohibitive memory costs. However, this approach is highly sensitive to initial conditions due to the chaotic nature of the dynamics. Robust methods to generate initial guesses include close recurrences generated from time integration (Viswanath 2007; Azimi, Ashtari & Schneider 2022; Redfern, Lazer & Lucas 2024),

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

project-then-search (Ahmed & Sharma 2020) or Reynolds continuation from known solutions (Ashtari & Schneider 2023), all of which become more difficult to use for longer periods and higher-dimensional problems. To address this, multiple-shooting methods (Christiansen *et al.* 1997) replace the single trajectory with shorter segments matched simultaneously, improving conditioning and enabling parallel computation (Sánchez *et al.* 2004). Further robustness is achieved by adding a hookstep to the Newton–GMRES iteration (Dennis & Schnabel 1996; Viswanath 2007, 2009), which, when combined with multiple-shooting methods, represents the most effective way of computing invariant solutions in fluid flows. Nevertheless, challenges remain as the system size grows (Veen, Vela-Martín & Kawahara 2019).

Global methods instead act on a temporally extended field that already obeys the time-periodicity constraint, rather than on a set of sequential field snapshots obtained from time integration. This space–time field is then iteratively modified until its spatial variation and temporal evolution together satisfy the governing equations to within a prescribed tolerance. The Newton flow method introduced by Lan & Cvitanović (2004) formulates the problem through variational dynamics where the solution residual decays exponentially (in the limit of continuous deformation) to invariant solutions, with an over-relaxation factor mitigating the erratic behaviour of standard Newton iterations. Unlike shooting methods, these global approaches avoid exponential trajectory divergence and are generally more robust. However, the Newton flow method is not naturally matrix free, as each iteration requires forming and solving an $N \times N$ Jacobian system, with N the degrees of freedom for both space and time. This makes the approach prohibitively costly for high-dimensional systems (Fazendeiro *et al.* 2010). Practical implementations must therefore employ GMRES, as in local methods. Despite these advantages, the method still suffers from sensitivity to initial guesses owing to the limited convergence radius of Newton iterations.

Fourier–Galerkin (or harmonic-balance) formulations offer an efficient global approach for computing time-periodic invariant solutions without long integrations. By projecting the Navier–Stokes equations onto a truncated Fourier basis in time, the problem reduces to a nonlinear algebraic system for the harmonic coefficients. Recent work has demonstrated the effectiveness of this strategy for periodic and quasi-periodic flows (Sierra-Ausin *et al.* 2022), its extension to adjoint-based sensitivity and stability analysis (Sierra *et al.* 2021), and its application to nonlinear input–output analysis capturing triadic energy transfers (Rigas, Sipp & Colonius 2021). These Fourier–Galerkin approaches provide a complementary route to invariant solutions, trading time marching for a frequency-domain nonlinear solve while naturally incorporating harmonic interactions and sensitivities.

An alternative global approach is to recast the search for invariant solutions as an optimisation problem rather than root finding. Here, the objective functional measures the total violation of the governing equations over the spatio-temporal domain. Introduced by Farazmand (2016) for equilibrium and travelling wave solutions in two-dimensional (2D) Kolmogorov flow, this method is formulated as the adjoint of the Newton descent of Lan & Cvitanović (2004). Crucially, it is naturally matrix free and, therefore, well suited to high-dimensional systems. The variational optimisation approach demonstrated markedly greater robustness to initial conditions, enabling convergence from a wider set of guesses and uncovering previously unknown solutions. More recently, Azimi *et al.* (2022) extended the framework to periodic flows, and Ashtari & Schneider (2023) applied it to finding equilibrium in wall-bounded problems by enforcing no-slip constraints through the influence matrix (IM) method (Kleiser & Schumann 1980), which maintains consistency with the pressure-Poisson equation and avoids boundary errors arising from the lack of an explicit pressure condition.

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

Despite its advantages, the variational optimisation method still faces challenges. While the IM method provides a clear strategy for equilibria in wall-bounded domains, its extension to temporally periodic fields has not yet been published in the literature, leaving open the problem of a general approach for periodic solutions in such settings. Moreover, the improved robustness compared with Newton-based methods comes at the cost of losing quadratic convergence. Convergence is initially rapid but slows considerably near the minimum, largely due to the use of gradient descent, which is notoriously inefficient in this regime (Nocedal & Wright 2006). Algorithms incorporating curvature information could improve performance. Farazmand (2016) addressed this by adopting a hybrid strategy, initially using optimisation to approach the solution after which Newton–GMRES–hookstep is employed to accelerate convergence. More recently, Ashtari & Schneider (2023) improved efficiency of variational optimisation by employing an extrapolation technique with dynamic mode decomposition. Nevertheless, the method remains hampered by a significantly slower convergence rate than Newton-based methods, particularly as problem dimension increases. Hence, it is desirable to develop modifications that improve convergence without switching methods, while retaining compatibility with complex no-slip geometries.

In this work the variational optimisation framework is extended with a Galerkin projection, providing a general methodology for wall-bounded periodic flows with no-slip boundaries. The form of this problem closely resembles the general harmonic-balance method described in Sierra-Ausin *et al.* (2022). This projection enforces incompressibility and no-slip constraints while decoupling pressure and velocity. The basis is constructed with resolvent modes, which form a divergence-free orthonormal set satisfying the boundary conditions (BCs). Our approach is similar to Li & Lasagna (2026), who used a projection-based optimisation for lid-driven cavity flow using modes derived from spectral proper orthogonal decomposition (SPOD). In this work, however, resolvent modes replace these SPOD modes, removing the need for fully resolved flow data and enabling the construction of exact solutions rather than reduced-order models. Resolvent analysis, a widely used tool for stability, control and modelling (Gayme 2010; Garnaud *et al.* 2013; Gómez *et al.* 2016; Beneddine *et al.* 2017; Symon *et al.* 2018; Gayme & Minnick 2019; Jin, Illingworth & Sandberg 2022), has been shown to provide dynamically significant bases for unstable solutions (Sharma *et al.* 2016; Rosenberg & McKeon 2019) and to enable new invariant solutions via 'projection-then-search' (Ahmed & Sharma 2020). It has also been applied to reduced-order modelling, e.g. Taylor vortex flow in Rosenberg & Mckeon (2019) and Barthel, Zhu & McKeon (2021). Our work builds on these ideas by employing resolvent modes within a variational framework, where the Galerkin projection both enforces boundary constraints and effectively preconditions the optimisation. In this setting, modal truncation is used not merely to approximate turbulence, as in Barthel *et al.* (2021), but to enable computation of exact invariant solutions while simultaneously improving the conditioning of the optimisation and thereby accelerating convergence. To further address the slow convergence inherent to variational optimisation, we utilise two strategies. First, we replace the gradient-descent schemes of Farazmand (2016), Azimi *et al.* (2022), Ashtari & Schneider (2023) with quasi-Newton algorithms such as the limited-memory BFGS (L-BFGS) and the conjugate gradient, which incorporate curvature information. We demonstrate a substantially faster convergence rate near the minimum compared with gradient descent. Second, we exploit the typically low-dimensional representation of ECS using resolvent modes (Sharma *et al.* 2016) to truncate the basis, restricting the optimisation to a smaller space. It is then demonstrated that the convergence rate of the optimisation is improved via this truncation, a result that is motivated by bounding the conditioning of the optimisation problem to the singular values

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

of the resolvent operator. This dual use of the projection framework thus addresses both the difficulty of enforcing no-slip constraints in wall-bounded flows and the slow convergence that has limited the practical utility of variational optimisation.

The remainder of the paper is organised as follows. Section 2 introduces the variational optimisation methodology and the Galerkin projection used to enforce incompressibility and no-slip BCs. Section 3 presents the derivation of the resolvent modes forming the projection basis. Section 4 introduces the two-dimensional three-component (2D3C) formulation of rotating plane Couette flow (RPCF) and features a basic analysis of the behaviour of the flow at various Reynolds number regimes. Section 5 then demonstrates the projected variational optimisation methodology using resolvent analysis, featuring equilibrium and periodic solutions. The underlying mechanics governing the conditioning of the optimisation is then discussed in § 6, providing a novel link between resolvent analysis and the convergence rate related to the formulation discussed in Mons & Marquet (2021). Finally, § 7 concludes with a discussion of the main results of this work and the possible avenues for future work.

2. Methodology

In § 2.1 the variational optimisation methodology for wall-bounded flows is described. The methodology is specialised for planar Couette flow. In § 2.2 the Galerkin projection is introduced as a way to solve issues that arise due to the presence of no-slip BCs. A more general derivation and discussion of the implications can be found in Burton (2025).

2.1. *Variational optimisation*

Consider a domain bounded by two parallel plates, denoted as Ω, with a single inhomogeneous wall-normal direction (y) normalised with respect to the half-channel height denoted as h, and the streamwise and spanwise directions (x and z, respectively) that possess periodic BCs. The unit vectors in each direction are denoted as $\boldsymbol{e}_x$, $\boldsymbol{e}_y$ and $\boldsymbol{e}_z$. The periodic BCs model the effect of an infinite domain in these directions that result in statistically homogeneous turbulence. The domain can then be defined as $\Omega = [0,\ L_x] \times [-1,\ 1] \times [0,\ L_z]$, with L_x and L_z denoting the streamwise and spanwise length of the domain, respectively, normalised with respect to the channel height. The velocity $\boldsymbol{u}$ is a three-dimensional (3D) vector field defined over the 3D space Ω, and the pressure field, denoted with p, is a scalar field over the same domain. The flow in this domain is governed by (non-dimensional) incompressible Navier–Stokes equations and continuity equation

$$\frac{\partial \boldsymbol{u}}{\partial t} = -\boldsymbol{\nabla} p - (\boldsymbol{u} \cdot \boldsymbol{\nabla})\, \boldsymbol{u} + \frac{1}{Re} \boldsymbol{\Delta} \boldsymbol{u}, \tag{2.1a}$$

$$\boldsymbol{\nabla} \cdot \boldsymbol{u} = 0, \tag{2.1b}$$

where $Re = U_w h/\nu$. With U_w the speed of the walls and ν the kinematic viscosity of the fluid. The Laplace operator of functions in this domain is defined as $\boldsymbol{\Delta} = \partial^2/\partial x^2 + \partial^2/\partial y^2 + \partial^2/\partial z^2$. No-slip BCs are imposed at the walls:

$$\boldsymbol{u}|_{y=\pm 1} = \pm \boldsymbol{e}_x. \tag{2.2}$$

The velocity field $\boldsymbol{u}$ is also defined over a finite-time horizon fixed by the orbit period T, thus, $t \in [0,\ T)$. The combination of the spatial and temporal domains is defined as

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

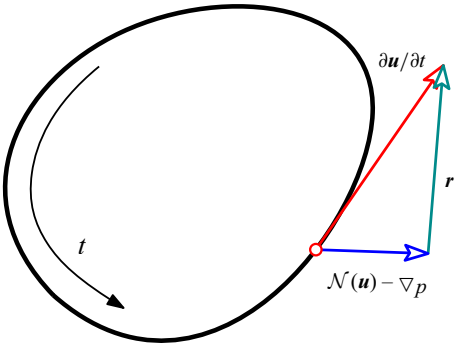


Figure 1. Schematic of an arbitrary loop in state space that does not satisfy the governing equations as its tangent vector $\partial \boldsymbol{u}/\partial t$ is not aligned with the evolution operator $\mathcal{N} = -(\boldsymbol{u} \cdot \nabla)\boldsymbol{u} + (1/Re)\boldsymbol{\Delta u}$.

$\Omega_t = \Omega \times [0,\ T)$. The inner product on this space is defined as

$$\langle \boldsymbol{u},\ \boldsymbol{v} \rangle_{\Omega_t} = \int_0^T \int_0^{L_z} \int_{-1}^1 \int_0^{L_x} \boldsymbol{u} \cdot \boldsymbol{v} \mathrm{d}x \mathrm{d}y \mathrm{d}z \mathrm{d}t. \tag{2.3}$$

The operator '$\cdot$' denotes the Euclidean dot product between two finite vectors. For real-valued vectors, this is given as $\boldsymbol{u} \cdot \boldsymbol{v} = \boldsymbol{u}^\top \boldsymbol{v}$, where $(\cdot)^\top$ denotes the transpose of the vector and, for complex-valued vectors, it is $\boldsymbol{u} \cdot \boldsymbol{v} = \boldsymbol{u}^\dagger \boldsymbol{v}$, where $(\cdot)^\dagger$ denotes the conjugate transpose of the vector. This inner product induces an associated norm $\|\boldsymbol{u}\|_{\Omega_t} = \sqrt{\langle \boldsymbol{u},\ \boldsymbol{u} \rangle_{\Omega_t}}$. Thus, any velocity field that has a finite norm under this definition is an element of the underlying Hilbert space, denoted with χ, and has finite kinetic energy. We now define the space of all incompressible flow fields, periodic in time, and that obey the spatial BCs as follows

$$\mathcal{P}_T = \left\{ \boldsymbol{u} \in \chi \ :\ \nabla \cdot \boldsymbol{u} = 0,\ \boldsymbol{u}|_{t=0} = \boldsymbol{u}|_{t=T}\ ,\ \boldsymbol{u}|_{y=\pm 1} = \pm \boldsymbol{e}_x,\ \boldsymbol{u} \text{ obeys periodic BCs} \right\}. \tag{2.4}$$

The elements of $\mathcal{P}_T$ are state-space loops that do not necessarily obey (2.1*a*). In general, T is not known *a priori* and should be included as part of the search for a solution. The goal is to derive a method that converges to elements of the subset of $\mathcal{P}_T$ that obey (2.1*a*). This can be done by introducing a mapping $\mathcal{R}\colon \mathcal{P}_T \to \mathbb{R}_{\geq 0}$ such that $\mathcal{R}[\boldsymbol{u}] = 0$ if and only if $\boldsymbol{u} \in \mathcal{P}_T$ and $\boldsymbol{u}$ satisfies (2.1*a*); otherwise, $\mathcal{R} > 0$. This is equivalent to defining an objective function whose global minima ($\mathcal{R} = 0$) correspond to periodic solutions of (2.1*a*) that are divergence-free and obey the no-slip and periodicity BCs, called here the global residual. To obtain the global residual, first a local residual that quantifies the violation of the Navier–Stokes equations at every point in Ω_t is defined as

$$\boldsymbol{r} = \frac{\partial \boldsymbol{u}}{\partial t} + \nabla p + (\boldsymbol{u} \cdot \nabla \boldsymbol{u}) - \frac{1}{Re} \boldsymbol{\Delta u}. \tag{2.5}$$

Figure 1 shows a representation of the local residual defined on all points of a state-space loop representing a periodic flow. The local residual spans the distance between the rate of change of a state-space loop and the forcing imposed by the Navier–Stokes equations. The task of finding periodic solutions to (2.1*a*) in the space $\mathcal{P}_T$ can now be stated in terms of the following optimisation problem:

$$\min_{\boldsymbol{u} \in \mathcal{P}_T,\, T} \quad \mathcal{R}\,[\boldsymbol{u}] = \frac{1}{2} \|\boldsymbol{r}\|^2_{\Omega_t}. \tag{2.6}$$

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

Geometrically, minimising $\mathcal{R}$ is equivalent to continuously deforming a state-space loop as depicted in figure 1 such that the vectors $\partial \boldsymbol{u}/\partial t$ and $\boldsymbol{\nabla} p + (\boldsymbol{u} \cdot \boldsymbol{\nabla} \boldsymbol{u}) - (1/Re)\boldsymbol{\Delta} \boldsymbol{u}$ align as closely as possible, while being constrained to the linear subspace that define the BCs and incompressibility conditions. This process terminates when the rate of change of the state vector and the right-hand side of (2.1*a*) are as closely aligned as possible.

To be able to solve this optimisation problem in practice, the gradient of $\mathcal{R}$ with respect to the field $\boldsymbol{u}$ is required, referred to as the functional derivative and denoted as $\delta \mathcal{R}/\delta \boldsymbol{u}$. This expression can be obtained by adding a perturbation to the velocity field denoted as $\boldsymbol{v}$ such that $\boldsymbol{u} \to \boldsymbol{u} + \epsilon \boldsymbol{v}$ where $\epsilon \in \mathbb{R}$. Substituting this into the definitions of the global residual relates the change in the residual, called the first variation of $\mathcal{R}$ and denoted as $\delta \mathcal{R}$, to the desired residual gradient by

$$\delta \mathcal{R} = \left\langle \frac{\delta \mathcal{R}}{\delta \boldsymbol{u}},\ \epsilon \boldsymbol{v} \right\rangle_{\Omega_t}. \tag{2.7}$$

Linearising the local residual in (2.5) and taking the adjoint of the resulting linearised operators gives the closed-form expression for the residual gradient

$$\frac{\delta \mathcal{R}}{\delta \boldsymbol{u}} = -\frac{\partial \boldsymbol{r}}{\partial t} - (\boldsymbol{u} \cdot \boldsymbol{\nabla})\, \boldsymbol{r} + (\boldsymbol{\nabla} \boldsymbol{u})^{\top} \boldsymbol{r} - \frac{1}{Re} \boldsymbol{\Delta} \boldsymbol{r} - \boldsymbol{\nabla} q, \tag{2.8}$$

with the additional constraints on the local residual

$$\boldsymbol{\nabla} \cdot \boldsymbol{r} = 0, \tag{2.9a}$$

$$\boldsymbol{r}|_{y=\pm 1} = 0, \tag{2.9b}$$

as well as $\boldsymbol{r}$ obeying the periodic BCs. The details of the derivation for (2.8) and the source of the extra constraints in (2.9) are given in Appendix A.1.

To implement a gradient-based optimisation of (2.6) it is also necessary to have a gradient of $\mathcal{R}$ with respect to T. In this work, however, the fundamental frequency, defined as $\omega = 2\pi/T$, is preferred. The fundamental frequency represents the lowest frequency oscillation that is permissible in the finite-time window defined by the period T. The gradient of $\mathcal{R}$ with respect to ω is given by

$$\frac{\partial \mathcal{R}}{\partial \omega} = \frac{1}{\omega} \left\langle \frac{\partial \boldsymbol{u}}{\partial t},\ \boldsymbol{r} \right\rangle_{\Omega_t}, \tag{2.10}$$

of which the detailed derivation is given in Appendix A.2.

The derivations described in Farazmand (2016), Azimi *et al.* (2022) and Ashtari & Schneider (2023) formulate the problem as a new set of dynamics that leads to taking the adjoint of the linearised dynamics. This is the primary reason that the method is typically known as the adjoint solver method rather than variational optimisation. The introduction of the new variable q in (2.8) is required to enforce the constraint that $\boldsymbol{\nabla} \cdot \delta \mathcal{R}/\delta \boldsymbol{u} = 0$, which ensures that as the optimisation progresses the flow remains incompressible. Since q performs an identical role to the pressure p in the primitive dynamics, we call it the adjoint pressure in line with its role in Ashtari & Schneider (2023). The expression (2.8) is the same as that derived in Ashtari & Schneider (2023), with the adjoint pressure q arising as a consequence of including the continuity equation as part of the constructed adjoint dynamics instead of an explicit constraint as it is treated here. Some of the general features of the derived variational/adjoint dynamics and its treatment in Ashtari & Schneider (2023) are discussed in Appendix B.

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

2.2. *Galerkin projection*

The difficulty in solving the optimisation problem in (2.6) comes from not having a simple way to compute the gradient in (2.8) while enforcing the constraints on the local residual given in (2.9). This stems from the lack of physical BCs for the pressures p and q. To provide a naïve example, if one initialises an optimisation with a candidate field $\boldsymbol{u} \in \mathcal{P}_T$ with period T that does not satisfy (2.1*a*), then to compute $\boldsymbol{r}$ in (2.5) it is required to first find the pressure field p. This would typically be done by solving a pressure-Poisson equation with Neumann BCs that ensures that $\boldsymbol{\nabla} \cdot \boldsymbol{r} = 0$. However, it does not guarantee $\boldsymbol{r}|_{y=\pm 1} = 0$ since the Dirichlet BCs are not satisfied by default, due to $\boldsymbol{u}$ not representing an actual solution to the Navier–Stokes equations. The result of this procedure is non-zero residual values at the wall, which invalidates the gradient $\delta\mathcal{R}/\delta\boldsymbol{u}$ as a guaranteed descent direction for $\mathcal{R}$, due to the constraints in (2.9) not being properly satisfied. In addition, the same problem exists for computing the adjoint pressure q, leading to a similar problem where it is very difficult to simultaneously enforce $\boldsymbol{\nabla} \cdot \boldsymbol{u} = 0$ and $\boldsymbol{u}|_{y=\pm 1} = \pm \boldsymbol{e}_x$. This leads to a velocity field at the next iteration not obeying the no-slip BCs. Put another way, if the system of equations that make up the variational dynamics is solved using the pressure-Poisson equations for p and q with standard Neumann BCs, then the no-slip BCs cannot be explicitly enforced and $\boldsymbol{u},\ \boldsymbol{r} \notin \mathcal{P}_T$.

