




Shape optimisation to enhance flow-induced vibration of a cylinder using Bayesian optimisation

Stephen Joel Terrington¹ , Mark C. Thompson¹  and Kerry Hourigan¹ 

¹Fluids Laboratory for Aeronautical and Industrial Research (FLAIR), Department of Mechanical and Aerospace Engineering, Monash University, Melbourne, VIC 3800, Australia

Corresponding author: Stephen Joel Terrington, stephen.terrington1@monash.edu

(Received 22 April 2025; revised 26 September 2025; accepted 24 October 2025)

Bayesian optimisation with Gaussian process regression was performed to optimise the shape of an elastically mounted cylinder undergoing transverse flow-induced vibration. The vibration amplitude and mean power coefficient were obtained from two-dimensional numerical simulations, with Reynolds number $Re = 100$. First, shape optimisation was performed to maximise the amplitude of undamped vibrations. The optimised shape was found to be a thin crescent cylinder aligned perpendicular to the oncoming flow. The optimised shapes exhibited simultaneous vortex-induced vibration and galloping, a response which was not observed for other cylinder geometries at the same Reynolds number. Shape optimisation was also performed to maximise the power coefficient, where the power generation device was modelled as a linear damper. The power-optimised cylinders were also thin crescents, but with greater curvature compared with the amplitude-optimised cylinders. Compared with a circular cylinder, improvements in the power coefficient and efficiency of up to 523 % and 152 %, respectively, were obtained.

Key words: flow-structure interactions, vortex shedding

1. Introduction

Flow-induced vibration (FIV) is a fluid–structure interaction where the flow of a fluid past a solid body induces a vibration response in the solid body. Flow-induced vibration often results in harmful structural damage, and is therefore minimised in many engineering applications. However, FIV may be used to extract clean energy from environmental fluid flows (Bernitsas *et al.* 2008; Wang, Fan & Lin 2020a; Sun *et al.* 2024), and in such applications one may wish to enhance FIVs to improve the power output and efficiency.

© The Author(s), 2025. 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.

The present study explores the use of Bayesian optimisation with Gaussian process regression (BO-GPR) to enhance the vibration amplitude and power output from FIV.

The use of FIV for energy harvesting was proposed by Bernitsas *et al.* (2008), who introduced the VIVACE (vortex-induced vibration for aquatic clean energy) converter design. The original VIVACE design uses a 1 degree of freedom (1-DOF) elastically mounted circular cylinder, attached to an electric generator. The generator extracts energy from the cylinder, and therefore applies mechanical damping to the system. The FIV response of a 1-DOF elastically mounted cylinder has been extensively studied (e.g. Sarpkaya 2004; Wang *et al.* 2020b). When the fluid's free-stream velocity is such that the natural vortex shedding frequency of the cylinder matches the natural vibration frequency of the cylinder, the fluctuating force due to vortex shedding results in a significant vibration response, known as vortex-induced vibration (VIV). In this regime, the VIVACE device was found to be a cost effective method for generating clean energy (Bernitsas *et al.* 2008).

One possible improvement to the VIVACE device is to alter the shape of the vibrating body to achieve a larger vibration response. Cylinders of various alternative shapes have been investigated, including rectangular (Zhang *et al.* 2018a,b), triangular (Sirohi & Mahadik 2011; Zhang *et al.* 2016), diamond (Lian *et al.* 2017), D-section (Chen & Li 2024) and elliptical (Lo *et al.* 2024) cylinders. In addition to VIV at lower flow speeds, asymmetrical shapes may exhibit a galloping response at high speeds, which results in increased power generation (Wang *et al.* 2020; Lv *et al.* 2021). Galloping arises from a movement-induced instability due to the varying incident angle relative to the oncoming flow. One promising shape for power generation is the thin elliptical cylinder (Lo *et al.* 2023), which exhibits a hyper-branch featuring a combination of VIV and a galloping-like movement-induced instability. This allows the thin elliptical cylinder to sustain large vibration amplitudes with significant damping, and at comparatively lower velocities compared with traditional galloping (Lo *et al.* 2023). The thin elliptical cylinder with aspect ratio 5 was found to achieve approximately one order of magnitude improvement in the mean power coefficient compared with the circular cylinder (Lo *et al.* 2024).

Several previous studies have compared the power generation for different cylinder cross-sections, including Yang, Zhao & Tang (2013), Ding *et al.* (2015), Lian *et al.* (2017), Zhang *et al.* (2017), Sun, Jo & Seok (2019a,b), Tamimi *et al.* (2019) and Huang *et al.* (2024). However, these studies typically select a small number of different cross-sections, with limited or no parametric variation of the different shapes. Table 1 reports the best-performing cross-section determined by several of these studies. Direct comparisons between different studies are difficult due to differences in conditions between the studies. In general, however, shapes with sharp corners at the upper and lower edges and with a convex afterbody appear to give the best performance. Huang *et al.* (2024) compared the performance of 36 different bluff body shapes, formed from combinations of 6 different shapes for the forebody and afterbody. They found that, within the VIV regimes, shapes with convex circular or convex triangular forebodies achieved the best performance, while in the galloping regime, shapes with convex triangular or convex circular afterbodies gave the best performance. The overall best-performing cross-section used a rectangular forebody and a triangular afterbody.