Ashtari & Schneider (2023) proposed a methodology specific to equilibrium solutions that avoids this issue by solving for both the update to the velocity field, $\boldsymbol{u}$, and the pressure field, p, simultaneously in a coupled fashion using the IM method, which ensures that both $\boldsymbol{\nabla} \cdot \boldsymbol{r} = 0$ and $\boldsymbol{r}|_{y=\pm 1} = 0$ are satisfied. The same treatment is provided for $\boldsymbol{r}$ and q in a staggered approach, providing a physically consistent evolution of both $\boldsymbol{u}$ and $\boldsymbol{r}$. However, this method was developed to solve for equilibria ($\partial/\partial t = 0$), and it is not obvious how this would work for temporally varying fields because the method used to obtain compatible velocity and pressure fields is designed as an update method for some time-stepping scheme, while we employ a Fourier basis in time.

With the difficulties in evolving the gradient-based optimisation while satisfying all the constraints now clear, we propose to use a Galerkin projection onto an orthogonal basis that satisfies the BCs and the incompressibility constraints. Some setup is required to arrive at the final procedure that yields a valid computation of the residuals and gradient that obey the constraints. First, the velocity field is expanded into a sum over a basis as

$$\boldsymbol{u}\,(x,\ y,\ z,\ t) = \boldsymbol{u}_{BC}(y) + \sum_{m=1}^{\infty} \sum_{\boldsymbol{k} \in \mathbb{Z}^3} a_{\boldsymbol{k}m} \boldsymbol{\psi}_{\boldsymbol{k}m}\,(y)\, e^{i\boldsymbol{k}\cdot\boldsymbol{\xi}}, \tag{2.11}$$

where $\boldsymbol{k} = (k_x,\ k_z,\ k_t)$ is the integer wavenumber vector and $\boldsymbol{\xi} = (\alpha x,\ \beta z,\ \omega t)$ is the scaled direction vector. The coefficients $\alpha = 2\pi/L_x,\ \beta = 2\pi/L_z$ and $\omega = 2\pi/T$ are determined by the domain size and period, and each represent the smallest frequency in each of their respective direction that can be accommodated in the finite space. To account for the inhomogeneous BCs on $\boldsymbol{u}$ in (2.2), the steady base flow $\boldsymbol{u}_{BC}(y)$ is introduced into (2.11). In this work the laminar state will be used for the base flow. The modes $\boldsymbol{\psi}_{\boldsymbol{k}m} = \boldsymbol{\psi}_{\boldsymbol{k}m}(y)$ are divergence-free and obey the no-slip BCs, i.e.

$$\boldsymbol{\nabla}_{\boldsymbol{k}} \cdot \boldsymbol{\psi}_{\boldsymbol{k}m} = 0, \tag{2.12}$$

$$\boldsymbol{\psi}_{\boldsymbol{k}m}\big|_{y=\pm 1} = 0, \tag{2.13}$$

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

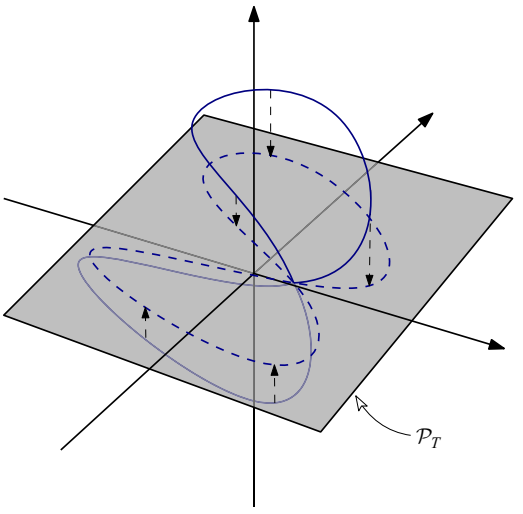


Figure 2. Schematic for a Galerkin projection of a state-space loop representing a velocity field onto the linear subspace in (2.4).

in addition to being orthonormal

$$\int_{-1}^{1} \boldsymbol{\psi}_{\boldsymbol{k}m} \cdot \boldsymbol{\psi}_{\boldsymbol{k}n} \mathrm{d}y = \delta_{mn}, \tag{2.14}$$

where δ_{mn} is the Kronecker delta. The operator $\boldsymbol{\nabla}_{\boldsymbol{k}} \cdot \boldsymbol{u}_{\boldsymbol{k}}(y) = i\alpha k_x \boldsymbol{u}_{\boldsymbol{k}} + \partial \boldsymbol{u}_{\boldsymbol{k}}/\partial y + i\beta k_z \boldsymbol{u}_{\boldsymbol{k}}$ is the divergence operator in spectral/Fourier space. The sum over the set of modes $\boldsymbol{\psi}_{\boldsymbol{k}m}$ is taken over a countably infinite set of modes for every frequency $\boldsymbol{k}$ to form a complete basis for the optimisation space $\mathcal{P}_T$. This includes the mean mode $\boldsymbol{k} = \boldsymbol{0}$, as this is part of the solution that needs to be determined during the optimisation. In practice, the sums in (2.11) are truncated to be finite. The key observation is that, for any combination of coefficients $a_{\boldsymbol{k}m} \in \mathbb{C}$, the velocity field is incompressible and obeys the no-slip BCs, restricting the velocity field to exist only within a finite-dimensional subspace of the complete function space.

The orthogonality of the modes allows for the following identity for the coefficients $a_{\boldsymbol{k}m}$ to be derived:

$$a_{\boldsymbol{k}m} = \frac{1}{2TL_zL_x} \left\langle \boldsymbol{u} - \boldsymbol{u}_{BC},\ e^{-i\boldsymbol{k}\cdot\boldsymbol{\xi}} \boldsymbol{\psi}_{\boldsymbol{k}m} \right\rangle_{\Omega_t}. \tag{2.15}$$

This projection is least squares, in the sense that the set of coefficients $a_{\boldsymbol{k}m}$ produces a trajectory that is the closest possible trajectory to $\boldsymbol{u}$ restricted to be within the linear subspace defined by $\boldsymbol{\psi}_{\boldsymbol{k}m}$. A low-dimensional schematic for such a projection is depicted in figure 2.

Next, a similar expansion can be performed for the local residual

$$\boldsymbol{r} = \sum_{m=1}^{\infty} \sum_{\boldsymbol{k}\in\mathbb{Z}^3} s_{\boldsymbol{k}m} \boldsymbol{\psi}_{\boldsymbol{k}m}(y)\, e^{i\boldsymbol{k}\cdot\boldsymbol{\xi}}. \tag{2.16}$$

Similar to (2.11), any combination of coefficients $s_{\boldsymbol{k}m} \in \mathbb{C}$ constructs a local residual field that obeys the constraints (2.9), making it the crucial step in this methodology. The coefficients $s_{\boldsymbol{k}m}$ can be computed in a similar way as the velocity:

$$s_{\boldsymbol{k}m} = \frac{1}{2TL_zL_x} \left\langle \boldsymbol{r},\ e^{-i\boldsymbol{k}\cdot\boldsymbol{\xi}} \boldsymbol{\psi}_{\boldsymbol{k}m} \right\rangle_{\Omega_t}. \tag{2.17}$$

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

If the expansions (2.11) and (2.16) are substituted into (2.8) and (2.5), and the results are projected using the same operation as in (2.15), we get the following expressions for the gradient of $\mathcal{R}$ with respect to the coefficients $a_{\boldsymbol{k}m}$, i.e.

$$\frac{\partial \mathcal{R}}{\partial a_{\boldsymbol{k}m}} = -ik_t \omega s_{\boldsymbol{k}m} - \frac{1}{2TL_zL_x}\left\langle (\boldsymbol{u}\cdot\nabla)\,\boldsymbol{r} - \left(\nabla \boldsymbol{u}^{\top}\right)\boldsymbol{r} + \frac{1}{Re}\Delta \boldsymbol{r},\ e^{-i\boldsymbol{k}\cdot\boldsymbol{\xi}}\boldsymbol{\psi}_{\boldsymbol{k}m}\right\rangle_{\Omega_t}, \tag{2.18}$$

and for the local residual coefficients $s_{\boldsymbol{k}m}$,

$$s_{\boldsymbol{k}m} = ik_t \omega a_{\boldsymbol{k}m} + \frac{1}{2TL_zL_x}\left\langle (\boldsymbol{u}\cdot\nabla)\,\boldsymbol{u} - \frac{1}{Re}\Delta \boldsymbol{u},\ e^{-i\boldsymbol{k}\cdot\boldsymbol{\xi}}\boldsymbol{\psi}_{\boldsymbol{k}m}\right\rangle_{\Omega_t}. \tag{2.19}$$

The key observation to be made on the expressions (2.18) and (2.19) is that the pressure gradient terms are not present. This is a result of the fact that the basis $\boldsymbol{\psi}_{\boldsymbol{k}m}$ is divergence-free, leading to the result

$$\left\langle \nabla p,\ e^{-i\boldsymbol{k}\cdot\boldsymbol{\xi}}\boldsymbol{\psi}_{\boldsymbol{k}m}\right\rangle_{\Omega_t} = \left\langle p,\ -\nabla\cdot\left(e^{-i\boldsymbol{k}\cdot\boldsymbol{\xi}}\boldsymbol{\psi}_{\boldsymbol{k}m}\right)\right\rangle_{\Omega_t} = \left\langle p,\ -e^{-i\boldsymbol{k}\cdot\boldsymbol{\xi}}\nabla_{\boldsymbol{k}}\cdot\boldsymbol{\psi}_{\boldsymbol{k}m}\right\rangle_{\Omega_t} = 0, \tag{2.20}$$

and similarly for the adjoint pressure gradient term ∇q.

Figure 3 is a flow diagram for the computations over a single iteration of a given optimisation loop to compute both $\mathcal{R}$ and $\partial\mathcal{R}/\partial a_{\boldsymbol{k}m}$. The computation begins with the coefficients $a_{\boldsymbol{k}m}$ that represent a given periodic and incompressible velocity field that obeys the periodicity and no-slip BCs, relative to a given set of modes $\boldsymbol{\psi}_{\boldsymbol{k}m}$ and laminar base flow $\boldsymbol{u}_{BC}$. The residual coefficients $s_{\boldsymbol{k}m}$ are then obtained by computing all the terms in (2.5), excluding the pressure gradient ∇p, the result of which is then projected onto the basis as in (2.17). The global residual can be obtained directly from $s_{\boldsymbol{k}m}$ by substituting (2.17) into (2.5) and evaluating the resulting integral, leading to the identity

$$\mathcal{R} = \frac{1}{2}\sum_{m=1}^{\infty}\sum_{\boldsymbol{k}\in\mathbb{Z}^3}|s_{\boldsymbol{k}m}|^2, \tag{2.21}$$

which is simply an application of Parseval's identity. To obtain the residual gradient, the velocity and residual coefficients are expanded back into physical space to facilitate the computation of all the terms on the right-hand side of (2.8), excluding the adjoint pressure gradient ∇q. Finally, the result is projected using (2.18) to obtain $\partial\mathcal{R}/\partial a_{\boldsymbol{k}m}$. With both $\mathcal{R}$ and its gradient, the chosen optimisation algorithm can be used to obtain an update to the coefficients $a_{\boldsymbol{k}m}$, denoted as $\delta a_{\boldsymbol{k}m}$ in figure 3. Convergence being determined by the size of $\mathcal{R}$. The process summarised here allows the direct computation of the required residual and gradient without resorting to solving a pressure-Poisson equation, avoiding the issues related to the BCs. The frequency gradient $\partial\mathcal{R}/\partial\omega$ does not require any extra special treatment as once the appropriate local residual $\boldsymbol{r}$ has been determined then (2.10) can be computed directly.

The wall-normal derivative terms in (2.18) and (2.19) can only be expressed in terms of the modes $\boldsymbol{\psi}_{\boldsymbol{k}m}$. Additionally, the computation of the nonlinear terms in projected space using $a_{\boldsymbol{k}m}$ and $s_{\boldsymbol{k}m}$ would require the evaluation of several convolution sums (Barthel *et al.* 2021) that scale poorly with the degrees of freedom of the discretised system. To avoid this, the pseudo-spectral method is used. The velocity and the residual are expanded back into their full physical space forms where the nonlinear terms are computed after which they are then projected onto the modes, replacing the heavy convolution sums with pointwise multiplication in physical space.

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

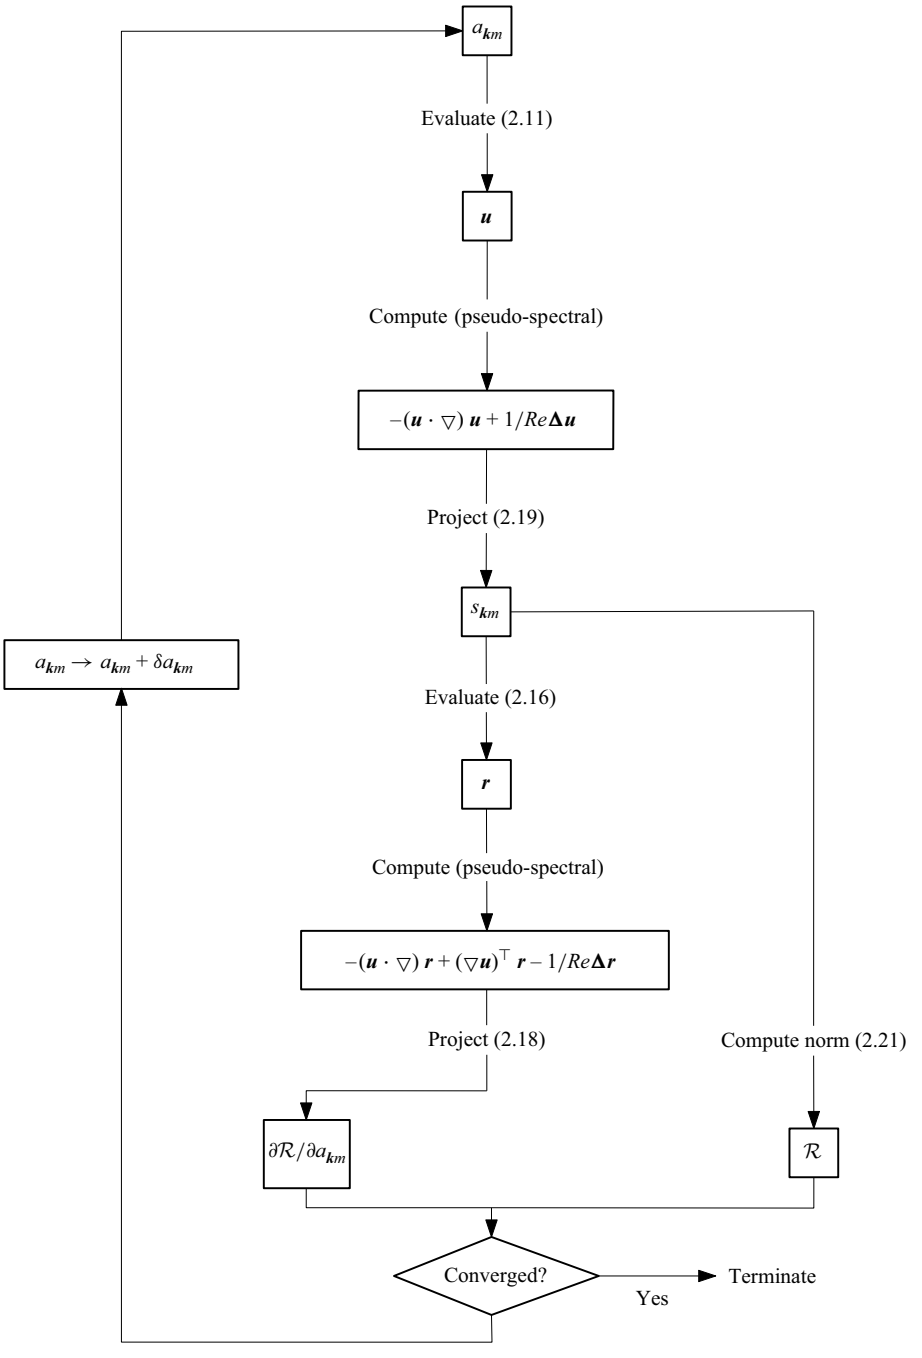


Figure 3. Flow diagram of the computations performed at each iteration of the optimisation.

3. Resolvent analysis

To implement the projected variational optimisation methodology, a method to generate the modes $\boldsymbol{\psi}_{\boldsymbol{km}}$ with the desired properties (2.12)–(2.14) is required. This choice is non-trivial with many possible solutions. For example, SPOD could be used to generate a set of modes from large data sets (Towne, Schmidt & Colonius 2018), or global stability analysis

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

can derive modes directly from the Navier–Stokes equations. Here resolvent analysis is used to generate the desired modes which form a complete basis as in (2.11) and (2.16). For completeness we provide an overview of resolvent analysis following closely the formulation provided in Sharma (2019).

First, decompose the velocity field into a steady base component and the fluctuation around this base flow as

$$\boldsymbol{u}\,(x,\ y,\ z,\ t) = \boldsymbol{u}_b\,(y) + \boldsymbol{u}'\,(x,\ y,\ z,\ t)\,, \tag{3.1}$$

where the base flow here is the flow about which the equations will be linearised, serving a distinct role from the base flow $\boldsymbol{u}_{BC}(y)$ used in (2.11). The standard choice in resolvent analysis is to take $\boldsymbol{u}_b$ to be the turbulent mean. The choice taken in this work is to use the laminar profile, primarily for pragmatic reasons, and is discussed further in §§ 5 and 6.

At this point it is useful to define the Leray projector, denoted as $\mathbb{P}$ (Temam 1984). This operator takes an instantaneous velocity field and orthogonally projects onto the divergence-free subspace, and can be derived from Helmholtz decomposition of the velocity field. It is used here as a notional short-hand which removes the pressure gradient term rendering the primitive formulation of the Navier–Stokes equation easier to manipulate. Applying the Leray projector to (2.1*a*) and then substituting in the decomposition (3.1) provides an evolution equation for the fluctuations

$$\frac{\partial \boldsymbol{u}'}{\partial t} = \mathbb{P}\left(\mathcal{N}\,(\boldsymbol{u}_b) + \mathcal{L}_{\boldsymbol{u}_b}\boldsymbol{u}' + \boldsymbol{f}\,(\boldsymbol{u}')\right). \tag{3.2}$$

The operator $\mathcal{N}(\boldsymbol{u})$ is the Navier–Stokes operator, given by the right-hand side of (2.1*a*). The operator $\mathcal{L}_{\boldsymbol{u}_b}$ is the linearised Navier–Stokes operator evaluated at the base flow $\boldsymbol{u}_b$. Finally, $\boldsymbol{f}(\boldsymbol{u}') = -(\boldsymbol{u}'\cdot\nabla)\boldsymbol{u}'$ is the nonlinear term for the fluctuations $\boldsymbol{u}'$. Physically $\boldsymbol{f}$ represents the nonlinear feedback that transports energy between scales and sustains any unsteady motion in the flow. The addition of the Leray projector allows us to ignore the pressure gradient terms in the operators $\mathcal{N}$ and $\mathcal{L}_{\boldsymbol{u}_b}$, as well as not needing to explicitly represent the continuity equation in (3.2).

Next we define the Fourier expansion of the velocity fluctuation $\boldsymbol{u}'$ in the homogeneous spatial and time directions as follows, following closely the formulation used in § 2.2:

$$\boldsymbol{u}'\,(x,\ y,\ z,\ t) = \sum_{\boldsymbol{k}\in\mathbb{Z}^3} \boldsymbol{u}'_{\boldsymbol{k}}\,(y)\,e^{i\boldsymbol{k}\cdot\boldsymbol{\xi}}. \tag{3.3}$$

Here $\boldsymbol{k}$ and $\boldsymbol{\xi}$ have the same definition as in § 2.2. The corresponding inverse operation to compute the Fourier coefficients is given by

$$\boldsymbol{u}'_{\boldsymbol{k}} = \frac{1}{8\pi^3}\int_0^{2\pi}\int_0^{2\pi}\int_0^{2\pi} \boldsymbol{u}' e^{i\boldsymbol{k}\cdot\boldsymbol{\xi}}\,\mathrm{d}\boldsymbol{\xi}. \tag{3.4}$$

By expanding (3.2) into its Fourier coefficients, the following is obtained for the fluctuations:

$$ik_t\omega\boldsymbol{u}'_{\boldsymbol{k}} = \mathbb{P}\left(\mathcal{L}_{\boldsymbol{k},\boldsymbol{u}_b}\boldsymbol{u}'_{\boldsymbol{k}} + \boldsymbol{f}_{\boldsymbol{k}}\right), \quad \boldsymbol{k}\in\mathbb{Z}^3\setminus\{\boldsymbol{0}\}\,. \tag{3.5}$$

Rearranging, one obtains the following linear relationship between the fluctuations and the nonlinear interactions resulting from the convective term:

$$\boldsymbol{u}'_{\boldsymbol{k}} = \boldsymbol{R}_{\boldsymbol{k}}\boldsymbol{f}_{\boldsymbol{k}}, \quad \boldsymbol{k}\in\mathbb{Z}^3\setminus\{\boldsymbol{0}\}\,. \tag{3.6}$$

The operator $\boldsymbol{R}_{\boldsymbol{k}} = (ik_t\omega\boldsymbol{I} - \mathbb{P}\mathcal{L}_{\boldsymbol{k},\boldsymbol{u}_b})^{-1}\mathbb{P}$ is the resolvent. The exact form of $\boldsymbol{R}_{\boldsymbol{k}}$ is given in Appendix C. The next step is to decompose the operator using a Schmidt decomposition, or equivalently, a singular value decomposition (SVD) when discretised, as follows:

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

$$\boldsymbol{R}_{\boldsymbol{k}}\left(\cdot\right)=\sum_{m=1}^{\infty}\sigma_{\boldsymbol{k}m}\boldsymbol{\psi}_{\boldsymbol{k}m}\left\langle\boldsymbol{\phi}_{\boldsymbol{k}m},\ \cdot\right\rangle. \tag{3.7}$$

This decomposition provides two sets of orthonormal modes, i.e.

$$\int_{-1}^{1}\boldsymbol{\psi}_{\boldsymbol{k}m}\cdot\boldsymbol{\psi}_{\boldsymbol{k}n}\mathrm{d}y=\delta_{mn}, \tag{3.8}$$

$$\int_{-1}^{1}\boldsymbol{\phi}_{\boldsymbol{k}m}\cdot\boldsymbol{\phi}_{\boldsymbol{k}n}\mathrm{d}y=\delta_{mn}, \tag{3.9}$$

ranked in order of the associated singular values $\sigma_m\geqslant\sigma_{m+1}\geq 0$ for all $i\in\mathbb{N}$. The modes $\boldsymbol{\psi}_{\boldsymbol{k}m}$ and $\boldsymbol{\phi}_{\boldsymbol{k}m}$ are the left and right singular modes and form a complete basis for the range (response) and domain (input) of the resolvent for every $\boldsymbol{k}\in\mathbb{Z}^3\setminus\{0\}$, respectively. Since the range of the resolvent is the space of divergence-free fluctuation velocity fields that obey the desired BC, the left singular modes, $\boldsymbol{\psi}_{\boldsymbol{k}m}$, also called the response modes, have these desired properties. Thus, the response modes $\boldsymbol{\psi}_{\boldsymbol{k}m}$ can be used as the basis for the Galerkin projection introduced in § 2.2.

To actually find exact solutions as discussed in § 2, a basis for the mean $\boldsymbol{k}=0$ is also required. The resolvent in (3.6) is only technically defined using the fluctuation equation at all non-zero frequencies, $\boldsymbol{k}\neq 0$. However, the authors found that the resolvent operator can regardless be evaluated at $\boldsymbol{k}=0$, which provided the required modes for the mean, completing the basis required to find an exact solution using the expansions (2.11) and (2.16). In other words, despite the technically undefined interpretation the authors found that taking the SVD of the resolvent at $\boldsymbol{k}=0$ yielded a valid orthonormal basis as desired. The exact physical interpretation of the basis used for the Galerkin projection is less important than its ability to provide a valid orthonormal basis in the space of divergence-free and no-slip fields. The literature discussing the evaluation of the resolvent at the mean frequencies and wavenumbers is extremely sparse, since it is not typically required as the mean is used as an input. The exception is Mons & Marquet (2021) in which resolvent analysis of the linearised RANS equations are used for the purpose of data assimilation. However, this does not directly apply to the formulation in this work as the resolvent in (3.6) is defined using the fluctuation equation not the RANS equation.

The rate at which the singular values σ_m decay determines how accurately the decomposition (3.7) can be represented with only a finite number of modes. It has been observed for fluid flows that the singular values decay rapidly (McKeon & Sharma 2010), meaning a low-rank approximation to the resolvent operator can be constructed with a partial sum of (3.7). This reduces the dimensionality of the amplification mechanisms in (3.6), only keeping the most significant contributions.

4. General features of RPCF

We now introduce RPCF as the main flow configuration being used as a test case for the variational optimisation methodology. Rotating plane Couette flow is defined on the same domain as the general planar wall-bounded flow in § 2.1 with the same Cartesian coordinate system. The channel has a height of $2h$, the top and bottom walls move in opposite directions at a speed of U_w, with the addition that the domain is rotated about the spanwise direction at a rotational rate of Ω. The (non-dimensional) governing equations for this flow are given by

$$\frac{\partial\boldsymbol{u}}{\partial t}=-\left(\boldsymbol{u}\cdot\boldsymbol{\nabla}\right)\boldsymbol{u}-\boldsymbol{\nabla}p+\frac{1}{Re}\boldsymbol{\Delta u}-Ro\left(\boldsymbol{e}_z\times\boldsymbol{u}\right), \tag{4.1a}$$

$$\boldsymbol{\nabla}\cdot\boldsymbol{u}=0, \tag{4.1b}$$

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

where the velocity is subject to the same BCs as the flow described at the beginning of § 2. The flow is characterised by two non-dimensional parameters, the Reynolds number Re as defined in § 2 as well as the rotation number Ro:

$$Ro = \frac{2\Omega h}{U_w}. \tag{4.2}$$

The rotation number is the ratio between the characteristic rotation and inertial forces of the system. In this work, only the rotation direction aligned with the shear is considered. This is known as anti-cyclonic rotation and acts to produce linear instabilities at finite Reynolds numbers (Tsukahara, Tillmark & Alfredsson 2010), as opposed to standard plane Couette flow that is linearly stable for all finite Reynolds numbers (Daviaud, Hegseth & Bergé 1992). The most interesting parameter regime of the rotation number in this study is $0 < Ro < 1$, as linear instabilities in the flow are present in this range. For any value outside this range for Ro, the flow is actually made more linearly stable due to the presence of system rotation (Lezius & Johnston 1976; Hiwatashi *et al.* 2007).

In this work we make an additional assumption that the flow is streamwise-independent ($\partial/\partial x = 0$) reducing it to a 2D3C formulation. The primary reason to use a 2D3C formulation is to reduce the dimensionality of the problem making the solutions easier to find while retaining some of the key features of the original dynamics. This constraint in fact makes the flow exactly analogous to 2D Rayleigh–Benard convection as shown in Eckhardt, Doering & Whitehead (2020). The analogy between rotating shear flows and convective thermal flows has been known for some decades (Chandrasekhar 1961; Yih 1965), and has been used to describe Taylor–Couette flows in turbulent regimes in relation to the Rayleigh–Benard convection (Eckhardt, Grossmann & Lohse 2007). The nature of some steady and streamwise-independent solutions to the RPCF at low Reynolds numbers (before bifurcating into more complex structures) is discussed in Nagata (2013) and Nagata *et al.* (2021). The streamwise-independence means that the domain size is given solely by the half-channel height h and the spanwise length L_z, with the aspect ratio defined as the ratio $\gamma = L_z/h$. In this work an aspect ratio of $\gamma = 4$ is used exclusively.

The flow starts with the stable laminar solution $\boldsymbol{u}(y) = y\boldsymbol{e}_x$. Linear stability analysis shows that the boundary of stability for the laminar flow has the relationship

$$Re_{crit} = \sqrt{\frac{107}{Ro\,(1 - Ro)}}, \tag{4.3}$$

given in Lezius & Johnston (1976). This initial bifurcation leads to a streamwise independent flow, and so it is also a valid relation for the 2D3C case. This implies that the laminar flow is most unstable for $Ro = 0.5$, which is the rotation number of choice for this work. At $Re_{crit} \approx 20.7$ the laminar solution bifurcates and becomes unstable. A new stable equilibrium solution is born out of this bifurcation that has the characteristic streamwise-independent rolls that are present throughout the total set of regimes accessible by varying Re.

We characterise the route to turbulence by analysing time series extrema of the kinetic energy from DNS data over a range of Reynolds numbers at $Ro = 0.5$, where the kinetic energy is defined as $\mathcal{K}(t) = (1/2)\|\boldsymbol{u}(t)\|_\Omega$. All DNS results in this work were obtained using a custom solver (Lasagna, Tutty & Chernyshenko 2016) that employs a vorticity–streamfunction formulation, details of which are provided in Appendix D. The results are shown in figure 4. The purpose of this analysis is to provide the necessary context for interpreting the results from the variational optimiser. A comprehensive study of the

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

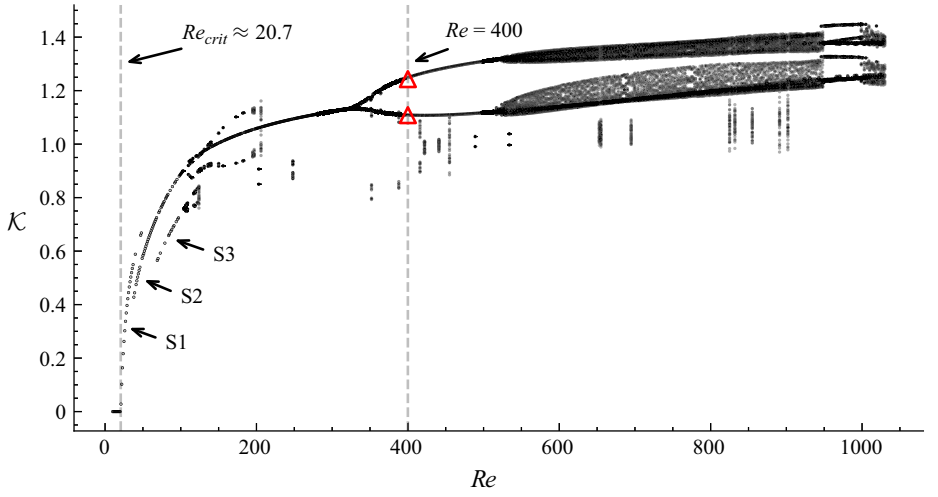


Figure 4. Bifurcation diagram of RPCF over a range of Reynolds numbers and $Ro = 0.5$ showing the transition from the stable laminar solution to turbulent flow. Also plotted are the corresponding kinetic energy extrema of the periodic solution obtained from optimisation at $Re = 400$ in § 5.3, denoted with red triangles.

transitionary behaviour of 2D3C RPCF lies beyond the present scope. For each Reynolds number, the DNS was initialised with a new random initial condition and then integrated until it converged to a final attractor, typically requiring thousands of time units. This procedure means that slightly different solutions may be obtained for different runs, depending on the initial condition.

The initial bifurcation from the laminar solution occurs at the expected value $Re \approx 20.7$, as predicted from (4.3). The kinetic energy of the resulting stable equilibrium grows rapidly as Re increases. Notably, instead of a single equilibrium branch being traced continuously, the DNS often settles on different attractors. These correspond to distinct stable equilibrium solutions, examples of which are labelled in figure 4. At $Re = 50$, three such branches, S1–S3, are shown in figure 5. They differ primarily in the number of streamwise rolls and exhibit kinetic energies of approximately $\mathcal{K} = 0.65$, 0.58 and 0.39 for S1, S2 and S3, respectively. The denser the rolls, the lower the kinetic energy, owing to higher dissipation at fixed Reynolds number. As Re increases, denser rolling structures are favoured, likely because sparser structures cannot sustain the required dissipation rates. All of these equilibria are linearly stable at $Re = 50$, since the DNS converges to them provided the initial condition lies within their basin of attraction. However, branches such as S3 are less commonly observed at this Reynolds number, likely because their basins of attraction are very small.

As Re increases further, each of these equilibria undergoes bifurcations to periodic solutions, sustained only over limited parameter ranges before seemingly losing stability and no longer acting as an attractor for the DNS. Around $Re = 200$, the S1 branch again emerges as the dominant equilibrium, before bifurcating near $Re = 350$ to a periodic orbit sustained up to a little above $Re = 500$. This orbit retains the large-scale streamwise roll structure of its parent equilibrium. Beyond $Re = 500$, the periodicity becomes increasingly irregular, with large amplitude fluctuations and a more broadband spectrum, indicating transition to turbulence. Finally, at the largest Reynolds numbers displayed in figure 4 the flow bifurcates again into a new apparently periodic solution before again transitioning to a chaotic evolution. These remaining bifurcations have not been investigated in any substantial way here.

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

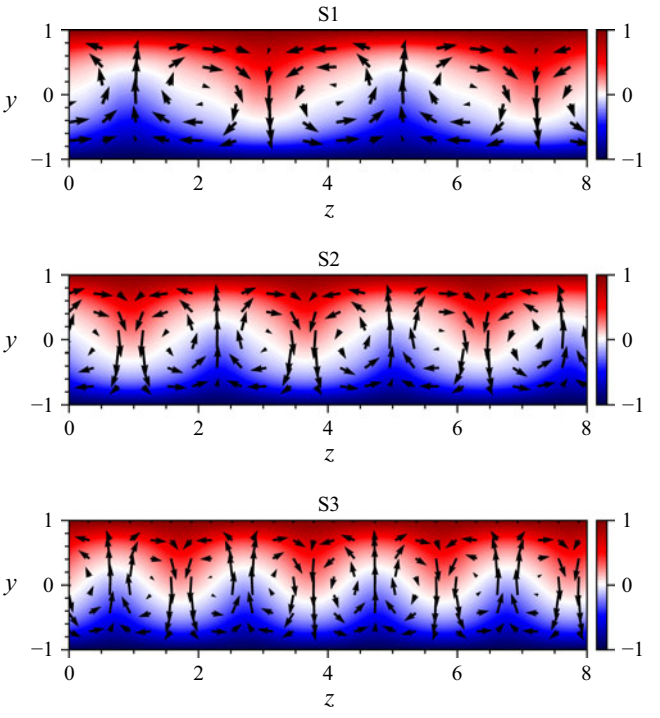


Figure 5. Equilibrium solutions labelled in figure 4 at $Re = 50$. Contours represent the streamwise velocity and vectors represent the wall-normal and spanwise velocities.

In addition to the main branches, figure 4 also shows intermittent deviations that appear only over very narrow ranges of Reynolds number. These most likely arise because each DNS run was initialised with a different random condition, leading to variations in the transient dynamics. The precise origin of these deviations has not been investigated here, but they could correspond either to additional stable branches with small basins of attraction, which are only rarely observed, or to trajectories becoming temporarily trapped in transient regions of state space before converging to the dominant attractor. A more systematic exploration of state space would be required to determine their exact nature.

5. Exact solutions found for RPCF

In this section the focus is a demonstration of the projected variational optimisation methodology applied to finding exact nonlinear solutions to the Navier–Stokes equations for RPCF.

5.1. *Implementation details*

This section discusses the numerics and programmatic strategies used to implement the procedure depicted in figure 3. All numerics are implemented in the Julia programming language, the majority of the code is available at The-ReSolver. A single spatially and temporally extended scalar field is represented as a 3D array, with each dimension representing the wall-normal, spanwise and time directions, respectively. The number of wall-normal points, spanwise points and temporal samples is denoted by N_y, N_z and N_t, respectively. The actual array being stored for most computations, $\boldsymbol{u_k}(y)$, has been transformed using a real fast Fourier transform (FFT) and, thus, has dimensions

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

of $N_y \times (\lfloor N_z/2 \rfloor + 1) \times N_t$ for each of the three vector components, where $\lfloor \cdot \rfloor$ denotes the integer floor operation. The spanwise and time directions are transformed using FFTW (fastest Fourier transform in the west, Frigo & Johnson 2005), and so the discretisations in these directions are necessarily uniform. In this work the wall-normal points are also uniformly distributed to allow for easy interoperability between the discretisations used by the DNS time-stepping code and the Julia optimiser code. The resolvent modes are represented by a four-dimensional array. At each frequency tuple for the streamwise independent flow, $\boldsymbol{k} = (k_z,\ k_t)$, there are a set of $M \in \mathbb{N}$ modes containing $3N_y$ elements that are required to be stored in memory. Here M is the finite number of resolvent modes used in the truncated version of the sums (2.11) and (2.16). This leads to an array with dimensions of $3N_y \times M \times (\lfloor N_z/2 \rfloor + 1) \times N_t$. The modal coefficients $a_{\boldsymbol{k}m}$ (and $s_{\boldsymbol{k}m}$) are represented very similarly to the spectral arrays, only with the array dimension corresponding to the wall-normal direction now representing the modal index m, with a length M for the number of modes used for the projection in (2.15) and (2.17). Thus, these arrays have a size of $M \times (\lfloor N_z/2 \rfloor + 1) \times N_t$. Since $M \sim N_y$ when the goal is to find exact solutions, the memory footprint of the coefficient arrays $a_{\boldsymbol{k}m}$ and $s_{\boldsymbol{k}m}$ is similar to the spectral arrays $\boldsymbol{u}_{\boldsymbol{k}}(y)$. The largest component of the memory requirements for finding exact solutions comes from the modes themselves, where all dimensions of the array grow with Reynolds number. This problem is highly parallelisable due to the global representation, which would be a necessary step to implementing these methods for fully 3D problems or for higher Reynolds numbers.

Derivatives in the spanwise and time directions can be computed using standard spectral methods. Derivatives in the wall-normal direction are computed using finite difference methods, generally of second-order accuracy to match as closely with the numerics of the DNS solver. The finite difference stencils are derived using the custom package FDGrids.jl. To compute the projection in (2.15) and (2.17), it is necessary to compute integrals over all the directions given the definition of the inner product in (2.3). In the spanwise and time directions this can be done by using Parseval's theorem, converting the integral over these directions to sums of the coefficients over the corresponding wavenumbers. The wall-normal integrals are computed using the method of undetermined coefficients for quadratures (Dahlquist & Björck 2008). The nonlinear terms in the computations are computed using a pseudo-spectral method, using a 3/2 padding rule in the spanwise and time directions to avoid aliasing errors.