Chen, Li & Yang (2024) performed shape optimisation to enhance the FIV response of a circular cylinder, using an unsteady adjoint-based method. They considered the subcritical regime, and found that shape optimisation reduces the critical Reynolds number from $Re_{cr} = 44$ to $Re_{cr} = 33$, reduces the minimum Re for energy harvesting from 20 to 14, and increases the power extraction by 420 % compared with a circular cylinder. Alhussein, Dalaq & Daqaq (2024) performed shape optimisation of a curved trapezoidal cylinder, using an artificial neural network as a surrogate model, and found that a modified trapezoid







	Shape		Method	Fluid	\bar{C}_p	Re
Yang <i>et al.</i> (2013)	Square		Experiment	Air	4.46×10^{-3}	2.2×10^4
Ding <i>et al.</i> (2015)	Quasi-trapezoid		Simulation	Water	0.245	1.2×10^5
Zhang <i>et al.</i> (2017)	Circle-triangle		Simulation	Water	0.265	1.3×10^4
Sun <i>et al.</i> (2019b)	Semi-circle		Experiment	Water	7.81×10^{-3}	4.0×10^4
Tamimi <i>et al.</i> (2019)	Triangular		Experiment	Water	0.234	3.8×10^4
Huang <i>et al.</i> (2024)	Rectangle-triangle		Experiment	Air	—	1.2×10^5

Table 1. Best-performing cylinder cross-sections determined from previous studies. A sketch of each shape is provided, with flow direction from left to right.

with curved upper and lower edges yields an increased power coefficient compared with a standard trapezoid.

While these studies have successfully implemented shape optimisation, they are limited in how they constrain the geometry. Chen *et al.* (2024) fixed the width of the cylinder and constrained the cross-sectional area to not exceed that of the circular cylinder. As a result, their analysis was restricted to aspect ratios below 1, and they could not represent thin cylinders aligned perpendicular to the oncoming flow. Likewise, Alhussein *et al.* (2024) restricted their geometry to an aspect ratio of 1, in addition to requiring flat front and back faces, so were also unable to represent thin cylinders. The present study finds that such thin cylinders result in a maximised power extraction from VIV. The geometry parametrisations adopted in previous studies cannot represent such cylinders.

The present study aims to determine the shape that optimises either the vibration amplitude or power coefficient for a 1-DOF elastically mounted cylinder. We consider a parametrised cylinder, that can represent a wide range of shapes including circular, elliptical, parabolic, triangular, rectangular, diamond and D-section cylinders. The vibration response of the parametrised shape is then determined using two-dimensional numerical simulations, performed in OpenFOAM. Finally, BO-GPR is used to tune the shape parameters to achieve an optimised FIV response.

Bayesian optimisation is a powerful framework for optimising expensive black-box functions – functions which are expensive to compute, and which do not have a known analytical expression (Diessner *et al.* 2022). Computational fluid dynamics simulations are an example of such expensive black-box functions, and the BO-GPR framework has proven useful for performing optimisations of computational fluid dynamics (CFD) simulations (Diessner *et al.* 2022; Morita *et al.* 2022). For the present study, the amplitude or power




Figure 1. Sketch of the problem set-up. A cylinder with height D and span W is immersed in a fluid flow with free-stream velocity U . The cylinder has a single degree of freedom, the displacement y in the transverse direction, and is attached to a spring–damper system with mass m , stiffness k and damping coefficient c . Finally, ρ , μ and ν are the density, dynamic viscosity and kinematic viscosity of the fluid, respectively.

coefficient obtained from CFD simulations is an expensive black-box function, which the BO-GPR methodology is well suited to optimise.

The structure of this article is as follows. First, in § 2 we briefly discuss the background theory for FIV and describe the numerical scheme used to predict the FIV response of the cylinder. Then, in § 3 we discuss the BO-GPR technique used to optimise the FIV response. Next, § 4 presents an optimisation study to maximise the amplitude of undamped vibrations. The combined VIV–galloping response of the optimised shape is then presented in § 5. Finally, an optimisation to maximise the power coefficient is investigated in § 6. Concluding remarks are made in § 7.

2. Flow-induced vibration

This section discusses the 1-DOF FIV of an elastically mounted cylinder, and the numerical method used to predict the vibration of the cylinder. The structure is as follows. First, in § 2.1, we discuss the problem set-up and non-dimensional equations of motion. Then, in § 2.2, we describe the geometry parametrisation used in this study. Next, in § 2.3, we discuss the numerical method used to predict the cylinder’s vibration response. Finally, in § 2.4 we validate the numerical methodology.

2.1. Problem set-up

The problem under consideration is as shown in figure 1. A cylinder with height D and span W is mounted to a spring–mass–damper system with total mass m , spring constant k and damping constant c . The cylinder is exposed to a fluid flow with free-stream velocity U , and the density, dynamic viscosity and kinematic viscosity of the fluid are ρ , μ and ν , respectively.

Non-dimensionalising by D , U and ρ , the equations of motion for the fluid are the Navier–Stokes equations

$$\frac{d\mathbf{u}}{dt} + \mathbf{u} \cdot \nabla \mathbf{u} = -\nabla p + \frac{1}{Re} \nabla^2 \mathbf{u}, \quad (2.1a)$$

$$\nabla \cdot \mathbf{u} = 0, \quad (2.1b)$$

where \mathbf{u} and p are the non-dimensional fluid velocity and pressure, respectively, while $Re = UD/\nu$ is the Reynolds number. The boundary conditions on the cylinder's surface are $\mathbf{u} = \dot{\mathbf{y}}$, where y is the dimensionless cylinder displacement, and the overdot denotes the non-dimensional time derivative.

The equation of motion for the cylinder displacement, in dimensionless form, is

$$\ddot{y} + \frac{4\pi\xi}{U^*}\dot{y} + \frac{4\pi^2}{U^{*2}}y = \frac{C_y}{m^*}, \quad (2.2)$$

where C_y is the transverse force coefficient, m^* is the dimensionless mass, U^* is the reduced velocity and ξ is the damping ratio. These dimensionless parameters are defined as

$$C_y = \frac{F_y}{\frac{1}{2}\rho U^2 DW}, \quad (2.3)$$

$$U^* = \frac{2\pi U}{D\sqrt{k/m}}, \quad (2.4)$$

$$m^* = \frac{m}{\frac{1}{2}\rho D^2 W}, \quad (2.5)$$

$$\xi = \frac{c}{2\sqrt{km}}, \quad (2.6)$$

where F_y is the transverse force applied to the body.