When constructing the modes $\boldsymbol{\psi}_{\boldsymbol{k}m}$ using resolvent analysis, the base flow used is the laminar solution $\boldsymbol{u}_b(y) = y\boldsymbol{e}_x$. This choice is made for its computational simplicity and numerical robustness; the laminar base flow is well defined, analytically simple and guarantees a tractable, well-conditioned SVD for all wavenumbers $\boldsymbol{k} \in \mathbb{Z}^2$ (including $\boldsymbol{k} = 0$). Consequently, it provides a reliably calculable set of modes $\boldsymbol{\psi}_{\boldsymbol{k}m}$ without requiring *a priori* knowledge of the flow or introducing the numerical complications associated with a turbulent mean profile. We have observed that the difference in the modes generated using the laminar profile and turbulent mean do not have significant qualitative differences, thus, the optimisation is not greatly affected by this choice as long as a sufficient number are used. To compute the resolvent modes, the resolvent operator in (3.6) has to be discretised. That is, for each wavenumber tuple $\boldsymbol{k}$, the operator $\boldsymbol{R}_{\boldsymbol{k}}$ has to be evaluated for the given discrete representation of the wall-normal derivative $\partial/\partial y$. This process is done using primitive variables ($\boldsymbol{u}$ and p) as depicted in Symon *et al.* (2018). The resolvent takes the form as a matrix, augmented with the incompressibility constraint and the pressure gradient term, for each $\boldsymbol{k}$, which is then inverted. Before performing the SVD the resulting matrix has to be transformed such that it obeys the correct orthogonality condition defined in (2.14) (Sipp & Marquet 2013; Mckeon 2023). To accomplish this, the

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

quadrature weights used to evaluate the wall-normal integrals are Cholesky decomposed, which provides a weighting matrix that ensures the resulting SVD will be orthogonal with respect to the correct inner-product space. The divide-and-conquer eigenvalue algorithm is used to perform the SVD.

We are free to choose the optimisation algorithm used for solving the optimisation problem (2.6), since any gradient-based method will work. Unless otherwise stated, the L-BFGS algorithm is used. The L-BFGS algorithm is a quasi-Newton method, incorporating approximate Hessian information into each iteration, which significantly improves the convergence properties of the optimisation near minima compared with gradient descent (Liu & Nocedal 1989; Fletcher 2000). This is particularly important here because the optimisation problem is non-convex, with many possible solutions potentially very close together. The problem is also generally quite poorly conditioned, a topic that will be further expanded upon in § 6. The ability of the L-BFGS algorithm to approximate the local curvature of the solution reduces the effect of this poor conditioning. Further discussions on each algorithm can be found in Nocedal & Wright (2006). The optimisation is performed using the Optim.jl package (Mogensen & Riseth 2018). Any algorithm used here is also coupled with the Hager–Zhang line search algorithm as described in Hager & Zhang (2005). This further improves convergence rates, albeit at the possibility of reducing the robustness of the convergence. In practice, however, this is not observed to be an issue. The threshold for convergence is defined as $\mathcal{R}[\boldsymbol{u}] < 10^{-12}$.

5.2. *Equilibria*

The Reynolds number is set to $Re = 50$ and the rotation number to $Ro = 0.5$ for the equilibrium solutions. To begin, the optimiser is validated by observing its ability to reconstruct a known solution to the flow. This was obtained using the custom DNS solver, with a grid discretisation of $N_y = 64$ on a uniform grid, and $N_z = 33$ corresponding to 17 spanwise Fourier modes. A set of $M = 64$ resolvent modes $\boldsymbol{\psi}_{\boldsymbol{km}}$ are used for this optimisation. To obtain the initial guess for the start of the optimisation, S1 from DNS in figure 5 is projected onto the resolvent basis and then perturbed with random Gaussian white noise at each coefficient $a_{\boldsymbol{km}}$. This results in the highly disordered field in figure 6(*a*). This perturbed flow is then optimised to try to recover the original solution. The result of this optimisation is shown in figure 6(*b*). This can be compared with S1 in figure 5, where the equilibria from the DNS and optimisation are qualitatively indistinguishable. Figure 6(*c*) shows the difference between the solution obtained from optimisation and from the DNS. The magnitude of the difference is a couple of orders of magnitude smaller than that of the actual solution. The discrepancy likely arises from differences between the optimisation framework and the DNS formulation, rather than from the optimisation algorithm itself. Specifically, the optimisation is based on a discretisation of the continuous adjoint operator. It is known that the continuous adjoint of the linearised operator being studied, when discretised, does not exactly correspond to the adjoint of the discretised linearised operator, producing errors that only disappear as the resolution of the discretisation increases (Nadarajah & Jameson 2000; Hekmat & Mirzaei 2016; Gao, Wu & Xia 2017). A limited convergence study was performed in which the finite difference stencil widths, as well as the wall-normal and spanwise resolutions, were systematically varied. Although both the Julia optimiser and DNS solver exhibited internal convergence, the relative discrepancy between them did not decrease monotonically with refinement, indicating that it is not solely attributable to insufficient spatial resolution. This validates the fact that the optimiser possesses minima that correspond to solutions of (4.1).

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

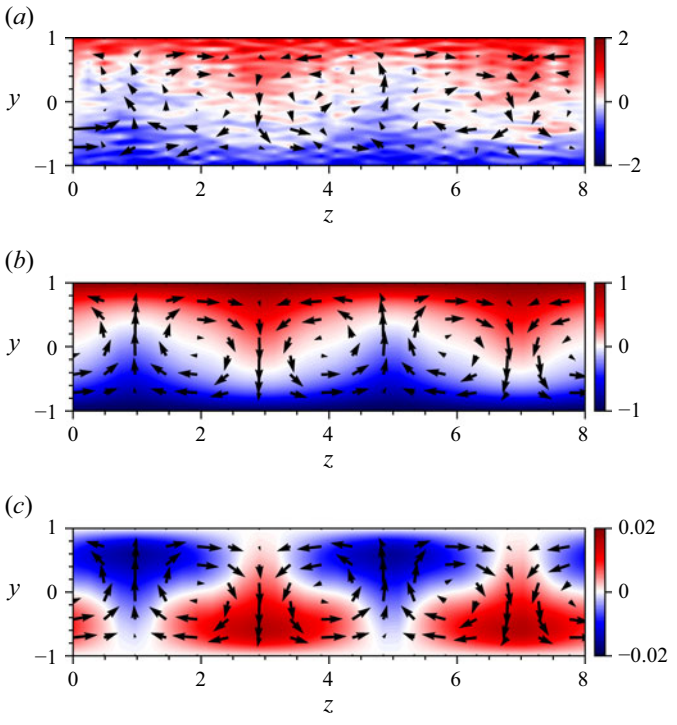


Figure 6. Snapshots of the flows before and after the optimisation at $Re = 50$ and $Ro = 0.5$, along with the solution obtained from DNS. Panel (*a*) shows the initial flow used for the optimisation, obtained by perturbing the stable solution obtained from DNS at the same Reynolds and rotation numbers. Panel (*b*) shows the result of the optimisation with a residual of $\mathcal{R} < 10^{-12}$ and panel (*c*) is the difference between the optimisation result and the DNS solution. Contours represent the streamwise velocity and vectors represent the wall-normal and spanwise velocities.

For comparison, the same initial condition was also optimised using the gradient descent and conjugate gradient algorithms (both still using Hager–Zhang line searches). The traces for each algorithm applied to the initial condition in figure 6(*a*) are shown in figure 7. Clearly L-BFGS performs the best, achieving the residual of $\mathcal{R} = 10^{-12}$ after a little over 4000 iterations. Gradient descent displays an initial large decrease in residual that quickly decreases as the convergence rate slows. This final and slow phase of the optimisation is a result of the optimiser approaching the minimum along the direction associated with the slowest growth direction in $\mathcal{R}$. The conjugate gradient also outperforms the gradient descent, achieving a smaller residual for the same number of iterations. However, after the initial large decrease in the residual at the beginning of the optimisation, the convergence rate slows, approaching a similar speed as the gradient descent. In addition, the iteration time of the conjugate gradient is observed to be larger than L-BFGS.

To investigate the robustness of the optimiser, it was initialised with initial conditions exciting certain spanwise wavenumbers with the intent of finding multiple new equilibrium solutions that are not observed in DNS at this Reynolds number. The same set of resolvent modes used for the previous optimisation are utilised here. The results of these optimisations are shown in figure 8. Figure 8(*a*) was initialised with the first third to fifth modal coefficients $a_{\boldsymbol{k}m}$ excited with random values at the spanwise wavenumber of $k_z = 4$, corresponding to streamwise rolls with half the wavelength of S1, along with the third mode at $k_z = 0$, with the rest of the coefficients left as zero. This synthetically constructed

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

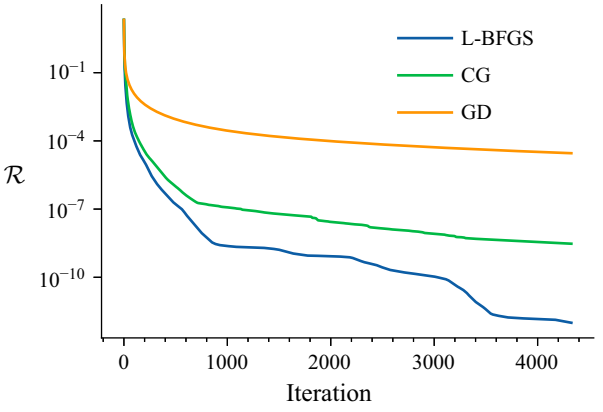


Figure 7. Residual trace for the optimisation of the initial flow given in figure 6(*a*), using gradient descent (GD), conjugate gradient (CG) and L-BFGS optimisation algorithms. All solutions converge towards the solution obtained in figure 6(*b*).

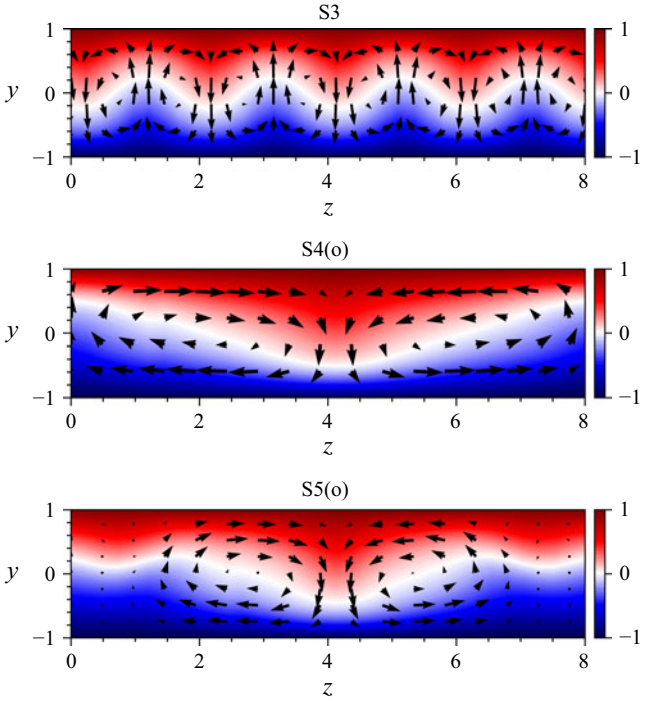


Figure 8. Final snapshots of the solutions obtained by optimising from various synthetic initial flow fields. Contours represent the streamwise velocity and vectors represent the wall-normal and spanwise velocities.

initial flow converges to S3, one of the stable equilibria shown in figure 5. Next, when the first five modal coefficients are excited randomly for the spanwise wavenumber $k_z = 1$, the solution in figure 8(*b*) is found, called S4(o). The addition of '(o)' in the name given to the equilibrium is in reference to the fact that it has been obtained from optimisation and is not a solution observed from the DNS, a result of them likely being linearly unstable. This solution again has the streamwise roll expected from the other equilibria solutions, but now

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

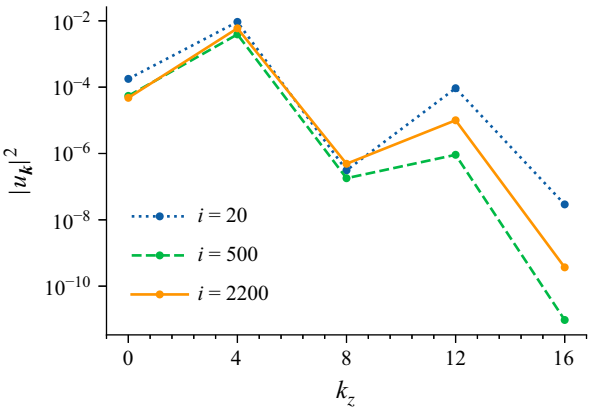


Figure 9. Spanwise power spectra of S3 of the solutions obtained at specific iterations of its optimisation, sampled at the channel mid-plane ($y = 0$).

with twice the spanwise wavelength. The highly symmetric structures of S3 and S4(o) closely mirror the solution of S1 and are ultimately expressing the same dynamics. To obtain a slightly more abnormal solution, the flow was initialised by exciting the first eight resolvent modes at the zeroth and first ($k_z = 0, 1$) spanwise wavenumbers with random values. The result of optimising such an initial condition is shown in figure 8(*c*), called S5(o). Here S5(o) shows a streamwise roll pattern with a width larger than in S1, roughly 3 spanwise units in length. This roll does not repeat over the remaining length of the domain since there is not sufficient room. Instead, the flow remains mostly stagnant in the remaining space, only expressing some very weak rolls that transport a small amount of the momentum from the top wall downwards.

The spanwise power spectra of the intermediate flow over the duration of the optimisation leading to S3 are shown in figure 9. That is, the squared magnitude of the Fourier coefficients of the solution S3 along the spanwise direction sampled at the mid-plane $y = 0$. For clarity, the wavenumbers that are orders of magnitude smaller than those shown are omitted. The omitted wavenumbers have negligible impact on the solution and are primarily a result of the spanwise domain size being large enough to contain multiple repetitions of the fundamental unit of the solution. The initial field used for the optimisation was excited at exactly one spanwise wavenumber, $k_z = 4$. The optimiser initially spreads out the spectral content from the initial condition to the other wavenumbers as can be seen at iteration $i = 20$. The large wavenumbers are then damped by viscous effects and the energy containing low wavenumbers are tuned towards the final solution by $i = 500$. The remaining iterations are focused on finely tuning the higher wavenumbers while gradually decreasing the residual, until a balanced flow is achieved at $i = 2\,200$. The effects of the individual terms in the gradient and how they affect the optimisation are further discussed in § 6

A preliminary natural continuation of equilibrium S1 implies that these equilibrium solutions do not exist at the larger-Reynolds-number regimes shown in § 4. It is not known, however, whether these equilibrium solutions exist in the higher-Reynolds-number regimes. The only definitive way to answer this question is to perform a rigorous continuation of the solutions to higher Reynolds numbers. With this comes the expected additional computational costs associated with the finer resolution required in more energetic regimes.

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

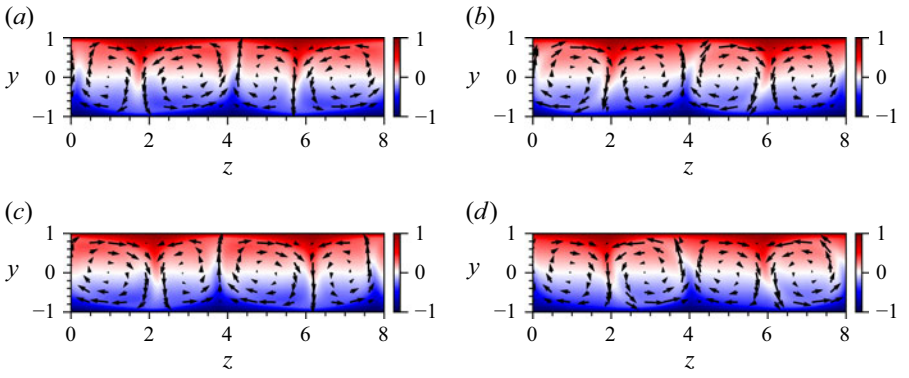


Figure 10. Snapshots of the periodic solution ($\mathcal{R} \approx 5 \times 10^{-11}$) obtained for $Re = 400$ with a period of $T \approx 25.05$. Results are shown for (*a*) $t = 0$, (*b*) $t = T/4$, (*c*) $t = T/2$, and (*d*) $t = 3T/4$. Contours represent the streamwise velocity and vectors represent the wall-normal and spanwise velocities.

5.3. *Periodic solutions*

The optimiser is now tasked with constructing a periodic solution at $Re = 400$, a regime where RPCF displays a stable periodic motion as seen in figure 4. The flow is discretised with $N_y = 128$, $N_z = 65$, corresponding to 33 spanwise modes, and $N_t = 35$ temporal modes. A set of $M = 64$ resolvent modes for each frequency are used for the optimisation. The smaller number of resolvent modes is used primarily to reduce computational effort. As will be seen in this section, the reduced number of modes compared with the degrees of freedom in the wall-normal direction has little effect on the accuracy of the final result. The initial guess for the optimiser is initialised in a similar way as in the first validation case in § 5.2. The L-BFGS algorithm is still used for the optimisation along with a Hager–Zhang line search. The periodic DNS data at $Re = 400$ is projected onto the resolvent modes. The resulting coefficients $a_{\boldsymbol{km}}$ are then perturbed with random Gaussian white noise and the resulting noisy flow is optimised.

A movie comparing the evolution of the solution obtained via optimisation and that from DNS can be found in the supplementary materials are available at https://doi.org/10.1017/jfm.2026.11722. Figure 10 shows a set of snapshots for the periodic solution obtained from optimisation at $Re = 400$, sampled at points along its trajectory. The primary streamwise rolls are clearly present, and evolve in a wavy motion as consecutive vortices contract and expand. Figure 11 shows the spectrum of the periodic solution sampled at the two wall-normal positions $y \approx -0.86$ and $y \approx 0$. Each spectrum is a slice of the total spectrum such that only the positive temporal frequencies are shown, since the negative temporal frequencies are a reflection. An imbalance in the positive and negative temporal frequencies would manifest as travelling waves moving in the spanwise direction, which is not a feature of this solution. The general lack of spanwise motions in solutions to channel flows is a noted feature in Cvitanović & Gibson (2010). A chequerboard pattern is observed in the spectrum near the centreline of the flow. In the temporal spectrum, this is a result of the periodic solution being made up of a pre-periodic solution repeating itself under a symmetry transformation. Pre-periodic solutions are able to exist in this flow due to the presence of certain discrete and continuous symmetries (Gibson *et al.* 2008); however, such symmetric solutions are not treated in this work. The pattern in the spanwise direction, present only near the centre plane is simply a result of the strong streamwise rolls that dominate the flow at this Reynolds number. The most energetic mode is located at $(k_z,\ k_t) = (1,\ 0)$, which is the mode that best fits the streamwise rolls

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

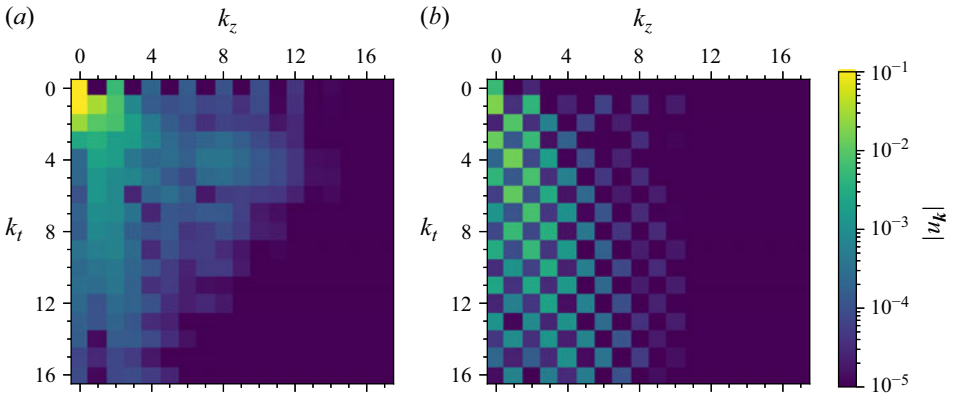


Figure 11. Spanwise and (positive) temporal power spectrum of the periodic solution in figure 10 at $y \approx -0.86$ in panel (*a*) and $y \approx 0$ in panel (*b*).