We define the reduced velocity with respect to the undamped natural frequency of the system in a vacuum, $U^* = U/(f_n D)$, where $f_n = 1/(2\pi)\sqrt{k/m}$. Note that a range of other frequencies are often used to define the reduced velocity, such as the natural frequency in still water, or the natural frequency including the added mass effect (Sarpkaya 2004). We define the reduced velocity as (2.4), since this form of the reduced velocity is determined completely by the system parameters alone.

For the dimensionless mass, previous studies typically consider the mass ratio, which is the ratio between the total oscillating mass and the mass of displaced fluid

$$m_r = \frac{m}{\rho W A_c}, \quad (2.7)$$

where A_c is the cross-sectional area of the cylinder. The dimensionless mass and mass ratio are related as $m^*/m_r = 2A_c/D^2$. Note that for a fixed total mass, m^* is independent of the cylinder's shape, while m_r depends on the cylinder's cross-sectional area. In the present study, optimisations are performed with either a fixed m^* , corresponding to a constant total mass, or for a fixed m_r , corresponding to a fixed density of the solid body.

Numerically solving the Navier–Stokes (2.1), coupled with the equation of motion for the cylinder (2.2), yields a predicted vibration response $y(t)$ and both velocity \mathbf{u} and pressure p fields, for a given cylinder shape and combination of the dimensionless parameters m^* , U^* , ξ and Re . The amplitude of vibration is characterised by either the maximum amplitude (y_{max}), or root mean square (r.m.s.) amplitude (y_{rms}).

For power generation applications, the attached generator must extract mechanical energy from the system. We model the energy extraction as a linear damper, with damping coefficient c . Assuming perfect efficiency of the generator, and no other sources of damping, the mean power coefficient is

$$\bar{C}_P = \frac{\bar{P}}{\frac{1}{2}\rho U^3 DW} = \frac{c}{\frac{1}{2}\rho U DW} \bar{\dot{y}}^2 = \frac{4\pi m^* \xi}{U^*} \bar{\dot{y}}^2, \quad (2.8)$$

where P is the dimensional power output, and overbars represent time-averaged quantities. In practice, there may be significant energy losses due to other sources of damping (e.g. friction in the bearing system), as well as losses in the transducer used to convert mechanical energy into electrical power. As a result, the power coefficients and efficiencies reported in this paper refer only to the performance of the cylinder (i.e. the transfer of energy from the fluid to the cylinder), not the performance of a realistic energy conversion system. Additionally, real power conversion devices may not impart a linear damping on the system.

Note that \bar{C}_P is non-dimensionalised by the projected area of the cylinder perpendicular to the flow direction. However, it is typical in power generation applications to non-dimensionalise the power extraction by the total swept area of the device. This represents an efficiency by which power is extracted from the total power available in the fluid. The total swept area for a single cylinder is $WD(1 + 2y_{\max})$, and therefore, the efficiency is

$$\eta = \frac{\bar{P}}{\frac{1}{2}\rho U^3 WD(1 + 2y_{\max})} = \frac{\bar{C}_P}{1 + 2y_{\max}}. \quad (2.9)$$

As discussed by Lo *et al.* (2024), increasing the maximum amplitude of the cylinder will generally increase the power coefficient. However, this may not lead to an increase in the efficiency, since the swept area also increases. The present study focuses on optimising the power coefficient (\bar{C}_P) of a single cylinder, representing the maximum power extraction from a single device. In practical applications, one may wish to extract the most energy from a given environmental fluid flow, in which case the efficiency (η) should be optimised instead. In such cases, one should consider the efficiency of an array containing multiple FIV devices, in either a tandem or parallel configuration, which is beyond the scope of the present study. Also note that the cost of a single VIVACE device is not likely to depend substantially on the vibration amplitude. Therefore, for a sufficiently abundant source of environmental energy, it may be more cost effective to maximise power coefficient, rather than the efficiency.

2.2. Parametrising the geometry

The Bayesian optimisation technique used in this work requires the geometry to be parametrised. In this section we introduce the various parameters that control the shape of the ellipse. These parameters are chosen to allow various elementary shapes, such as circles, ellipses, squares, D-section cylinders, diamonds, triangles to be represented.

We first consider the following generalisation of the ellipse (superellipse):

$$|x/a|^{2/m} + |y/b|^{2/n} = 1, \quad (2.10)$$

which can be defined parametrically as

$$x = a \operatorname{sign}(\sin(\theta)) |\sin(\theta)|^m, \quad y = b \operatorname{sign}(\cos(\theta)) |\cos(\theta)|^n, \quad 0 \leq \theta \leq 2\pi. \quad (2.11)$$

For $n = m = 1$, (2.10) describes an ellipse with width $2a$ and height $2b$. For the present study, the cylinder has unit height in dimensionless coordinates ($b = 0.5$), while $a = 1/(2AR)$, where AR is the cylinder aspect ratio. Varying the parameters n and m varies the shape of the cylinder, while maintaining a constant width and height. As shown in

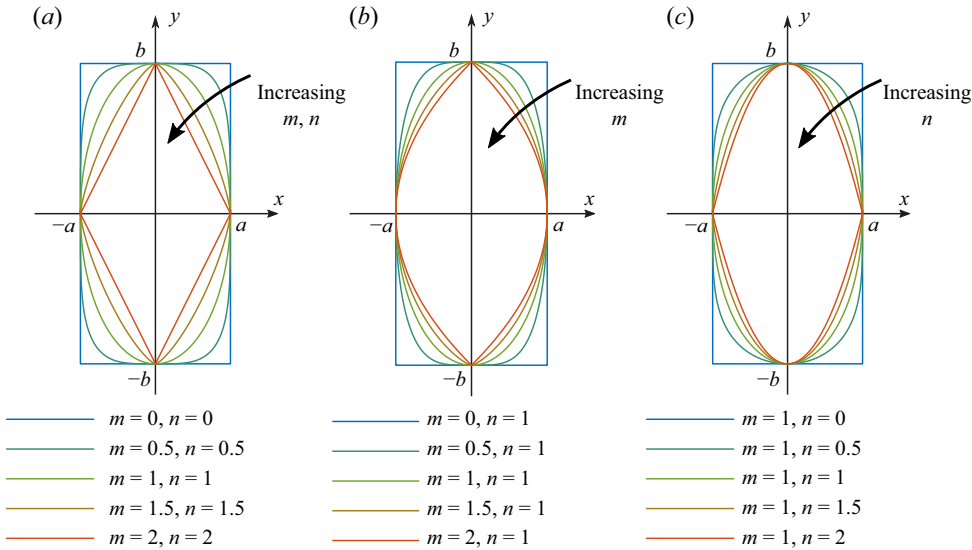


Figure 2. Effect of the parameters m and n on the shape of the generalised superellipse. (a) Effect of varying both m and n , with $m = n$; (b) effect of varying m , with $n = 1$; and (c) effect of varying n , with $m = 1$.

figure 2, the superellipse can represent a diamond ($n = m = 2$), a rectangle ($n = m = 0$), or a shape with parabolic edges ($m = 2$ and $n = 1$ or $m = 2$ and $n = 1$).

Previous studies have identified that the forebody and afterbody play different roles in generating FIV motion (Zhao, Hourigan & Thompson 2018; Chen & Li 2023; Gupta *et al.* 2025). Therefore, we allow the shape parameters to differ for the front and rear surfaces. First, we restrict $b = 0.5$ for both the forebody and afterbody. Then, we define a_f , m_f and n_f as the parameters for the forebody ($x < 0$), and a_a , m_a and n_a as the parameters for the afterbody. We note that a_a and a_f are constrained by the cylinder aspect ratio, $a_f + a_a = 1/AR$. We also define a parameter α , which controls the relative width of the forebody and afterbody

$$(1 - \alpha)a_f = (1 + \alpha)a_a, \quad -1 \leq \alpha \leq 1. \quad (2.12)$$

As shown in figure 3, using either $\alpha = -1$ or $\alpha = 1$ results in a shape with no forebody or no afterbody, respectively, including D-section ($n = m = 1$) and triangular ($n = m = 2$) cylinders.

Finally, we introduce a parameter to control the camber of the cylinder. We define $x_{1,f}(y)$ and $x_{1,a}(y)$ as the forebody and afterbody of the uncambered cylinder, which are computed using the six parameters AR , α , n_f , m_f , n_a , m_a as discussed above. We then define the following cambered curve:

$$x_c = \beta(1 - 2y_c)^2, \quad -0.5 \leq y_c \leq 0.5, \quad (2.13)$$

where β represents the amount of camber. We then compute a normalised arc-length parametrisation of this curve, $x_c(s)$, $y_c(s)$, with $-0.5 \leq s \leq 0.5$. Finally, the cylinder shape is given by

$$x = x_c(s) + n_x(s)x_1(s), \quad y = y_c(s) + n_y(s)x_1(s), \quad (2.14)$$

where n_x and n_y are components of a unit vector. For thin cylinders ($AR > 3$), this unit vector is chosen to be the unit normal to the camber line, so that the uncambered cylinder represents a thickness distribution about the camber line. To avoid self-intersections,




Figure 3. Effect of the parameter α on the cylinder shape. For (a,c) $\alpha = -1$, the cylinder has only an afterbody, while for $\alpha = 1$, the cylinder has only a forebody. This allows the cylinder to be either a D-section (a,b) or triangular cylinder (c,d).




Figure 4. Addition of camber to the cylinder. (a) shows the uncambered cylinder (x_1, y_1) , while (b) shows the camber curve (x_c, y_c) and the cambered cylinder (x, y) .

we use $n_x = 1$ and $n_y = 0$ for thick cylinders ($AR \leq 1$). For $1 \leq AR \leq 3$, the angle of the unit vector is linearly interpolated between these values.

Figure 4 demonstrates the application of camber to a cylinder. Panel (a) shows the uncambered cylinder, with parameters $AR = 5$, $\alpha = 0$, $r_{m,f} = 0.5$, $r_{n,f} = 0.5$, $r_{m,a} = 2$ and $r_{n,a} = 1$. Panel (b) shows both the camber curve (x_c, y_c) , and the cambered cylinder (x, y) .

2.3. OpenFOAM simulations

Numerical simulations are performed using the open-source software OpenFOAM 10.0 (Weller *et al.* 1998), with the PIMPLE algorithm (a combination of the Pressure-Implicit with Splitting of Operators (PISO) and Semi-Implicit Method for Pressure Linked

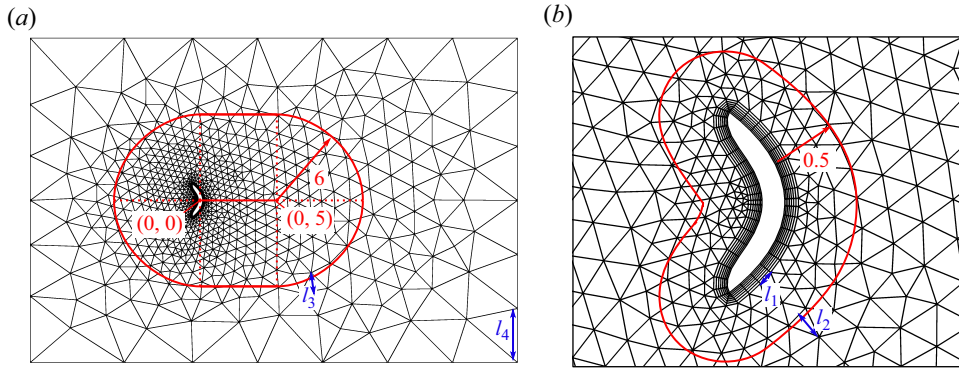


Figure 5. Sketch illustrating the unstructured mesh used in this study (not to scale). The parameters l_1 , l_2 , l_3 and l_4 represent approximate cell sizes (side lengths of the triangular cells) at various distances from the cylinder.