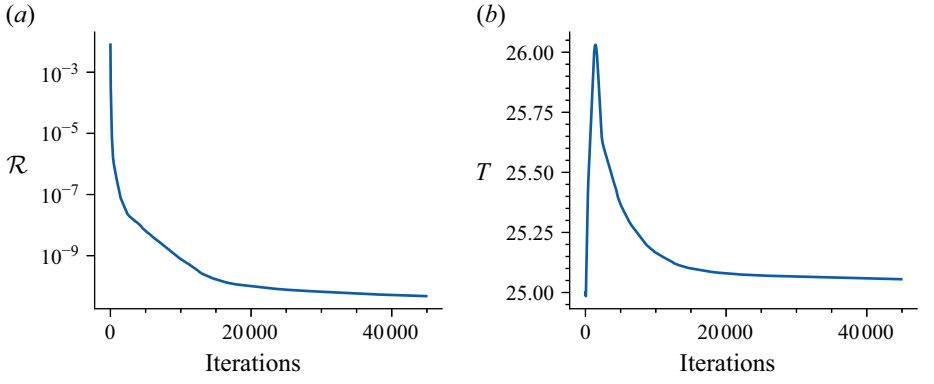


Figure 12. Traces of the global residual (*a*) and the solution period (*b*) of the periodic solution in figure 10 over the duration of the optimisation.

that do not vary significantly with time. The spectrum at the location nearer the wall has a more continuous decay indicative of its less obvious spatial and temporal structure compared with the centre of the channel.

Figure 12 shows the trace of the global residual and period in panels (*a*) and (*b*), respectively. The final residual achieved is roughly $\mathcal{R} \approx 5 \times 10^{-11}$. Compared with figure 7, the convergence rate of the problem is considerably slower after the initial rapid drop in residual. This slower convergence compared with that observed in figure 7 can be explained by the increased degrees of freedom of the problem, both in terms of the inclusion of an extra temporal dimension and in terms of the added resolution required at higher Reynolds numbers. It is also likely that this particular solution at the Reynolds number of $Re = 400$ is near marginal stability; however, analysis of the stability of these solutions has not been undertaken in this work. The impact of marginal stability is discussed further in § 6. As a final point, the maximal and minimal kinetic energies of the optimised periodic solution are plotted on figure 4 (denoted with red triangles) with the corresponding kinetic energy extrema from DNS, which shows good agreement between the solutions from the two different numerical sources.

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

5.4. *Performance of variational optimisation*

The performance of the variational optimisation is best understood in comparison with the leading methods used to compute time-periodic solutions of nonlinear dynamical systems, namely the variational Newton flow method of Lan & Cvitanović (2004) and the Newton–GMRES–hookstep shooting method (Viswanath 2007, 2009). Both are Newton-based approaches and, therefore, require the solution of a linear system at each iteration. For the variational Newton flow method, the associated linear operator has dimension $N \times N$, where N is the total number of degrees of freedom. For realistic fluid flow problems, this renders the method computationally intractable without any modification.

The Newton–GMRES–hookstep shooting method circumvents the explicit formation of this operator by modifying the Newton iteration with a Krylov subspace method, yielding a matrix-free implementation suitable for high-dimensional turbulent flows. Nevertheless, each Newton step typically requires tens to hundreds of matrix–vector products at least in order to solve the linear system to sufficient accuracy. In contrast, the variational optimisation method evaluates the gradient, and hence, the update direction, using a single matrix–vector product per iteration. It is therefore significantly cheaper per iteration. Additionally, as demonstrated here and in Farazmand (2016), Azimi *et al.* (2022), Ashtari & Schneider (2023), the variational optimisation is also substantially more robust to the choice of initial guess. More formally, convergence to a minimum is guaranteed, whereas Newton-based methods may fail to converge unless the initial guess lies sufficiently close to a solution.

These advantages, however, come at the expense of linear convergence near a minimum, in contrast to the quadratic convergence attainable with Newton methods. Consequently, once the residual has been reduced below a moderate threshold, it becomes more efficient to switch to Newton–GMRES–hookstep in order to drive the residual to machine precision. This observation motivates the hybrid strategy proposed by Farazmand (2016). In the present computations, the first 40 000 iterations shown in figure 12 required approximately three days of wall time, with each iteration taking $\sim$7 s. However, the first 2000 iterations, sufficient to reduce the residual below $\mathcal{R} < 10^{-7}$, required only around four hours.

In the specific formulation of the problem here, the memory cost of the algorithm is dominated by the storage of the modes $\boldsymbol{\psi}_{\boldsymbol{km}}$. The memory consumed by these modes scales with $\mathcal{O}(MN)$, where M is the total number of resolvent modes used and N here represents the combination of the spatial and temporal degrees of the freedom of the problem. For 2D3C RPCF, this is given as $N = 3N_y(\lfloor N_z/2 \rfloor + 1)N_t$. For reference, the modes used for the optimisation of the periodic solution shown in figure 10 consumed $\sim$430 MB. Using the Kolmogorov hypothesis it can be estimated that the number of grid points required in each direction scales as $Re^{3/4}$ (Pope 2001). Using this scaling, and assuming that the number of modes M scales in a similar way to the grid points, it can be estimated that the required memory to store the modes at a fully turbulent Reynolds number of $Re = 4000$ (up from $Re = 400$) is $\sim$420 GB. Clearly this becomes a prohibitively expensive endeavour as the Reynolds number grows. The most practical solution is to distribute the data over multiple processes to reduce the local memory requirements. Further discussions on the performance of the variational optimiser in general can be found in Azimi *et al.* (2022). The time complexity of the algorithm is dominated by the Fourier transforms performed to compute the nonlinear terms in (2.18) and (2.19), scaling as $\mathcal{O}(N_y N_{hom} \log(N_{hom}))$, where $N_{hom} = (\lfloor N_z/2 + 1 \rfloor)N_t$ are the degrees of freedom of the homogenous spatial directions and time. This cannot be made any more efficient in time than its current form.

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

6. Conditioning of the optimisation

In the literature the convergence rate of variational optimisation algorithms slows considerably near a minimum, a problem governed by the local curvature of the residual function (Nocedal & Wright 2006). This curvature is characterised by the Hessian operator, which contains all second-derivative information at a point in optimisation space. The condition number of the Hessian measures the scale separation between the slowest and most rapid growth rates experienced by the objective when moving away from the minimum. A lower condition number implies improved convergence rates for gradient-based optimisation algorithms. A condition number of one signifies a perfectly quadratic local neighbourhood, while larger values indicate a long, narrow 'valley', which impedes convergence. The goal of this section is to link this conditioning directly to the underlying flow dynamics and to demonstrate that the resolvent modes provide a basis optimally constructed to improve the conditioning of the optimisation problem.

The first step in the analysis is to derive a closed-form expression for the Hessian operator of the residual as defined in (2.6). To do this, first consider a minimum $\mathcal{R}[\boldsymbol{u}^*] = 0$, with $\boldsymbol{u}^*$ denoting the flow field that solves the Navier–Stokes equations. Since $\boldsymbol{u}^*$ is a minimum of (2.6), then we also know that $\delta\mathcal{R}/\delta\boldsymbol{u}^* = 0$. Assume that $\boldsymbol{u}^*$ is an equilibrium. This does not change the content of the discussions that follow, only serving to simplify the mathematics by removing the need to use Floquet analysis. Adding an infinitesimal (potentially unsteady) perturbation to $\boldsymbol{u}^*$ in the direction $\boldsymbol{v}$, where $\boldsymbol{v}$ is incompressible and obeys a set of homogeneous BCs similar to the local residual in (2.9), and expanding in terms of a Taylor series, only retaining the second-order term, leads to the relation

$$\langle \boldsymbol{v},\ \boldsymbol{H}\boldsymbol{v}\rangle_{\Omega_t} = \left\| \frac{\partial \boldsymbol{v}}{\partial t} - \mathbb{P}\mathcal{L}_{\boldsymbol{u}^*}\boldsymbol{v} \right\|^2_{\Omega_t}, \tag{6.1}$$

where $\boldsymbol{H}$ denotes the Hessian operator, $\mathcal{L}_{\boldsymbol{u}^*}$ denotes the linearised Navier–Stokes operator evaluated at the minimum $\boldsymbol{u}^*$ and $\mathbb{P}$ is the Leray projector from § 3. The inclusion of the Leray projector is as a notational shorthand to avoid the added difficulty in having to explicitly account for the pressure gradient term in the Navier–Stokes equations. The detailed derivation of this equation is given in Appendix E. Let $\boldsymbol{v}$ be a unimodal perturbation to that minimum, i.e. let $\boldsymbol{v} = \boldsymbol{v}_0 e^{i\omega t} + \text{c.c.}$ for some arbitrary frequency where $\omega \in \mathbb{R}$ and some steady flow field $\boldsymbol{v}_0$, where c.c. denotes the complex conjugate. This perturbation can be substituted into (6.1), effectively performing a Fourier transform, which gives

$$\langle \boldsymbol{v}_0,\ \boldsymbol{H}_\omega \boldsymbol{v}_0\rangle_{\Omega} = \| \left(i\omega \boldsymbol{I} - \mathbb{P}\mathcal{L}_{\boldsymbol{u}^*}\right) \boldsymbol{v}_0 \|^2_{\Omega}. \tag{6.2}$$

The Hessian operator is self-adjoint, as is apparent from (6.1), which means it is also normal. As such, the eigenvalues obey a strict ordering, i.e.

$$0 \leqslant \mu_{\omega 1} \leqslant \mu_{\omega 2} \leqslant \cdots \leqslant \mu_{\omega n} \leqslant \cdots . \tag{6.3}$$

This sequence of eigenvalues is unbounded since the Hessian itself is an unbounded operator. This is equivalent to saying that there exists no positive number, L, such that $\mu_{\omega n} \leqslant L$ for all $n \in \mathbb{N}$. In practical terms, when the governing equations are discretised in space this leads to a new discretised Hessian operator, denoted as $\hat{\boldsymbol{H}}_\omega$, which does have a bounded spectrum. The relationship between the spectra of $\boldsymbol{H}_\omega$ and its discretised counterpart $\hat{\boldsymbol{H}}_\omega$ depends on the particular discretisation used and is not the focus of the present analysis. Instead, the focus is that of the largest eigenvalue of the Hessian that can be resolved from $\hat{\boldsymbol{H}}_\omega$. Defining this most extreme element of the spectrum as $\mu_{\omega N}$ it is possible to then define a condition number of the (discrete) Hessian: $\kappa(\hat{\boldsymbol{H}}_\omega) = \mu_{\omega N}/\mu_{\omega 1}$.

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

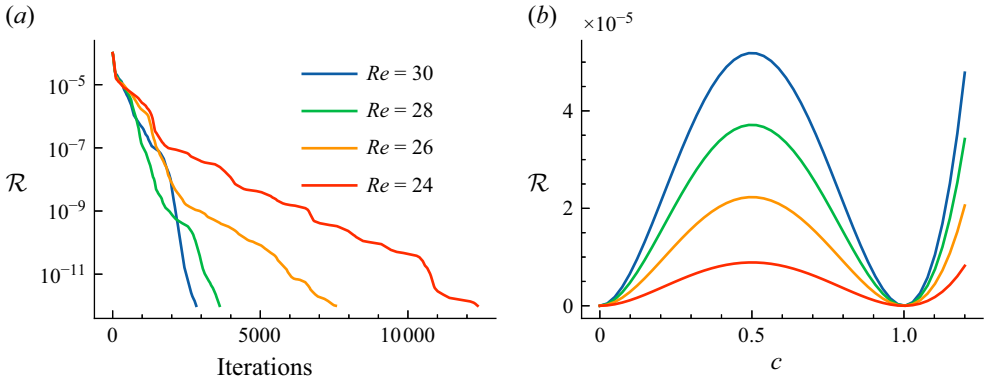


Figure 13. Residuals from optimisations of at Reynolds numbers of $Re = 30,\ 28,\ 26,\ 24$. Panel (a) shows the residuals achieved during an optimisation of the equilibrium at each Reynolds number, and panel (b) shows the residual values of the linearly interpolated velocity field between the final equilibrium solution and the laminar solution, $\boldsymbol{u}(c) = \boldsymbol{u}_{lam} + c(\boldsymbol{u}_{EQ} - \boldsymbol{u}_{lam})$.

It has been documented that the convergence rate of variational optimisers degrades near marginally stable solutions or bifurcation points (Lakoba & Yang 2007; Farazmand 2016). This phenomenon can be understood by considering the linear stability of the equilibrium solution $\boldsymbol{u}^*$, governed by the spectrum of $\mathcal{L}_{\boldsymbol{u}^*}$. Inspecting the derived expression for the Hessian in (6.2), it is clear that in the limit of any eigenvalue of $\mathcal{L}_{\boldsymbol{u}^*}$ approaching the imaginary axis (neutral stability), the action of the operator $i\omega\boldsymbol{I} - \mathcal{L}_{\boldsymbol{u}^*}$ becomes arbitrarily small near the frequency ω where that eigenvalue crosses. This implies that any perturbation $\boldsymbol{v}_0$ aligned with this marginally stable direction has a vanishingly small action from the Hessian, meaning the residual grows very slowly in this direction. In relation to the spectrum of $\boldsymbol{H}_\omega$, this is equivalent to the smallest Hessian eigenvalue $\mu_{\omega 1}$ approaching zero, which forces the condition number $\kappa(\hat{\boldsymbol{H}}_\omega) \to \infty$. This mathematical insight is the basis for acceleration methods that assume the optimiser is traversing along the most marginally stable mode (Azimi *et al.* 2022; Ashtari & Schneider 2023), and several techniques have been devised to try and remove these problematic directions (Yang & Lakoba 2007), although they have not been applied to fluid dynamics problems.

This behaviour is demonstrated in figure 13(a), which shows the global residual trace obtained from a series of optimisations performed at difference Reynolds numbers. As mentioned in § 4 at $Re_{crit} \approx 20.7$, a bifurcation occurs where a new stable equilibrium solution is born out of the laminar solution. Thus, as the Reynolds number gets closer to this bifurcation, the stable equilibrium gradually becomes more neutrally stable, leading to degrading convergence rates. Figure 13(b) plots the global residual of the linearly interpolated flow field between the laminar solution and the equilibrium solution obtained from the optimisation at the corresponding Reynolds number, given by the expression $\boldsymbol{u}(c) = \boldsymbol{u}_{lam} + c(\boldsymbol{u}_{EQ} - \boldsymbol{u}_{lam})$, where $c \in [0,\ 1]$. The variation of global residual between the laminar solution and the equilibrium solution decreases as the Reynolds number approaches Re_{crit}. This demonstrates the emergence of a near neutral direction for the optimisation. The rate at which the convergence degrades as the Reynolds number approaches neutrality depends on the type of bifurcation that is occurring and how fast the particular eigenvalue that is switching stability approaches the imaginary axis.

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

The relationship in (6.2) can be rearranged to provide a closed-form expression for the continuous Hessian operator, i.e.

$$\boldsymbol{H}_{\omega}\boldsymbol{v}_0 = (i\omega\boldsymbol{I} - \mathbb{P}\mathcal{L}_{\boldsymbol{u}^*})^{+}\,(i\omega\boldsymbol{I} - \mathbb{P}\mathcal{L}_{\boldsymbol{u}^*})\,\boldsymbol{v}_0, \tag{6.4}$$

where $(\cdot)^{+}$ denotes the adjoint of an operator. Taking the definition of the resolvent operator $\boldsymbol{R}_{\omega} = (i\omega\boldsymbol{I} - \mathbb{P}\mathcal{L}_{\boldsymbol{u}^*})^{-1}\mathbb{P}$ from (3.6) and applying the SVD defined in (3.7), the non-normal growth mechanisms of $\mathcal{L}_{\boldsymbol{u}^*}$ can be related to the spectrum of $\boldsymbol{H}_{\omega}$:

$$\boldsymbol{H}_{\omega}\boldsymbol{v}_0 = \sum_{m=1}^{\infty} \sigma_{\omega m}^{-2}\boldsymbol{\psi}_{\omega m}\left\langle \boldsymbol{\psi}_{\omega m},\ \boldsymbol{v}_0\right\rangle_{\Omega}. \tag{6.5}$$

Equation (6.5) represents an eigendecomposition of the Hessian, directly linking its spectrum to the resolvent of $\mathcal{L}_{\boldsymbol{u}^*}$. The eigenvalues of $\boldsymbol{H}_{\omega}$ are the squared inverse of the resolvent's singular values ($\mu_{\omega m} = \sigma_{\omega m}^{-2}$), and their corresponding eigenmodes are identical to the resolvent response modes. Therefore, the condition number of the discretised operator $\kappa(\hat{\boldsymbol{H}}_{\omega})$ is primarily governed by the smallest retained singular value when the expansion is discretised. This relationship provides an optimal strategy for preconditioning since by truncating the resolvent expansion to remove the modes with the smallest singular values, one selectively eliminates the largest, most problematic eigenvalues from the Hessian. This optimally reduces the condition number, as it precisely targets and removes the directions associated with the strongest growth in the residual. In Mons & Marquet (2021) the authors attempted to derive a methodology to determine optimal sensor placement for data assimilation. Their formulation lead to a similar relationship to (6.4), relating the Hessian operator of their problem to the resolvent of their linearised dynamics. The fact that this relationship appears in both this work and Mons & Marquet (2021), despite the different objectives, is indicative of the role the resolvent has in modelling the higher-order sensitivities in these types of optimisation problems.

A key limitation is that these optimal resolvent modes depend on the linearised operator $\mathcal{L}_{\boldsymbol{u}^*}$ evaluated at an *a priori* unknown minimum. Thus, to be able to derive the modes used in (6.5), the solution has to be known before the optimisation. Therefore, it appears that this does not provide a useful method for improving convergence rates. We posit, however, that the higher-order modes exhibit a universal structure, independent of the base flow, due to the dominance of viscous dissipation at small scales. This is favourable, as (6.5) indicates that rejecting these higher-order modes is the principal method for improving the Hessian's condition number. It should be noted that the sensitivity of resolvent modes to the base flow used for the linearisation remains poorly characterised, beyond the established dependence of the singular values in Brandt *et al.* (2011).

To demonstrate that rejecting higher-order resolvent modes improves the convergence rate of the optimisation, we perform a series of optimisations starting from a perturbed state close to S3 from §§ 4 and 5.2 at $Re = 50$. The optimisation is performed using progressively fewer resolvent modes to generate the low-order model. The modes used are the same as those used in § 5.2, derived by linearising about the laminar base flow. The results of these optimisations are shown in figure 14.

Figure 14(*a*) shows the global residual traces for the perturbed field using 8, 16, 32 and 64 resolvent modes. All the residuals reduce at a high rate initially, indicating the power of the variational optimiser to quickly seek out the solution primarily by modifying the large-scale structures. The initially fast decrease, however, eventually gives way to a slower convergence rate governed by a power law, as is typically observed in the literature (Azimi *et al.* 2022; Ashtari & Schneider 2023). The primary point to note from figure 14(*a*) is that both the degree of the eventual slowdown in convergence rate and the iteration count at

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

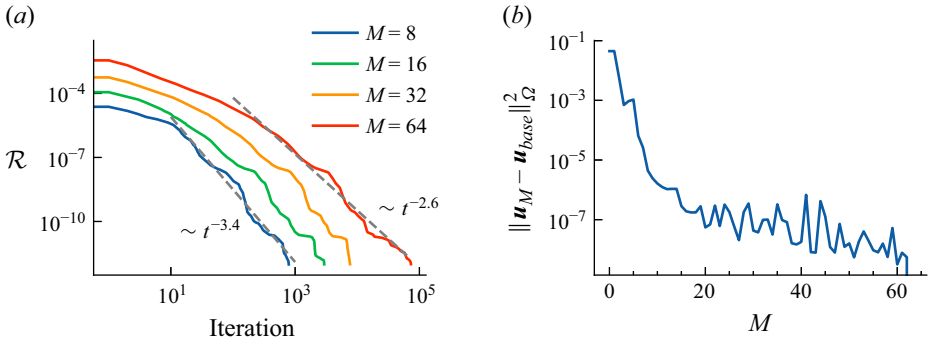


Figure 14. (*a*) Global residual traces for the optimisation of a perturbed S3 solution, performed using $M = 8$, 16, 32 and 64 resolvent modes, all with the same set of initial coefficients, using the L-BFGS algorithm. (*b*) Accuracy of the resulting solutions found by the optimiser relative to the 'base' case obtained for 64 resolvent modes, plotted against the number of modes used for the projection, each solution being converged such that $\mathcal{R} < 10^{-12}$.

which it occurs are linked to the degree of truncation of the sums in (2.11) and (2.16), i.e. how many dimensions the resulting optimiser has to navigate during its run. This is shown in figure 14(*a*) with two straight line estimating the rate of decrease of the residual for $M = 64$ and $M = 8$. The case with more modes has a power-law scaling equal to roughly -2.6 and the case with fewer modes has a power law of roughly -3.4. The larger value for $M = 8$ indicates a faster rate of decrease of the residual. The case of $M = 64$ resolvent modes achieved its final residual after roughly 10^5 iterations, whereas the $M = 8$ case shows the same residual after two orders of magnitude fewer iterations. The figure clearly demonstrates that the truncation of the resolvent modes, derived about the laminar base flow, associated with the smallest singular values improves convergence. This provides confirmation that the condition number of the Hessian at the minimum is improved, even though the resolvent modes are derived relative to a different base flow. A final note is that an immediate result of truncating the resolvent modes is to reduce the global residual of the initial field. This is a result of the noise introduced in the perturbation of S3 being truncated in the projection, reducing the smaller-scale noise present in the starting flow. The initially smaller residual should not be the reason attributed to why the cases with fewer modes converge earlier, as the faster convergence is primarily associated with the faster convergence rate as discussed earlier.