Equations (SIMPLE) algorithms) used for pressure–velocity coupling. The equations of motion are discretised using the finite-volume scheme. The linear interpolation scheme (equivalent to second-order centred differencing) was used for all spatial derivatives, aside from the advection term, for which a second-order linear upwind scheme was used. The first-order backwards Euler scheme was used for temporal derivatives. An adaptive time-step scheme was used, with a maximum Courant number of $Co_{max} = 0.9$.

A dynamic mesh, using the `sixDoFRigidBodyMotion` utility in `openFOAM`, was used to account for the motion of the cylinder. First, the viscous and pressure forces applied to the cylinder are computed from the finite-volume discretisations of the velocity and pressure fields. Then, the position of the cylinder is updated using (2.2). Grid points within 30 diameters of the cylinder move rigidly with the cylinder, while the position of the remaining grid points are updated using a Laplacian mesh motion solver.

We consider a two-dimensional square computational domain, with boundaries located at a distance L_B from the cylinder’s rest position. Specifically, there is an inlet located at $x = -L_B$, an outlet at $x = L_B$ and upper and lower free-slip walls at $y = \pm L_B$. Boundary conditions are $\mathbf{u} = (0, \dot{y}, 0)$ at the cylinder’s surface, $\mathbf{u} = (1, 0, 0)$ at the inlet, $p = 0$ at the outlet and $\partial u_x / \partial y = 0$ on the free-slip walls.

To achieve reasonable mesh quality for an arbitrary cylinder shape, unstructured meshes are generated programmatically using the open-source software GMSH (Geuzaine, Remacle & Dular 2009). An example mesh is shown in figure 5. Increased resolution is used near the body and in the wake region behind the body. The algorithm used by GMSH attempts to produce elements with a pre-defined target cell size at each location in the computational domain. These target cell sizes are illustrated in figure 5: l_1 is the element size near the solid body; l_2 is the target element size at a distance 0.5 from the body; l_3 is the target element size in the wake (at a distance 6 from a line segment passing through the points (0, 0) and (0, 5)); and l_4 is the target element size in the far-field. Linear interpolation is used to determine the target element size at an arbitrary location in the domain.

An inflation layer containing high aspect ratio quadrilateral elements is used near the boundary of the body, while triangular elements are used for the remaining region. The thickness of the first layer is $0.1l_1$, and the growth rate is 1.1. Curvature-based mesh adaptation is used with n_c boundary elements per radian of curvature. A boundary-layer fan is employed at sharp corners (with angle greater than 0.2 radians). The number of boundary-layer fan elements per radian of angle is set to n_f .




Figure 6. Time histories of cylinder displacement for a variety of different cylinders: (a) a circular cylinder at $U^* = 6$ and $\xi = 0$, (b) the OpU10 cylinder (see figure 9) at $U^* = 11$ and $\xi = 0$, (c) the semi-elliptical cylinder SE6 (see table 7) at $U^* = 6$ and $\xi = 0.0293$ and (d) the triangular cylinder T7 (see table 7) with $U^* = 7$ and $\xi = 0.193$. For all cases, $Re = 100$ and $m^* = 4.7124$.

To ensure accurate determination of the objective function, it is important that the flow reaches a statistically steady state. Figure 6 plots time histories of cylinder displacement for several different cylinder shapes. While the time taken to reach a steady state differs for each cylinder, in all cases a steady state is approximately reached after $t = 200$, and this steady state is close to periodic. Temporally averaged statistics (mean power coefficient, r.m.s. displacements) are obtained over a time interval spanning the first local maximum of cylinder displacement after $t = 200$, to the last local maximum before $t = 500$. For the maximum cylinder displacement, the average value of each local maximum between $t = 200$ and $t = 500$ is reported.

To reduce the startup time, the cylinder is given an initial velocity of $\dot{y} = 2$ at $t = 0$. In some cases, this may trigger ‘hard’-oscillating behaviour (Novak & Tanaka 1974), in which the cylinder only oscillates given a sufficiently large initial displacement. For cases where the oncoming flow is relatively steady, it is reasonable to assume that a power generating device could be manually started in this manner. For devices which experience substantial gust loads, the transient behaviour must be considered (Alhussein *et al.* 2024), which could result in a different optimal shape to those obtained in the present study.

The present study considers two-dimensional laminar flow at low Reynolds numbers ($Re \approx 100$), to keep the computation cost per simulation low. The assumption of two-dimensionality restricts the results to wide cylinders ($W/D \gg 1$). For future research,

	l_1	l_2	l_3	l_4	n_c	n_f	N_c
Mesh 1	8×10^{-2}	2×10^{-1}	8×10^{-1}	2×10^1	5	2	5000
Mesh 2	4×10^{-2}	1×10^{-1}	4×10^{-1}	1×10^1	10	4	20 000
Mesh 3	2×10^{-2}	5×10^{-2}	2×10^{-1}	5×10^0	20	7	80 000

Table 2. Parameters for the meshes used in the grid resolution study. The parameters l_1 , l_2 , l_3 and l_4 are the target cells sizes shown in figure 5, while n_c and n_f are the minimum number of cells per radian on the curved boundary and at sharp corners, respectively. The approximate number of cells, N_c , is also shown.

		y_{max}		$\overline{\dot{y}^2}$		St_y	
$U^* = 6$	Mesh 1	2.16	(3.62 %)	1.55	(3.96 %)	0.162	(0.50 %)
	Mesh 2	2.24	(0.11 %)	1.61	(0.17 %)	0.162	(0.06 %)
	Mesh 3	2.24		1.61		0.162	
$U^* = 10$	Mesh 1	4.88	(4.08 %)	2.15	(3.81 %)	0.0991	(0.59 %)
	Mesh 2	4.73	(0.91 %)	2.08	(0.91 %)	0.0997	(0.05 %)
	Mesh 3	4.68		2.07		0.0997	

Table 3. Mesh resolution study comparing the oscillation amplitude (y_{max}), the r.m.s. velocity of the cylinder ($\overline{\dot{y}^2}$) and the Strouhal number based on the cylinder displacement St_y , obtained on the three different meshes listed in table 2, at various reduced velocities (U^*). Simulations are performed for a preliminary optimised cylinder at $Re = 100$ and $m^* = 4.7124$. Parentheses indicate the relative error compared with mesh 3.

it would be interesting to perform the optimisation at larger Reynolds numbers (i.e. $Re \approx 10^4$), using an appropriate turbulence model, and validate the results against experimental measurements. Additionally, other effects relevant to practical applications, such as the effects of gust loads or the angle of oncoming flow, should be considered in the future.

2.4. Validation of numerical scheme

The mesh resolution and domain size are validated using an optimised cylinder obtained in the preliminary stages of this project. Table 2 provides the target element sizes for three meshes with increasing resolution, while table 3 reports the amplitude obtained using each cylinder at $U^* = 6$ and $U^* = 10$, as well as the r.m.s. velocity of the cylinder $\overline{\dot{y}^2}$ and Strouhal number based on cylinder displacement St_y . The differences between meshes 2 and 3 are below 1 %, and therefore mesh 2 is satisfactory for the present study.

Simulations were performed using the Setonix supercomputer located at the Pawsey Supercomputing Research Centre in Perth, Western Australia. Simulations using mesh 2 require approximately 33 hours of wall clock time, using 8 CPU cores in parallel. Due to the small size of this problem, further increasing the number of CPU cores does not substantially improve the time. Simulations using mesh 1 typically require 6 hours of wall time on a single CPU core. To speed up the optimisation process, an initial optimisation is performed using mesh 1 to efficiently explore the entire design space. Then, a second optimisation is performed over a reduced design space, using mesh 2, to obtain a more accurate solution. Further details are provided in § 4.

Table 4 reports the variation of oscillation amplitude (y_{max}), r.m.s. velocity of the cylinder $\overline{\dot{y}^2}$ and Strouhal number based on cylinder displacement St_y obtained using various domain sizes (L_B). The difference between $L_B = 200$ and $L_B = 400$ are below

	L_B	y_{max}		$\overline{\dot{y}^2}$		St_y	
$U^* = 6$	50	2.46	(11.43 %)	1.78	(11.88 %)	0.163	(0.48 %)
	100	2.33	(5.79 %)	1.68	(6.06 %)	0.162	(0.27 %)
	200	2.24	(1.73 %)	1.61	(1.83 %)	0.162	(0.10 %)
	400	2.2		1.58		0.162	
$U^* = 10$	50	5.03	(6.34 %)	2.21	(6.08 %)	0.0992	(−0.52 %)
	100	4.79	(1.53 %)	2.11	(1.51 %)	0.0997	(0.04 %)
	200	4.73	(0.17 %)	2.08	(0.18 %)	0.0997	(0.03 %)
	400	4.72		2.08		0.0997	

Table 4. Domain size study comparing the oscillation amplitude (y_{max}), the r.m.s. velocity of the cylinder ($\overline{\dot{y}^2}$) and the Strouhal number based on the cylinder displacement St_y obtained for different domain sizes L_B using the mesh 2 resolution, reduced velocities ($U^* = 6$ and $U^* = 10$). Simulations are performed for a preliminary optimised cylinder at $Re = 100$ and $m^* = 4.7124$. Parentheses indicate the relative error compared with $L_B = 400$.




Figure 7. Comparison between the present numerical simulations and previous studies of Leontini *et al.* (2006) and Rajamuni *et al.* (2020), showing the oscillation amplitude y_{max} against U^* , for a circular cylinder with $m_r = 10$, $\xi = 0.01$ and $Re = 200$.

2 %, and therefore $L_B = 200$ is considered satisfactory for this project. Note that $L_B = 50$ was found to be satisfactory for circular cylinders, which oscillate at a much lower amplitude. The large domain size needed to achieve a domain-independent solution suggests that the effective blockage ratio is related to the amplitude of oscillations, rather than the cylinder diameter.

Figure 7 presents a validation study, comparing the present numerical scheme with previous studies of Leontini, Thompson & Hourigan (2006) and Rajamuni, Thompson & Hourigan (2020), for the VIV of a circular cylinder with $m_r = 10$, $\xi = 0.01$ and $Re = 200$. Reasonable agreement between the present simulations and the prior studies is observed, validating the numerical scheme.

3. Bayesian optimisation with Gaussian process regression

This section discusses the BO-GPR method used to perform the shape optimisation. The BO-GPR method is effective for optimisations involving ‘expensive black-box’ functions – i.e. functions that are expensive to evaluate, and whose mathematical expressions are unknown (Diessner *et al.* 2022). Optimisation problems in fluid mechanics are often expensive black-box functions (Diessner *et al.* 2022; Morita *et al.* 2022), as either numerical simulations or experiments are performed to evaluate the function. In the present study, either the amplitude (y_{max}) or power coefficient (\bar{C}_p) is the output of a function over the shape parameters (AR , α , β , m_f , n_f , m_a , n_a) and the FIV parameters (m^* , U^* and ξ). To evaluate this function requires performing an unsteady CFD simulation, which requires either 6 h (coarse mesh) or 33 h (fine mesh) per function evaluation, as discussed previously. This certainly fits the definition of an expensive black-box function.