In addition to the convergence rates, figure 14(*b*) shows the accuracy of the resulting solutions obtained from the projected optimisations plotted against the number of resolvent modes used for the expansions in (2.11) and (2.16). The accuracy was computed as the norm of the difference between the field obtained from the projected optimisation, denoted with $\boldsymbol{u}_M$, and a base solution, denoted with $\boldsymbol{u}_{base}$, which was taken as the solution obtained using the full set of resolvent modes, i.e. $M = 64$. As the number of modes used for the projection is increased, the error between the obtained solution and the base decreases, initially quite quickly when the number of modes used is small, but then reducing slower and saturating at around $\|\boldsymbol{u}_M - \boldsymbol{u}_{base}\|^2_{\Omega_t} \approx 10^{-7}$ at roughly $M = 20$. The initial decrease in the error is a result of the extra modes being added having a relatively large contribution to the solution S3, with the observed exponential decrease in error typical of spectral methods. The saturation of the solution error as more modes are added is due to tolerances associated with the small changes in the minimum position and initial starting point for the optimisation for each value M.

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

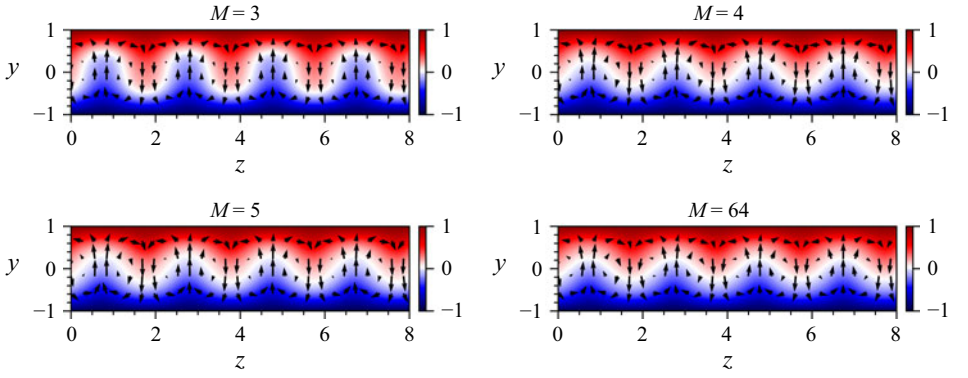


Figure 15. Snapshots of the solutions obtained from the projected optimisation of the perturbed S3 solution. Contours represent the streamwise velocity and vectors represent the wall-normal and spanwise velocities.

Figure 15 demonstrates that even a small number of resolvent modes can faithfully reconstruct the large-scale structure of a solution, despite potentially poor quantitative accuracy. It shows snapshots of the solutions obtained using $M = 3$, $M = 4$, $M = 5$ and the base case $M = 64$. Each solution displays the desired streamwise rolls with the correct spanwise wavelength. The solution for $M = 3$ has noticeable qualitative differences from the base solution, consistent with its $\mathcal{O}(1)$ error in figure 14(*b*). However, the solutions for $M = 4$ and $M = 5$ are far more similar to the base solution visually, even though their quantitative error remains significant.

Taking these results together, we conclude that retaining a sufficiently small number of modes allows the projected optimisation to reconstruct the dominant structures of the desired solution, albeit at the cost of final accuracy, but with a significantly improved convergence rate. This presents a practical trade-off. If only the large-scale structures are sought, aggressively truncating the number of modes yields a solution far more rapidly. This could be used to initialise searches for ECS, using a truncated set of modes to quickly reduce the residual and capture the large-scale flow. The number of modes could then be increased to resolve smaller scales, or the output could be handed off to a Newton-GMRES-hookstep method for final convergence, akin to the approach in Farazmand (2016).

7. Conclusions

This work makes two primary contributions to the variational optimisation of ECS. First, we introduce a methodology for treating wall-bounded time periodic flows by projecting the optimisation onto a basis of divergence-free resolvent modes that inherently enforce no-slip BCs and incompressibility. This approach, which extends the methods of Farazmand (2016) and Azimi *et al.* (2022), directly embeds the required constraints, alleviating a key theoretical difficulty. This is in contrast to the approach taken in Ashtari & Schneider (2023) using the IM method. In doing so, it unifies the optimisation framework of these earlier works for finding ECS with the resolvent-based modelling framework of McKeon & Sharma (2010), providing a direct method to 'close the loop' in resolvent analysis (Barthel *et al.* 2021) by solving for the self-consistent, finite-amplitude velocity field that sustains the chosen forcing modes. Second, we establish a formal link between the conditioning of the optimisation problem and resolvent analysis via the Hessian operator. We demonstrate that truncating the resolvent basis not only acts as an effective preconditioner for the optimiser, but also benefits from the resulting reduced-order model

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

reducing the dimensionality of the problem while still leading to physically meaningful results.

The projected optimisation methodology is applied to 2D3C RPCF, with all analysis performed with a rotation number of $Ro = 0.5$. Both equilibria and periodic solutions are sought in this work. To ensure that only minimal *a priori* knowledge about the flow is required, the laminar profile is used to generate the required resolvent modes instead of the turbulent mean as is more standard for typical resolvent analysis applications. The implications of this choice and the effect it has on the modes has not been investigated here and warrants further study. The optimiser successfully identified equilibria that are not observed in DNS at $Re = 50$ from various initial conditions. At $Re = 400$, a stable periodic solution was obtained by initialising the optimiser with a perturbed DNS solution. In this case the residual decreased from $\mathcal{R} > 10^{-3}$ to $\mathcal{R} < 10^{-12}$ in roughly 40 000 iterations, over half of which occurred in the first 5000. The slower convergence for the optimisation of the periodic solution is argued to be primarily a result of the increase in the degrees of freedom of the system and the possible weak linear stability of the solution. A notable omission is that of any time-periodic solutions obtained from optimisation that were not observed in DNS. Such results are omitted due to the very large cost of obtaining such solutions as the Reynolds number and computational domain size increase, specifically in terms of memory usage. To make these larger problems more feasible, a distributed implementation is required, which is discussed further below.

We also investigated the factors that affect the asymptotic convergence rate of the optimisation via the conditioning of the discretised Hessian operator at a global minimum. It is shown that the closer a solution is to being marginally stable the larger the condition number and, thus, the slower the convergence, a behaviour which has been discussed in Farazmand (2016). In addition, a direct link is established between the condition number of the Hessian and the resolvent expansion at a global minimum. It is shown that truncating the resolvent expansion by excluding the smallest singular values is equivalent to removing the fastest growing directions in the residual, which optimally reduces the condition number of the Hessian operator. The improved convergence is demonstrated by optimising to an equilibrium solution using a reduced number of resolvent modes. Physically, this equivalence arises because the Laplace operator dominates the higher-order behaviour of the Hessian eigenmodes and the resolvent modes regardless of the base flow around which the linearisations are performed. This operator is primarily responsible for the degradation in convergence rates, owing to its rapid dissipation of high-frequency components, particularly at higher Reynolds numbers where small-scale motions must be resolved. Removing the highest-order resolvent modes is equivalent to removing these modes of growth from the optimisation. The resulting reduced-order solutions obtained from the optimisation on a truncated set of modes were shown to closely approximate the original solution derived from the full basis set. Specifically, only a handful of modes were shown to be very close to the original solution. This means that even though the solution obtained using a very reduced set of modes does not accurately reconstruct the true solution in state space, it could be used to initialise a more accurate search with superior convergence properties such as the Newton-GMRES-hookstep method.

Ultimately, the Galerkin projection not only enables the construction of solutions for general wall-bounded flows but also enhances the capability of variational optimisation. Coupled with its robustness to initial conditions, the improved convergence makes the method more practical both as a stand-alone solver and as a preconditioning step for root-finding approaches. The ability to truncate the modal basis further extends applicability to larger problems than are typically tractable with global solvers.

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

The projected variational optimiser presented in this work has two primary limitations. The most immediate is its high memory requirement for storing the high-dimensional state vector, which can become prohibitive for extending this method to fully 3D turbulent flows. A promising path to mitigate this is leveraging parallel computing architectures, for which this framework is exceptionally well suited since, unlike traditional DNS of turbulence, the optimisation can be parallelised in the temporal dimension in addition to the spatial ones. A highly parallel implementation of the methodology would also permit an investigation into the so-called 'quasi-trajectories' from Burton *et al.* (2025) for turbulent flows. These quasi-trajectories would constitute very low-order representations of the flow, similar in spirit to those explored by Li & Lasagna (2026). However, whereas Li & Lasagna (2026) considered cases where the final objective was driven to very small values, sometimes near machine precision, the quasi-trajectories discussed by Burton *et al.* (2025) would not aim to satisfy the low-order projection of the governing equations to such an extent. The second, more fundamental limitation is the efficacy of the resolvent basis being intrinsically linked to the chosen base flow. Determining a general principle for selecting the most effective base flow to yield useful, dynamically relevant modes requires extra investigation and represents an important direction for future work.

Supplementary movie. Supplementary movie is available at https://doi.org/10.1017/jfm.2026.11722.

Funding. This work was supported using the Antony Wright Scholarship provided by the Department of Aeronautics and Astronautics at the University of Southampton.

Competing interests. The authors declare none.

Appendix A. Residual gradient derivation

A.1. *Velocity field derivative*

Before beginning the derivation, we define a modified optimisation problem that enforces the incompressibility constraint through a Lagrange multiplier,

$$\min_{\boldsymbol{u}\in\mathcal{P}_{T,BC},\,T} \quad \mathcal{R}\left[\boldsymbol{u},\ q\right] = \frac{1}{2}\|\boldsymbol{r}\|^2_{\Omega_t} + \langle q,\ \boldsymbol{\nabla}\cdot\boldsymbol{u}\rangle_{\Omega_t}\,, \tag{A1}$$

where

$$\mathcal{P}_{T,BC} = \left\{\boldsymbol{u}\in\chi \;\middle|\; \boldsymbol{u}|_{t=0} = \boldsymbol{u}|_{t=T}\,,\ \boldsymbol{u}|_{y=\pm1} = \pm\boldsymbol{e}_x,\ \boldsymbol{u} \text{ obeys periodic BCs}\right\} \tag{A2}$$

is the space of velocity fields that obey the BCs, but are not necessarily incompressible. Also define the space of fields that obey the homogeneous BCs

$$\mathcal{P}_{T,HBC} = \left\{\boldsymbol{u}\in\chi \;\middle|\; \boldsymbol{u}|_{t=0} = \boldsymbol{u}|_{t=T}\,,\ \boldsymbol{u}|_{y=\pm1} = 0,\ \boldsymbol{u} \text{ obeys periodic BCs}\right\}, \tag{A3}$$

which will be useful for the later parts of this derivation. To derive the gradient, first define a perturbation to the velocity field $\boldsymbol{u}\to\boldsymbol{u}+\epsilon\boldsymbol{v}$ where $0<\epsilon\in\mathbb{R}$ is small enough such that $\epsilon^2\approx 0$. The perturbation is incompressible and obeys the homogeneous BCs, i.e. $\boldsymbol{v}\in\mathcal{P}_{T,HBC}$ and $\boldsymbol{\nabla}\cdot\boldsymbol{v}=0$. Calculus of variations provides the following identity:

$$\left[\frac{d}{\mathrm{d}\epsilon}\mathcal{R}\left[\boldsymbol{u}+\epsilon\boldsymbol{v}\right]\right]_{\epsilon=0} = \left\langle\frac{\delta\mathcal{R}}{\delta\boldsymbol{u}},\ \boldsymbol{v}\right\rangle_{\Omega_t}. \tag{A4}$$

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

To derive the expression for $\delta\mathcal{R}/\delta\boldsymbol{u}$, it is necessary to propagate the perturbation $\boldsymbol{v}$ through all the variables that depend on $\boldsymbol{u}$. The local residual perturbation is

$$\boldsymbol{r}\left(\boldsymbol{u}+\epsilon\boldsymbol{v}\right)=\boldsymbol{r}+\boldsymbol{\delta r}=\frac{\partial\boldsymbol{u}}{\partial t}+\epsilon\frac{\partial\boldsymbol{v}}{\partial t}-\mathcal{N}\left(\boldsymbol{u}\right)-\delta\mathcal{N}-\boldsymbol{\nabla}p-\epsilon\boldsymbol{\nabla}\delta p, \tag{A5}$$

where

$$\mathcal{N}\left(\boldsymbol{u}\right)=-\left(\boldsymbol{u}\cdot\boldsymbol{\nabla}\right)\boldsymbol{u}+\frac{1}{Re}\boldsymbol{\Delta u}. \tag{A6}$$

Rearranging for $\boldsymbol{\delta r}$ gives

$$\boldsymbol{\delta r}=\epsilon\left(\frac{\partial\boldsymbol{v}}{\partial t}-\boldsymbol{\nabla}\delta p\right)-\delta\mathcal{N}. \tag{A7}$$

The perturbation term $\delta\mathcal{N}$ can be computed as

$$\begin{aligned}\mathcal{N}\left(\boldsymbol{u}+\epsilon\boldsymbol{v}\right)&=\mathcal{N}+\delta\mathcal{N}\\&=-\left(\left(\boldsymbol{u}+\epsilon\boldsymbol{v}\right)\cdot\boldsymbol{\nabla}\right)\left(\boldsymbol{u}+\epsilon\boldsymbol{v}\right)+\frac{1}{Re}\boldsymbol{\Delta}\left(\boldsymbol{u}+\epsilon\boldsymbol{v}\right)\\&=-\left(\boldsymbol{u}\cdot\boldsymbol{\nabla}\right)\boldsymbol{u}+\frac{1}{Re}\boldsymbol{\Delta u}-\epsilon\left(-\left(\boldsymbol{u}\cdot\boldsymbol{\nabla}\right)\boldsymbol{v}-\left(\boldsymbol{v}\cdot\boldsymbol{\nabla}\right)\boldsymbol{u}+\frac{1}{Re}\boldsymbol{\Delta v}\right)+\mathcal{O}\left(\epsilon^2\right),\end{aligned} \tag{A8}$$

which when rearranged for $\delta\mathcal{N}$, and neglecting the ϵ^2 term, gives

$$\delta\mathcal{N}=\epsilon\left(-\left(\boldsymbol{u}\cdot\boldsymbol{\nabla}\right)\boldsymbol{v}-\left(\boldsymbol{v}\cdot\boldsymbol{\nabla}\right)\boldsymbol{u}+\frac{1}{Re}\boldsymbol{\Delta v}\right)+\mathcal{O}\left(\epsilon^2\right). \tag{A9}$$

The perturbation in the global residual defined in (A1), called the first variation in $\mathcal{R}$, is given by

$$\begin{aligned}\mathcal{R}\left[\boldsymbol{u}+\epsilon\boldsymbol{v}\right]&=\frac{1}{2}\|\boldsymbol{r}+\boldsymbol{\delta r}\|^2_{\Omega_t}+\langle q,\ \boldsymbol{\nabla}\cdot\left(\boldsymbol{u}+\epsilon\boldsymbol{v}\right)\rangle_{\Omega_t}\\&=\frac{1}{2}\|r\|^2_{\Omega_t}+\langle\boldsymbol{r},\ \boldsymbol{\delta r}\rangle_{\Omega_t}+\frac{1}{2}\langle\boldsymbol{\delta r},\ \boldsymbol{\delta r}\rangle_{\Omega_t}+\langle q,\ \boldsymbol{\nabla}\cdot\boldsymbol{u}\rangle_{\Omega_t}+\epsilon\langle q,\ \boldsymbol{\nabla}\cdot\boldsymbol{v}\rangle_{\Omega_t}.\end{aligned} \tag{A10}$$

Substituting (A7) and (A9) into (A10), neglecting all the ϵ^2 terms, gives

$$\begin{aligned}\mathcal{R}\left[\boldsymbol{u}+\epsilon\boldsymbol{v}\right]=&\frac{1}{2}\|\boldsymbol{r}\|^2_{\Omega_t}+\langle q,\ \boldsymbol{\nabla}\cdot\boldsymbol{u}\rangle_{\Omega_t}\\&+\epsilon\left(\left\langle\boldsymbol{r},\ \frac{\partial\boldsymbol{v}}{\partial t}-\left(\boldsymbol{u}\cdot\boldsymbol{\nabla}\right)\boldsymbol{v}-\left(\boldsymbol{v}\cdot\boldsymbol{\nabla}\right)\boldsymbol{u}+\frac{1}{Re}\boldsymbol{\Delta v}-\boldsymbol{\nabla}\delta p\right\rangle_{\Omega_t}+\langle q,\ \boldsymbol{\nabla}\cdot\boldsymbol{v}\rangle_{\Omega_t}\right)\\&+\mathcal{O}\left(\epsilon^2\right).\end{aligned} \tag{A11}$$

Then, substituting (A11) into (A4) provides a relationship for the functional derivative:

$$\left\langle\frac{\delta\mathcal{R}}{\delta\boldsymbol{u}},\ \boldsymbol{v}\right\rangle_{\Omega_t}=\left\langle\boldsymbol{r},\ \frac{\partial\boldsymbol{v}}{\partial t}-\left(\boldsymbol{u}\cdot\boldsymbol{\nabla}\right)\boldsymbol{v}-\left(\boldsymbol{v}\cdot\boldsymbol{\nabla}\right)\boldsymbol{u}+\frac{1}{Re}\boldsymbol{\Delta v}-\boldsymbol{\nabla}\delta p\right\rangle_{\Omega_t}+\langle q,\ \boldsymbol{\nabla}\cdot\boldsymbol{v}\rangle_{\Omega_t}. \tag{A12}$$

To obtain a closed-form expression for $\delta\mathcal{R}/\delta\boldsymbol{u}$, it is necessary to derive the adjoint of all the operators on the right-hand side of (A12) such that it resembles the left-hand side.