Detailed discussions of the application of BO-GPR to fluid mechanics problems can be found in Morita *et al.* (2022) and Diessner *et al.* (2022). Morita *et al.* (2022) consider three example optimisation problems: shape optimisation of a lid-driven cavity; shape optimisation of a diverging channel; and optimisation of a spoiler-ice airfoil model. They show that BO-GPR efficiently explores the parameter space, and converges towards the optimum solution after relatively few evaluations. Diessner *et al.* (2022) first consider several synthetic test functions with various challenging features – including oscillatory, multi-modal, steep edges and mostly flat test functions – to compare the performance of various BO-GPR methods. Then, they apply BO-GPR to optimise the control law for a boundary-layer control device to minimise drag while achieving a net energy saving. The BO-GPR method has been used to solve a wide range of fluid mechanics related optimisation problems (Talnikar *et al.* 2014; Mahfoze *et al.* 2019; Nabae & Fukagata 2021; Begall *et al.* 2023; Kumar *et al.* 2024; Mallor *et al.* 2024).

3.1. Overview of Bayesian optimisation with Gaussian process regression

We now give a brief overview of the BO-GPR method, based on the presentations given in Morita *et al.* (2022) and Diessner *et al.* (2022). We let \mathbf{q} be a vector of design parameters, and $r(\mathbf{q})$ be the objective function to be minimised. The first step in Bayesian optimisation (BO) is, given a set of n measurement data $D_n = (\mathbf{q}_i, r_i)_{i=1}^n$, to estimate the posterior probability distribution for the measurement at the next sample point \mathbf{q}_{n+1} . The value of \mathbf{q}_{n+1} that maximises some acquisition function is then selected and the response value $r_{n+1} = r(\mathbf{q}_{n+1})$ is obtained. This process is then repeated until an optimum solution is found.

In BO-GPR, the objective function is approximated as a Gaussian process (GP). A GP is a stochastic process in which every finite collection of random points is represented by a multivariate normal distribution. The GP is defined by a mean function $m(\mathbf{q})$, and a covariance function (kernel) $c(\mathbf{q}, \mathbf{q}')$. For the set of sample points $\mathbf{q}_{1,2,\dots,n}$ (where \mathbf{q}_i represents the i th measurement), we have a vector of measurement points \mathbf{r} , a vector of mean values \mathbf{m} and a covariance tensor \mathbf{C} , with components defined as $r_i = f(\mathbf{q}_i)$, $m_i = m(\mathbf{q}_i)$ and $C_{i,j} = c(\mathbf{q}_i, \mathbf{q}_j)$. The prior probability distribution for the measurement points (i.e. the assumed probability distribution prior to obtaining the measurement data) is a multivariate normal distribution

$$\mathbf{r} \sim N(\mathbf{m}, \mathbf{C}). \quad (3.1)$$

The posterior predictive distribution for the next sample point $r_{n+1} = f(\mathbf{q}_{n+1})$ is also a normal distribution

$$r_{n+1}|D_n = N(\mu_n(\mathbf{q}_{n+1}), \sigma_n(\mathbf{q}_{n+1})), \quad (3.2)$$

with the posterior mean and variance given by

$$\mu_n(\mathbf{q}_{n+1}) = \mathbf{m}(\mathbf{q}_{n+1}) + \mathbf{c}(\mathbf{q}_{n+1})^T (\mathbf{C} + \sigma^2 \mathbf{I})^{-1} (\mathbf{r} - \mathbf{m}), \quad (3.3)$$

$$\sigma_n^2(\mathbf{q}_{n+1}) = c(\mathbf{q}_{n+1}, \mathbf{q}_{n+1}) - \mathbf{c}(\mathbf{q}_{n+1})^T (\mathbf{C} + \sigma^2 \mathbf{I})^{-1} \mathbf{c}(\mathbf{q}_{n+1}), \quad (3.4)$$

where $\mathbf{c}(\mathbf{q}_{n+1})$ is a vector of covariances between the candidate point and all design points ($c_i = c(\mathbf{q}_{n+1}, \mathbf{q}_i)$). The result is a probability distribution for the objective function, sampled at the candidate point \mathbf{q}_{n+1} .

In this study, we use the zero mean function $\mathbf{m}(\mathbf{q}) = 0$ for the prior distribution. The sample data r_i are also normalised to have zero mean, and unit standard deviation. For the covariance function c , the Matérn 5/2 kernel is used, following Diessner *et al.* (2022),

$$c(\mathbf{q}, \mathbf{q}') = \sigma_f^2 \left(1 + \frac{\sqrt{5}}{\rho} + \frac{5d^2}{4\rho^2} \right) \exp \left(-\frac{\sqrt{5}d}{\rho} \right). \quad (3.5)$$

Here, $d = \sqrt{(\mathbf{q} - \mathbf{q}') \cdot (\mathbf{q} - \mathbf{q}')}$ is the distance between the two sample points, σ_f^2 is a scaling factor and ρ is the length scale. Both σ_f^2 and ρ are hyperparameters that are estimated from the sample data using maximum likelihood estimation. Note that (3.5) has a single length scale for all parameters. To account for this, we normalise each of our design parameters to lie in the range $0 \leq q_i \leq 1$.