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

Starting with the time derivative,

$$\left\langle \boldsymbol{r},\ \frac{\partial \boldsymbol{v}}{\partial t}\right\rangle_{\Omega_t} = \left\langle -\frac{\partial \boldsymbol{r}}{\partial t},\ \boldsymbol{v}\right\rangle_{\Omega_t}, \tag{A13}$$

where integration by parts has been used, with the boundary terms (at $t=0$ and $t=T$) cancelling due to the flow being periodic in time. To simplify the derivation, we make use of the product rule for divergence

$$\boldsymbol{\nabla}\cdot(u\boldsymbol{v}) = u\boldsymbol{\nabla}\cdot\boldsymbol{v} + \boldsymbol{v}\cdot\boldsymbol{\nabla}u, \tag{A14}$$

as well as the divergence theorem

$$\int_{\Omega} \boldsymbol{\nabla}\cdot\boldsymbol{u}\,\mathrm{d}V = \int_{\partial\Omega} \boldsymbol{u}\cdot\hat{\boldsymbol{n}}\,\mathrm{d}S, \tag{A15}$$

where $\hat{\mathbf{n}}$ is an outwardly pointing normal vector on the surface $\partial\Omega$, and $\mathrm{d}V$ and $\mathrm{d}S$ are infinitesimal volume and surface elements, respectively. Note that, for any field $\boldsymbol{u}\in\mathcal{P}_{T,HBC}$, the surface integral vanishes, i.e.

$$\int_{\partial\Omega} \boldsymbol{u}\cdot\hat{\boldsymbol{n}}\,\mathrm{d}S = 0, \tag{A16}$$

due to all the periodic boundary terms cancelling and the homogeneous no-slip boundary terms being equal to zero. Thus, using (A14) and (A15) for the first of the convective terms, we have

$$\langle \boldsymbol{r},\ (\boldsymbol{u}\cdot\boldsymbol{\nabla})\,\boldsymbol{v}\rangle_{\Omega_t} = \int_0^T \left(\int_{\partial\Omega} (\boldsymbol{r}\cdot\boldsymbol{v})\,\boldsymbol{u}\cdot\hat{\boldsymbol{n}}\,\mathrm{d}S - \int_{\Omega} ((\boldsymbol{u}\cdot\boldsymbol{\nabla})\,\boldsymbol{r})\cdot\boldsymbol{v}\,\mathrm{d}V\right)\mathrm{d}t. \tag{A17}$$

Using the fact that $\boldsymbol{v}\in\mathcal{P}_{T,HBC}$ and $\boldsymbol{u}\in\mathcal{P}_{T,BC}$, the surface integral vanishes, leaving the following:

$$\langle \boldsymbol{r},\ (\boldsymbol{u}\cdot\boldsymbol{\nabla})\,\boldsymbol{v}\rangle_{\Omega_t} = \langle -(\boldsymbol{u}\cdot\boldsymbol{\nabla})\,\boldsymbol{r},\ \boldsymbol{v}\rangle_{\Omega_t}. \tag{A18}$$

The second convective term is slightly simpler due to $\boldsymbol{v}$ not appearing in any derivatives, i.e.

$$\langle \boldsymbol{r},\ (\boldsymbol{v}\cdot\boldsymbol{\nabla})\,\boldsymbol{u}\rangle_{\Omega_t} = \langle \boldsymbol{r},\ (\boldsymbol{\nabla}\boldsymbol{u})\,\boldsymbol{v}\rangle_{\Omega_t} = \left\langle (\boldsymbol{\nabla}\boldsymbol{u})^{\top}\,\boldsymbol{r},\ \boldsymbol{v}\right\rangle_{\Omega_t}, \tag{A19}$$

where $\boldsymbol{\nabla}\boldsymbol{u}$ is the gradient of the vector field $\boldsymbol{u}$. In Cartesian coordinates this looks like the gradient of the scalar components on $\boldsymbol{u}$ stacked on top of each other to form a matrix. Next, the diffusion term

$$\langle \boldsymbol{r},\ \boldsymbol{\Delta}\boldsymbol{v}\rangle_{\Omega_t} = \int_0^T \left(\int_{\partial\Omega} (\boldsymbol{r}\cdot\boldsymbol{\nabla}\boldsymbol{v} - \boldsymbol{\nabla}\boldsymbol{r}\cdot\boldsymbol{v})\cdot\hat{\boldsymbol{n}}\,\mathrm{d}S + \int_{\Omega} \boldsymbol{\Delta}\boldsymbol{r}\cdot\boldsymbol{v}\,\mathrm{d}V\right)\mathrm{d}t, \tag{A20}$$

where the product rule and divergence theorem have been used twice to effectively integrate by parts the second-order derivatives in the Laplacian. The second of the boundary integrals vanish due to $\boldsymbol{v}\in\mathcal{P}_{T,HMB}$, however, the other boundary integral does not by default since we have not imposed any extra restrictions on the gradients of $\boldsymbol{v}$. To resolve this, it is necessary to impose $\boldsymbol{r}\in\mathcal{P}_{T,HBC}$, enforcing the fact that $\boldsymbol{r}$ vanishes at the walls, which means the first boundary integral also vanishes, giving

$$\langle \boldsymbol{r},\ \boldsymbol{\Delta}\boldsymbol{v}\rangle_{\Omega_t} = \langle \boldsymbol{\Delta}\boldsymbol{r},\ \boldsymbol{v}\rangle_{\Omega_t}. \tag{A21}$$

Moving to the pressure gradient we have

$$\langle \boldsymbol{r},\ \boldsymbol{\nabla}\delta p\rangle_{\Omega_t} = \int_0^T \left(\int_{\partial\Omega} \delta p\,(\boldsymbol{r}\cdot\hat{\boldsymbol{n}})\,\mathrm{d}S - \int_{\Omega} (\boldsymbol{\nabla}\cdot\boldsymbol{r})\,\delta p\,\mathrm{d}V\right)\mathrm{d}t. \tag{A22}$$

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

Since we know that $\boldsymbol{r} \in \mathcal{P}_{T,HBC}$ as prescribed by the adjoint of the diffusion term, the boundary term here must vanish, leaving

$$\langle \boldsymbol{r},\ \nabla \delta p \rangle_{\Omega_t} = \langle -\nabla \cdot \boldsymbol{r},\ \delta p \rangle_{\Omega_t}\,. \tag{A23}$$

The divergence of the residual is actually the pressure-Poisson equation, and since the pressure field is assumed to satisfy this equation we also know that $\nabla \cdot \boldsymbol{r} = 0$, and thus, (A23) is also zero. This just leaves the divergence term for the Lagrange multiplier, enforcing the divergence-free evolution of the velocity field $\boldsymbol{u}$ under the optimisation problem given in (A1). Due to the symmetry of the inner product (in the case of real fields), it is easy to see that the adjoint of the divergence operator is the negative gradient operator. Applying this result gives

$$\langle q,\ \nabla \cdot \boldsymbol{v} \rangle_{\Omega_t} = \int_0^T \int_\Omega -\nabla q \cdot \boldsymbol{v} \mathrm{d}V \mathrm{d}t = \langle -\nabla q,\ \boldsymbol{v} \rangle_{\Omega_t}\,. \tag{A24}$$

Combining all of these results and substituting them into the right-hand side of (A12) results in

$$\left\langle \frac{\delta \mathcal{R}}{\delta \boldsymbol{u}},\ \boldsymbol{v} \right\rangle_{\Omega_t} = \left\langle -\frac{\partial \boldsymbol{r}}{\partial t} - (\boldsymbol{u} \cdot \nabla)\,\boldsymbol{r} + (\nabla \boldsymbol{u})^\top \boldsymbol{r} - \frac{1}{Re}\boldsymbol{\Delta r} - \nabla q,\ \boldsymbol{v} \right\rangle_{\Omega_t}, \tag{A25}$$

which provides the desired closed-form expression for the functional derivative of $\mathcal{R}$ with respect to $\boldsymbol{u}$, i.e.

$$\frac{\delta \mathcal{R}}{\delta \boldsymbol{u}} = -\frac{\partial \boldsymbol{r}}{\partial t} - (\boldsymbol{u} \cdot \nabla)\,\boldsymbol{r} + (\nabla \boldsymbol{u})^\top \boldsymbol{r} - \frac{1}{Re}\boldsymbol{\Delta r} - \nabla q, \tag{A26}$$

with the additional constraints $\boldsymbol{r} \in \mathcal{P}_{T,HBC}$ and $\nabla \cdot \boldsymbol{r} = 0$. The boundary constraints on $\boldsymbol{r}$ at the boundaries is an instance of natural BCs, and are necessary to enforce. If they are not strictly enforced then the residual gradient given in (A26) is not guaranteed to be a descent direction. A similar interpretation applies to $\nabla \cdot \boldsymbol{r} = 0$, except it applies throughout the whole domain rather than just the boundaries.

A.2. *Fundamental frequency derivative*

The derivative of $\mathcal{R}$ with respect to the fundamental frequency ω is simpler to derive since $\mathcal{R}$ is an ordinary function of ω rather than a functional such as with $\boldsymbol{u}$. First, define a new scaled time $t' = \omega t$ that then leads to the following relationship between the time derivatives $\partial/\partial t = \omega \partial/\partial t'$. Using this scaled time in the definition of the local residual in (2.5) and then substituting the resulting expression into the definition of the global residual gives

$$\begin{aligned}\mathcal{R} = \frac{\omega^2}{2}\left\| \frac{\partial \boldsymbol{u}}{\partial t'} \right\|_{\Omega_t}^2 - \omega \left\langle \frac{\partial \boldsymbol{u}}{\partial t'},\ -(\boldsymbol{u} \cdot \nabla \boldsymbol{u}) + \frac{1}{Re}\boldsymbol{v} - \nabla p \right\rangle_{\Omega_t} \\ + \frac{1}{2}\left\| -(\boldsymbol{u} \cdot \nabla \boldsymbol{u}) + \frac{1}{Re}\boldsymbol{v} - \nabla p \right\|_{\Omega_t}^2.\end{aligned} \tag{A27}$$

Fixing the particular velocity field $\boldsymbol{u}$ under consideration makes this a quadratic function, with a single global minimum at some ω^*. This global minimum could be derived rather simply, however using it directly during the optimisation was found to lead to very erratic behaviour, especially when far away from a solution. Taking the derivative of (A27) with

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

respect to ω gives

$$\frac{\partial \mathcal{R}}{\partial \omega} = \omega \left\| \frac{\partial \boldsymbol{u}}{\partial t'} \right\|^2_{\Omega_t} - \left\langle \frac{\partial \boldsymbol{u}}{\partial t'}, \, -(\boldsymbol{u} \cdot \boldsymbol{\nabla} \boldsymbol{u}) + \frac{1}{Re} \boldsymbol{v} - \boldsymbol{\nabla} p \right\rangle_{\Omega_t}. \tag{A28}$$

With a slight rearrangement the following simplified expression is obtained:

$$\frac{\partial \mathcal{R}}{\partial \omega} = \frac{1}{\omega} \left\langle \frac{\partial \boldsymbol{u}}{\partial t}, \, \boldsymbol{r} \right\rangle_{\Omega_t}. \tag{A29}$$

A similar expression can be obtained for the derivative with respect to the solution period $T = 2\pi/\omega$ using the chain rule

$$\frac{\partial \mathcal{R}}{\partial T} = -\frac{1}{T} \left\langle \frac{\partial \boldsymbol{u}}{\partial t}, \, \boldsymbol{r} \right\rangle_{\Omega_t}. \tag{A30}$$

Appendix B. Differential-algebraic equation form of variational optimisation

In the literature (Farazmand 2016; Azimi *et al.* 2022; Ashtari & Schneider 2023) the variational optimisation is derived using an adjoint variational dynamics formulation. Using the same notation, where τ represents the fictitious time introduced to parametrise the evolution of the variational dynamics, we can say that

$$\frac{\partial \boldsymbol{u}}{\partial \tau} := -\frac{\delta \mathcal{R}}{\delta \boldsymbol{u}}, \quad \frac{\mathrm{d}\omega}{\mathrm{d}\tau} := -\frac{\partial \mathcal{R}}{\partial \omega}, \tag{B1}$$

Using this definition the gradient-based optimisation problem can be restated in the form of a differential-algebraic equation for the evolution of the spatio-temporal flow field:

$$\frac{\partial \boldsymbol{u}}{\partial \tau} = \frac{\partial \boldsymbol{r}}{\partial t} + (\boldsymbol{u} \cdot \boldsymbol{\nabla}) \, \boldsymbol{r} - (\boldsymbol{\nabla} \boldsymbol{u})^\top \boldsymbol{r} + \frac{1}{Re} \boldsymbol{\Delta} \boldsymbol{r} - \boldsymbol{\nabla} q, \tag{B2a}$$

$$\frac{\partial \omega}{\partial \tau} = -\frac{1}{\omega} \left\langle \frac{\partial \boldsymbol{u}}{\partial t}, \, \boldsymbol{r} \right\rangle_{\Omega_t}, \tag{B2b}$$

$$0 = \frac{\partial \boldsymbol{u}}{\partial t} + \boldsymbol{\nabla} p + (\boldsymbol{u} \cdot \boldsymbol{\nabla}) \, \boldsymbol{u} - \frac{1}{Re} \boldsymbol{\Delta} \boldsymbol{u} - \boldsymbol{r}, \tag{B2c}$$

$$0 = \boldsymbol{\nabla} \cdot \boldsymbol{u}, \tag{B2d}$$

$$0 = \boldsymbol{\nabla} \cdot \boldsymbol{r}, \tag{B2e}$$

$$0 = \boldsymbol{u}|_{y=\pm 1} \mp \boldsymbol{e}_x, \tag{B2f}$$

$$0 = \boldsymbol{r}|_{y=\pm 1}, \tag{B2g}$$

In the work of Azimi *et al.* (2022), Ashtari & Schneider (2023), (B2) is solved using a simple first-order Euler scheme to explicitly step forward the solution by $\Delta\tau$. The use of the Euler method was justified as the accuracy of the intermediate result does not matter, just the accuracy of the solution at the end once $\mathcal{R}$ becomes sufficiently small. This allowed for larger step sizes in τ. Farazmand (2016) uses a Runge–Kutta time-stepping method in τ to improve accuracy, but this still suffers from the same slow convergence expected from gradient-descent methods. However, as shown here, the problem can be viewed simply as an optimisation problem, where the goal is to minimise $\mathcal{R}$ as much as possible regardless of the route taken in the loop space to do so. This means the intermediate accuracy of any time-stepping scheme in $\Delta\tau$ is not productive. Instead, more broad optimisation algorithms can be used that trade the exact evolution prescribed by (B2*a*) and (B2*b*) for convergence to smaller residuals in far fewer iterations. This is

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

displayed in § 5.2 where gradient descent (mathematically equivalent to the Euler method in τ) is compared with common quasi-Newton optimisation algorithms.

Appendix C. Resolvent for 2D3C RPCF

In this work we use resolvent analysis to generate a divergence-free and no-slip orthonormal basis set for the velocity and local residual fields. The exact form of the operators defined in § 3 for 2D3C RPCF is given here. First, the (nonlinear) Navier–Stokes operator is defined as

$$\mathcal{N}(\boldsymbol{u}_b) = -(\boldsymbol{u}_b \cdot \boldsymbol{\nabla})\,\boldsymbol{u}_b + \frac{1}{Re}\boldsymbol{\Delta}\boldsymbol{u}_b - Ro\,(\boldsymbol{e}_z \times \boldsymbol{u}_b)\,, \tag{C1}$$

where the base pressure field p_b is not included since it is projected away by the Leray projector $\mathbb{P}$ in (3.2). The linearised Navier–Stokes operator evaluated at the base flow $\boldsymbol{u}_b$, $\mathcal{L}_{\boldsymbol{u}_b}$, is given by

$$\mathcal{L}_{\boldsymbol{u}_b}\boldsymbol{u}' = -(\boldsymbol{u}_b \cdot \boldsymbol{\nabla})\,\boldsymbol{u}' - (\boldsymbol{\nabla}\boldsymbol{u}_b)\,\boldsymbol{u}' + \frac{1}{Re}\boldsymbol{\Delta}\boldsymbol{u}' - Ro\,\boldsymbol{e}_z \times \boldsymbol{u}', \tag{C2}$$

where again the pressure is not included due to the Leray projector in (3.2). Assuming that $\boldsymbol{u}_b(y) = u_b(y)\boldsymbol{e}_x$, and using the definition of the resolvent in § 3, the final expression for the operator for 2D3C RPCF is expressed as

$$\boldsymbol{R}_{\boldsymbol{k}} = \begin{pmatrix} ik_t\omega - \dfrac{1}{Re}\boldsymbol{\Delta}_{\boldsymbol{k}} & \dfrac{\partial u_b}{\partial y} - Ro & 0 \\ Ro & ik_t\omega - \dfrac{1}{Re}\boldsymbol{\Delta}_{\boldsymbol{k}} & 0 \\ 0 & 0 & ik_t\omega - \dfrac{1}{Re}\boldsymbol{\Delta}_{\boldsymbol{k}} \end{pmatrix}^{-1} \mathbb{P}, \tag{C3}$$

where $\boldsymbol{\Delta}_{\boldsymbol{k}} = \partial^2/\partial y^2 - (\beta k_z)^2$.

Appendix D. DNS details

The velocity vector field is conveniently decomposed into the laminar base state and the velocity fluctuation, denoted in this appendix as $\boldsymbol{u}(y,\ z,\ t) = (u,\ v,\ w)$. A formulation using the streamwise component of the vorticity perturbation $\omega_x = \partial w/\partial y - \partial v/\partial z$, and the streamwise (fluctuation) velocity u, is used to numerically solve the problem (4.1). This formulation eliminates complications arising from the pressure term. Because the domain is 2D the streamfunction ψ, related to the transverse and wall-normal (fluctuation) velocity components by the relations $w = -\partial\psi/\partial y$ and $v = \partial\psi/\partial z$, and to the vorticity by the Poisson-type equation $\boldsymbol{\Delta}\psi = -\omega$, is introduced. Introduction of the streamfunction ensures that the continuity equation is satisfied. With these definitions, the coupled system of nonlinear partial differential equations (PDEs)

$$\frac{\partial u}{\partial t} = -\frac{\partial\psi}{\partial z}\frac{\partial u}{\partial y} + \frac{\partial\psi}{\partial y}\frac{\partial u}{\partial z} + \frac{\partial\psi}{\partial z}(Ro - 1) + \frac{1}{Re}\boldsymbol{\Delta}u, \tag{D1a}$$

$$\frac{\partial\omega}{\partial t} = -\frac{\partial\psi}{\partial z}\frac{\partial\omega}{\partial y} + \frac{\partial\psi}{\partial y}\frac{\partial\omega}{\partial z} - \frac{\partial u}{\partial z}Ro + \frac{1}{Re}\boldsymbol{\Delta}\omega \tag{D1b}$$

governs the evolution of the streamwise (fluctuation) velocity and vorticity. A spectral spatial discretisation of (D1) is adopted along the axial direction by introducing the

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

truncated Fourier expansion

$$f(z,\, y,\, t) = \sum_{k=-N/2}^{N/2-1} f_k(y,\, t)\exp(ik_z\beta z) \tag{D2}$$

for u, ω and ψ, with $\beta = 2\pi/L_z$, where k_z is the spanwise wavenumber. Projection of the governing equations (D1) onto (D2) leads to a system of sets of three one-dimensional coupled PDEs in the wall-normal direction y for each wavenumber k_z. This system of PDEs is then discretised in time using a fairly standard splitting approach. The viscous term is treated implicitly using a Crank–Nicholson method, while the nonlinear term and the term arising from the rotation and the mean flow shear are integrated explicitly via a second-order accurate Adams–Bashforth method. The nonlinear term is treated pseudo-spectrally, see Canuto *et al.* (1988), whereas de-aliasing was not found necessary, given the low-resolution requirements. Temporal discretisation results in a system of ordinary differential equations in the space variable y, having the form of Helmoltz problems for each wavenumber k_z and for the three flow variables. These are solved at each time step using a second-order centred finite difference scheme, leading to a sparse banded system solved directly by a Gaussian elimination method with pivoting.

Appendix E. Hessian operator derivation

To derive the Hessian operator, first take the definition of the local residual in (2.5) and project it onto the divergence-free subspace using the Leray projector from § 3,

$$\boldsymbol{r} = \frac{\partial \boldsymbol{u}}{\partial t} - \mathbb{P}\mathcal{N}\left(\boldsymbol{u}\right), \tag{E1}$$

where the local residual itself is unaffected by the projection since it is divergence free by construction. Consider a solution to the optimisation problem denoted as $\boldsymbol{u}^*$. A perturbation to the solution $\boldsymbol{u}^*$ by some arbitrary velocity field $\boldsymbol{V}$ that is incompressible and obeys the appropriate BCs is applied. The global residual $\mathcal{R}$ can then be expanded at the minimum using a Taylor series

$$\begin{aligned}\mathcal{R}\left[\boldsymbol{u}^* + \epsilon\boldsymbol{v}\right] &= \mathcal{R}\left[\boldsymbol{u}^*\right] + \epsilon\left\langle\frac{\delta\mathcal{R}}{\delta\boldsymbol{u}^*},\ \boldsymbol{v}\right\rangle_{\Omega_t} + \frac{\epsilon^2}{2}\left\langle\boldsymbol{v},\ \boldsymbol{H}\boldsymbol{v}\right\rangle_{\Omega_t} + \mathcal{O}\left(\epsilon^3\right)\\ &= \frac{\epsilon^2}{2}\left\langle\boldsymbol{v},\ \boldsymbol{H}\boldsymbol{v}\right\rangle_{\Omega_t} + \mathcal{O}\left(\epsilon^3\right),\end{aligned} \tag{E2}$$

where the fact that $\boldsymbol{u}^*$ is a minimum means that $\mathcal{R}[\boldsymbol{u}^*] = \|\delta\mathcal{R}/\delta\boldsymbol{u}^*\| = 0$. We then have a local quadratic model for the variation of $\mathcal{R}$ that is governed by the Hessian operator $\boldsymbol{H}$. The second-order variation of the global residual is given by $\mathcal{R}[\boldsymbol{u}^* + \epsilon\boldsymbol{v}] = \|\delta\boldsymbol{r}\|^2_{\Omega_t}/2$. Linearising the local residual in (E1) and substituting this into the previous expressions gives

$$\left\langle\boldsymbol{v},\ \boldsymbol{H}\boldsymbol{v}\right\rangle_{\Omega_t} = \left\|\frac{\partial\boldsymbol{v}}{\partial t} - \mathbb{P}\mathcal{L}_{\boldsymbol{u}^*}\boldsymbol{v}\right\|^2_{\Omega_t}. \tag{E3}$$

References

AHMED, M.A. & SHARMA, A.S. 2020 Basis for finding exact coherent states. *Phys. Rev. E* **101**, 12213.

AIZAWA, Y. 1982 Global aspects of the dissipative dynamical systems. I: — statistical identification and fractal properties of the Lorenz chaos. *Prog. Theor. Phys.* **68**, 64–84.

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

ARTUSO, R., AURELL, E. & CVITANOVIC, P. 1990 Recycling of strange sets: I. Cycle expansions. *Nonlinearity* **3**, 325–359.
ASHTARI, O. & SCHNEIDER, T.M. 2023 Identifying invariant solutions of wall-bounded three-dimensional shear flows using robust adjoint-based variational techniques. *J. Fluid Mech.* **977**, A7.
AUERBACH, D., CVITANOVIC, P., ECKMANN, J.P., GUNARATNE, G. & PROCACCIA, I. 1987 Exploring chaotic motion through periodic orbits. *Phys. Rev. Lett.* **58**, 2387–2389.
AZIMI, S., ASHTARI, O. & SCHNEIDER, T.M. 2022 Constructing periodic orbits of high-dimensional chaotic systems by an adjoint-based variational method. *Phys. Rev. E* **105**, 14217.
BARTHEL, B., ZHU, X. & MCKEON, B.J. 2021 Closing the loop: nonlinear Taylor vortex flow through the lens of resolvent analysis. *J. Fluid Mech.* **924**, A9.
BENEDDINE, S., YEGAVIAN, R., SIPP, D. & LECLAIRE, B. 2017 Unsteady flow dynamics reconstruction from mean flow and point sensors: an experimental study. *J. Fluid Mech.* **824**, 174–201.
BRANDT, L., SIPP, D., PRALITS, J.O. & MARQUET, O. 2011 Effect of base-flow variation in noise amplifiers: the flat-plate boundary layer. *J. Fluid Mech.* **687**, 503–528.
BURTON, T. 2025 Resolvent-based models for nonlinear solutions of wall-bounded flows and statistical estimation of chaotic systems. PhD thesis. University of Southampton, Southampton.
BURTON, T., SYMON, S., SHARMA, A.S. & LASAGNA, D. 2025 Resolvent-based optimization for approximating the statistics of a chaotic Lorenz system. *Phys. Rev. E* **111**, 25104.
CANUTO, C., HUSSAINI, M.Y., QUARTERONI, A. & ZANG, T.A. 1988 *Spectral Methods in Fluid Dynamics*. Springer.
CHANDLER, G.J. & KERSWELL, R.R. 2013 Invariant recurrent solutions embedded in a turbulent two-dimensional Kolmogorov flow. *J. Fluid Mech.* **722**, 554–595.
CHANDRASEKHAR, S. 1961 *Hydrodynamic and Hydromagnetic Stability*. Oxford University Press.
CHRISTIANSEN, F., CVITANOVIC, P. & PUTKARADZE, V. 1997 Spatiotemporal chaos in terms of unstable recurrent patterns. *Nonlinearity* **10**, 55–70.
CROWLEY, C.J., PUGHE-SANFORD, J.L., TOLER, W., KRYGIER, M.C., GRIGORIEV, R.O. & SCHATZ, M.F. 2022 Turbulence tracks recurrent solutions. *Proc. Natl Acad. Sci. USA* **119**, e2120665119.
CVITANOVIĆ, P. 1988 Invariant measurement of strange sets in terms of cycles. *Phys. Rev. Lett.* **61**, 2729–2732.
CVITANOVIĆ, P. 1995 Dynamical averaging in terms of periodic orbits. *Phys. D: Nonlinear Phenom.* **83**, 109–123.
CVITANOVIĆ, P. & GIBSON, J.F. 2010 Geometry of the turbulence in wall-bounded shear flows: periodic orbits. *Phys. Scr.* **T142**, 014007.
DAHLQUIST, G. & BJÖRCK, Å. 2008 *Numerical Methods in Scientific Computing*, vol. I. Society for Industrial and Applied Mathematics.
DAVIAUD, F., HEGSETH, J. & BERGÉ, P. 1992 Subcritical transition to turbulence in plane Couette flow. *Phys. Rev. Lett.* **69**, 2511–2514.
DENNIS, J.E. & SCHNABEL, B. 1996 *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*. Society for Industrial and Applied Mathematics.
ECKHARDT, B., DOERING, C.R. & WHITEHEAD, J.P. 2020 Exact relations between Rayleigh–Bénard and rotating plane Couette flow in two dimensions. *J. Fluid Mech.* **903**, R4.
ECKHARDT, B., GROSSMANN, S. & LOHSE, D. 2007 Torque scaling in turbulent Taylor–Couette flow between independently rotating cylinders. *J. Fluid Mech.* **581**, 221–250.
FAISST, H. & ECKHARDT, B. 2003 Traveling waves in pipe flow. *Phys. Rev. Lett.* **91**, 224502.
FARAZMAND, M. 2016 An adjoint-based approach for finding invariant solutions of Navier–Stokes equations. *J. Fluid Mech.* **795**, 278–312.
FAZENDEIRO, L., BOGHOSIAN, B.M., COVENEY, P.V. & LÄTT, J. 2010 Unstable periodic orbits in weak turbulence. *J. Comput. Sci.* **1**, 13–23.
FLETCHER, R. 2000 *Structure of Methods*, pp. 12–43. John Wiley & Sons, Ltd.
FRIGO, M. & JOHNSON, S.G. 2005 The design and implementation of FFTW3. *Proc. IEEE* **93**, 216–231.
GAO, Y., WU, Y. & XIA, J. 2017 Automatic differentiation based discrete adjoint method for aerodynamic design optimization on unstructured meshes. *Chin. J. Aeronaut.* **30** (2), 611–627.
GARNAUD, X., LESSHAFFT, L., SCHMID, P.J. & HUERRE, P. 2013 The preferred mode of incompressible jets: linear frequency response analysis. *J. Fluid Mech.* **716**, 189–202.
GAYME, D. 2010 A robust control approach to understanding nonlinear mechanisms in shear flow turbulence. PhD thesis, California Institute of Technology, CA, USA.
GAYME, D.F. & MINNICK, B.A. 2019 Coherent structure-based approach to modeling wall turbulence. *Phys. Rev. Fluids* **4**, 110505.
GIBSON, J.F., HALCROW, J. & CVITANOVIC, P. 2008 Visualizing the geometry of state space in plane Couette flow. *J. Fluid Mech.* **611**, 107–130.

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

GÓMEZ, F., BLACKBURN, H.M., RUDMAN, M., SHARMA, A.S. & MCKEON, B.J. 2016 A reduced-order model of three-dimensional unsteady flow in a cavity based on the resolvent operator. *J. Fluid Mech.* **798**, R2.

HAGER, W.W. & ZHANG, H. 2005 A new conjugate gradient method with guaranteed descent and an efficient line search. *SIAM J. Optim.* **16**, 170–192.

HEKMAT, M.H. & MIRZAEI, M. 2016 A comparison of the continuous and discrete adjoint approach extended based on the standard lattice Boltzmann method in flow field inverse optimization problems. *Acta Mechanica* **227** (4), 1025–1050.

HIWATASHI, K., ALFREDSSON, P.H., TILLMARK, N. & NAGATA, M. 2007 Experimental observations of instabilities in rotating plane Couette flow. *Phys. Fluids* **19**, 048103.

HOPF, E. 1942 Abzweigung einer periodischen lösung von einer stationären lösung eines differentialsystems. *Berichte Verhandlungen Sächsischen Akad. Wissensch. Leipzig Math.-Naturwissenschaftliche Klasse* **94**, 1–22.

HOPF, E. 1948 A mathematical example displaying features of turbulence. *Commun. Pure Appl. Maths* **1**, 303–322.

ITANO, T. & GENERALIS, S.C. 2009 Hairpin vortex solution in planar Couette flow: a tapestry of knotted vortices. *Phys. Rev. Lett.* **102**, 114501.

JIN, B., ILLINGWORTH, S.J. & SANDBERG, R.D. 2022 Resolvent-based approach for $\{H_2\}$-optimal estimation and control: an application to the cylinder flow. *Theor. Comput. Fluid Dyn.* **36**, 491–515.

KAWAHARA, G. & KIDA, S. 2001 Periodic motion embedded in plane Couette turbulence: regeneration cycle and burst. *J. Fluid Mech.* **449**, 291–300.

KAWAHARA, G., UHLMANN, M. & VAN VEEN, L. 2012 The significance of simple invariant solutions in turbulent flows. *Annu. Rev. Fluid Mech.* **44**, 203–225.

KLEISER, L. & SCHUMANN, U. 1980 *Treatment of Incompressibility and Boundary Conditions in 3-D Numerical Spectral Simulations of Plane Channel Flows*, pp. 165–173. Vieweg+Teubner Verlag.

KRYGIER, M.C., PUGHE-SANFORD, J.L. & GRIGORIEV, R.O. 2021 Exact coherent structures and shadowing in turbulent Taylor–Couette flow. *J. Fluid Mech.* **923**, A7.

LAKOBA, T.I. & YANG, J. 2007 A mode elimination technique to improve convergence of iteration methods for finding solitary waves. *J. Comput. Phys.* **226**, 1693–1709.

LAN, Y. & CVITANOVIĆ, P. 2004 Variational method for finding periodic orbits in a general flow. *Phys. Rev. E - Stat. Phys. Plasmas Fluids Related Interdiscip. Topics* **69**, 10.

LASAGNA, D., TUTTY, O.R. & CHERNYSHENKO, S. 2016 Flow regimes in a simplified Taylor–Couette-type flow model. *Eur. J. Mech. - B/Fluids* **57**, 176–191.

LEZIUS, D.K. & JOHNSTON, J.P. 1976 Roll-cell instabilities in rotating laminar and turbulent channel flows. *J. Fluid Mech.* **77**, 153–174.

LI, X. & LASAGNA, D. 2026 Space-time nonlinear reduced-order modelling for unsteady flows. *J. Fluid Mech.* **1037**, A32.

LIU, D.C. & NOCEDAL, J. 1989 On the limited memory BFGS method for large scale optimization. *Math. Programming* **45**, 503–528.

MCKEON, B. 2023 *Applications of Resolvent Analysis in Fluid Mechanics*. California Institute of Technology.

MCKEON, B.J. & SHARMA, A.S. 2010 A critical-layer framework for turbulent pipe flow. *J. Fluid Mech.* **658**, 336–382.

MOGENSEN, P.K. & RISETH, A.N. 2018 Optim: a mathematical optimization package for Julia. *J. Open Source Softw.* **3**, 615.

MONS, V. & MARQUET, O. 2021 Linear and nonlinear sensor placement strategies for mean-flow reconstruction via data assimilation. *J. Fluid Mech.* **923**, A1.

NADARAJAH, S. & JAMESON, A. 2000 A comparison of the continuous and discrete adjoint approach to automatic aerodynamic optimization. In *38th Aerospace Sciences Meeting and Exhibit, Aerospace Sciences Meetings*. American Institute of Aeronautics and Astronautics.

NAGATA, M. 1990 Three-dimensional finite-amplitude solutions in plane Couette flow: bifurcation from infinity. *J. Fluid Mech.* **217**, 519–527.

NAGATA, M. 1997 Three-dimensional traveling-wave solutions in plane Couette flow. *Phys. Rev. E* **55**, 2023–2025.

NAGATA, M. 2013 A note on the mirror-symmetric coherent structure in plane Couette flow. *J. Fluid Mech.* **727**, R1.

NAGATA, M., SONG, B. & WALL, D.P. 2021 Onset of vortex structures in rotating plane Couette flow. *J. Fluid Mech.* **918**, A2.

NOCEDAL, J. & WRIGHT, S. 2006 *Numerical Optimization*, 2nd edn. Springer-Verlag New York.

OKINO, S., NAGATA, M., WEDIN, H. & BOTTARO, A. 2010 A new nonlinear vortex state in square-duct flow. *J. Fluid Mech.* **657**, 413–429.

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

PAGE, J., NORGAARD, P., BRENNER, M.P. & KERSWELL, R.R. 2024 Recurrent flow patterns as a basis for two-dimensional turbulence: predicting statistics from structures. *Proc. Natl Acad. Sci. USA* **121**, e2320007121.

POPE, S.B. 2001 Turbulent flows. *Meas. Sci. Technol.* **12**, 2020.

PRINGLE, C.C.T. & KERSWELL, R.R. 2007 Asymmetric, helical, and mirror-symmetric traveling waves in pipe flow. *Phys. Rev. Lett.* **99**, 74502.

REDFERN, E.M., LAZER, A.L. & LUCAS, D. 2024 Dynamically relevant recurrent flows obtained via a nonlinear recurrence function from two-dimensional turbulence. *Phys. Rev. Fluids* **9** (12), 124401.

RIGAS, G., SIPP, D. & COLONIUS, T. 2021 Nonlinear input/output analysis: application to boundary layer transition. *J. Fluid Mech.* **911**, A15.

ROSENBERG, K. & MCKEON, B.J. 2019 Computing exact coherent states in channels starting from the laminar profile: a resolvent-based approach. *Phys. Rev. E* **100**, 21101.

ROSENBERG, K. & MCKEON, B.J. 2019 Efficient representation of exact coherent states of the Navier–Stokes equations using resolvent analysis. *Fluid Dyn. Res.* **51**, 011401.

RUELLE, D. & TAKENS, F. 1971 On the nature of turbulence. *Commun. Math. Phys.* **20**, 167–192.

SAAD, Y. & SCHULTZ, M.H. 1986 GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM J. Sci. Stat. Comput.* **7**, 856–869.

SHARMA, A.S. 2019 Elements of resolvent methods in fluid mechanics: notes for an introductory short course v0.3. arXiv pp. 1–24.

SHARMA, A.S., MOARREF, R., MCKEON, B.J., PARK, J.S., GRAHAM, M.D. & WILLIS, A.P. 2016 Low-dimensional representations of exact coherent states of the Navier–Stokes equations from the resolvent model of wall turbulence. *Phys. Rev. E* **93**, 21102.

SIERRA, J., JOLIVET, P., GIANNETTI, F. & CITRO, V. 2021 Adjoint-based sensitivity analysis of periodic orbits by the Fourier–Galerkin method. *J. Comput. Phys.* **440**, 110403.

SIERRA-AUSIN, J., CITRO, V., GIANNETTI, F. & FABRE, D. 2022 Efficient computation of time-periodic compressible flows with spectral techniques. *Comput. Meth. Appl. Mech. Engng* **393**, 114736.

SIPP, D. & MARQUET, O. 2013 Characterization of noise amplifiers with global singular modes: the case of the leading-edge flat-plate boundary layer. *Theor. Comput. Fluid Dyn.* **27** (5), 617–635.

SÁNCHEZ, J., NET, M., GARCÍA-ARCHILLA, B. & SIMÓ, C. 2004 Newton–Krylov continuation of periodic orbits for Navier–Stokes flows. *J. Comput. Phys.* **201**, 13–33.

SURI, B., KAGEORGE, L., GRIGORIEV, R.O. & SCHATZ, M.F. 2020 Capturing turbulent dynamics and statistics in experiments with unstable periodic orbits. *Phys. Rev. Lett.* **125**, 64501.

SYMON, S., ROSENBERG, K., DAWSON, S.T.M. & MCKEON, B.J. 2018 Non-normality and classification of amplification mechanisms in stability and resolvent analysis. *Phys. Rev. Fluids* **3**, 053902.

TEMAM, R. 1984 *Navier–Stokes Equations: Theory and Numerical Analysis*. AMS Chelsea Publishing.

TOMOAKI, I. & SADAYOSHI, T. 2001 The dynamics of bursting process in wall turbulence. *J. Phys. Soc. Japan* **70**, 703–716.

TOWNE, A., SCHMIDT, O.T. & COLONIUS, T. 2018 Spectral proper orthogonal decomposition and its relationship to dynamic mode decomposition and resolvent analysis. *J. Fluid Mech.* **847**, 821–867.

TSUKAHARA, T., TILLMARK, N. & ALFREDSSON, P.H. 2010 Flow regimes in a plane Couette flow with system rotation. *J. Fluid Mech.* **648**, 5–33.

UHLMANN, M., KAWAHARA, G. & PINELLI, A. 2010 Traveling-waves consistent with turbulence-driven secondary flow in a square duct. *Phys. Fluids* **22**, 084102.

VAN VEEN, L., KIDA, S. & KAWAHARA, G. 2006 Periodic motion representing isotropic turbulence. *Fluid Dyn. Res.* **38**, 19–46.

VEEN, L.V., VELA-MARTÍN, A. & KAWAHARA, G. 2019 Time-periodic inertial range dynamics. *Phys. Rev. Lett.* **123**, 134502.

VISWANATH, D. 2007 Recurrent motions within plane Couette turbulence. *J. Fluid Mech.* **580**, 339–358.

VISWANATH, D. 2009 The critical layer in pipe flow at high Reynolds number. *Phil. Trans. R. Soc. Lond. A: Math. Phys. Engng Sci.* **367**, 561–576.

WALEFFE, F. 1998 Three-dimensional coherent states in plane shear flows. *Phys. Rev. Lett.* **81**, 4140–4143.

WALEFFE, F. 2001 Exact coherent structures in channel flow. *J. Fluid Mech.* **435**, 93–102.

WALEFFE, F. 2009 Exact coherent structures in turbulent shear flows. In *Turbulence and Interactions* (ed. M. DEVILLE, T.-H. LÊ & P. SAGAUT), pp. 139–158. Springer Berlin Heidelberg.

WANG, B., AYATS, R., DEGUCHI, K., MESEGUER, A. & MELLIBOVSKY, F. 2025 Mathematically established chaos and forecast of statistics with recurrent patterns in Taylor–Couette flow. *J. Fluid Mech.* **1011**, R2.

WEDIN, H., BOTTARO, A. & NAGATA, M. 2009 Three-dimensional traveling waves in a square duct. *Phys. Rev. E* **79**, 65305.

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press

WEDIN, H. & KERSWELL, R.R. 2004 Exact coherent structures in pipe flow: travelling wave solutions. *J. Fluid Mech.* **508**, 333–371.

WILLIS, A.P., CVITANOVIĆ, P. & AVILA, M. 2013 Revealing the state space of turbulent pipe flow by symmetry reduction. *J. Fluid Mech.* **721**, 514–540.

YALNIZ, G., HOF, B. & BUDANUR, N.B. 2021 Coarse graining the state space of a turbulent flow using periodic orbits. *Phys. Rev. Lett.* **126**, 244502.

YANG, J. & LAKOBA, T.I. 2007 Universally-convergent squared-operator iteration methods for solitary waves in general nonlinear wave equations. *Stud. Appl. Maths* **118**, 153–197.

YIH, C.-S. 1965 *Dynamics of Nonhomogeneous Fluids*. Macmillan.

https://doi.org/10.1017/jfm.2026.11722 Published online by Cambridge University Press