Given the posterior probability function for the surrogate objective function, an acquisition function is then applied to select the next sample point(s). In this study, we use the expected improvement (EI) acquisition function, which is one of the most commonly used acquisition functions (Morita *et al.* 2022). Given the best point found so far, $r^* = f(\mathbf{q}^*)$, the EI is defined as

$$\alpha_{EI} = E(\max(f(\mathbf{q}^*) - f(\mathbf{q}))), \quad (3.6)$$

where $E()$ represents the expected value. The value of \mathbf{q} that gives the maximum value of α_{EI} – i.e. the design point that gives the greatest EI – is selected as the next design point.

In this study, batch optimisation with local penalisation (Gonzalez *et al.* 2016) is used. The acquisition function is initially defined as in (3.6). The first batch member, \mathbf{q}_{n+1} , is selected based on the maximum of this acquisition. Then, the acquisition function is penalised in the neighbourhood of this point, and the next batch element selected based on the penalised acquisition function. This is repeated until all points in the batch are selected. Numerical simulations for each batch are then performed simultaneously.

The BO-GPR method was implemented using the Python libraries GPy (<http://github.com/SheffieldML/GPy>) and GPyOpt (<http://github.com/SheffieldML/GPyOpt>). The BO-GPR method was coupled non-intrusively to the OpenFOAM simulations using Bash scripts. The overall process is as follows: first, the next batch of design points are obtained using the BO-GPR method implemented in GPyOPT. Then, meshes are procedurally generated using the Python GMSH package (Geuzaine *et al.* 2009). Next, simulations are performed using OpenFOAM. Finally, the results of the OpenFOAM simulations are post-processed in Python to compute the objective function. These results are then supplied to GPyOpt to select the next batch of design points.

4. Optimising amplitude of undamped vibrations

In this section, the BO-GPR method is applied to maximise the oscillation amplitude of an undamped ($\xi = 0$) cylinder. Unless otherwise stated, optimisations are performed for fixed values of U^* , Re and m^* . The Reynolds number is set to $Re = 100$ to keep the computational cost per simulation low. The dimensionless mass is set to $m^* = 4.7124$,

which is the same as the thin elliptical cylinder ($AR = 5$, $m_r = 15$) studied by Lo *et al.* (2024), which is known to produce substantial oscillations. Finally, optimisations are performed for a range of different U^* .

For this case, the objective function to be maximised is the r.m.s. amplitude ($r(\mathbf{q}) = y_{rms}$), while the design parameter vector \mathbf{q} is a vector of all shape parameters, $\mathbf{q} = (AR, \alpha, \beta, m_f, n_f, m_a, n_a)$. The r.m.s. amplitude is selected as the optimisation target, rather than the peak amplitude y_{max} , as it is less sensitive to isolated large peaks. The batch size is set to 10, with the first 10 design points randomly sampled using Latin hypercube sampling. The shape parameters are constrained to the following domains: $1 \leq AR \leq 20$, $-1 \leq \alpha, \beta \leq 1$ and $0.1 \leq m_f, n_f, m_a, n_a \leq 2$. The use of Latin hypercube sampling for the initial condition results in significant randomness in the number of evaluations required to optimise the objective function. Thus, a large number of optimisations would be required to tune various hyperparameters (such as batch size). Combined with the high computational cost of evaluating the objective function, this was not considered cost effective. The hyperparameters used in this study are sufficient to substantially improve the cylinder performance using reasonable computational effort, but are not likely to be the optimal parameters for fastest optimisation.

To speed up the training process, the optimisation is completed in two stages. First, an optimisation is performed over the entire design space using the coarse mesh (mesh 1 from table 2). Then, a second optimisation is performed with a high-resolution mesh (mesh 2) over a reduced domain, with a domain size determined using the 10 best-performing cylinders in the coarse-mesh optimisation, as described in detail below. Simulations on the coarse mesh typically require approximately 6 h per batch, while the fine mesh requires approximately 33 h per batch. The optimisation over the full domain typically requires approximately 50 batches, while less than 10 batches are needed on the reduced domain. Therefore, the two-stage approach allows an accurate optimum to be obtained in approximately 600 h, while 1600 h would be required if a single optimisation were performed using the fine mesh.

Figure 8 plots the maximum value of the objective function against number of batches, for $U^* = 6$ and $U^* = 10$, for both the initial optimisation with a coarse mesh, and the final optimisation with the fine mesh. One difficulty with the BO-GPR method is the determination of when to terminate the optimisation (Morita *et al.* 2022). For the coarse-mesh optimisations (panels (a) and (c)), the objective function increases regularly until 41 batches (panel (a)) and 24 batches (panel (c)), respectively, after which no further improvements are observed. The coarse-mesh optimisations are terminated once less than 1 % improvement was observed over 20 batches. For the fine-mesh optimisations (panels (b,c)) improvements of the objective function are observed over the first 5 batches, and optimisations are terminated once less than 1 % improvement is observed for 5 consecutive batches. In general, however, one cannot guarantee that a global optimum has been reached.

Figure 8 also plots the variation of the best-performing shape obtained at various stages of the optimisation. As the optimisation is performed, the optimised shapes for different U^* converge towards a similar design. The final optimised shape for each U^* is plotted in figure 9, and the shape parameters are provided in table 5. The label OpUX indicates the optimised shape for $U^* = X$. For all U^* , the optimised shape is a thin crescent cylinder, with sharp vertices. The aspect ratio of the optimised shape is generally close to 20, which is the maximum AR considered in this study. Hence, the thickness of the cylinder should be minimised to maximise the amplitude of FIV oscillations. The cylinders all have a slight positive camber. The extent of camber is characterised by the parameter $x_{cam} = \beta - \alpha / (2AR)$, which represents the maximum x -coordinate of the mean camber line. The camber of the optimised cylinders generally lies between $x_{cam} = 0.09$ and $x_{cam} = 0.11$ for

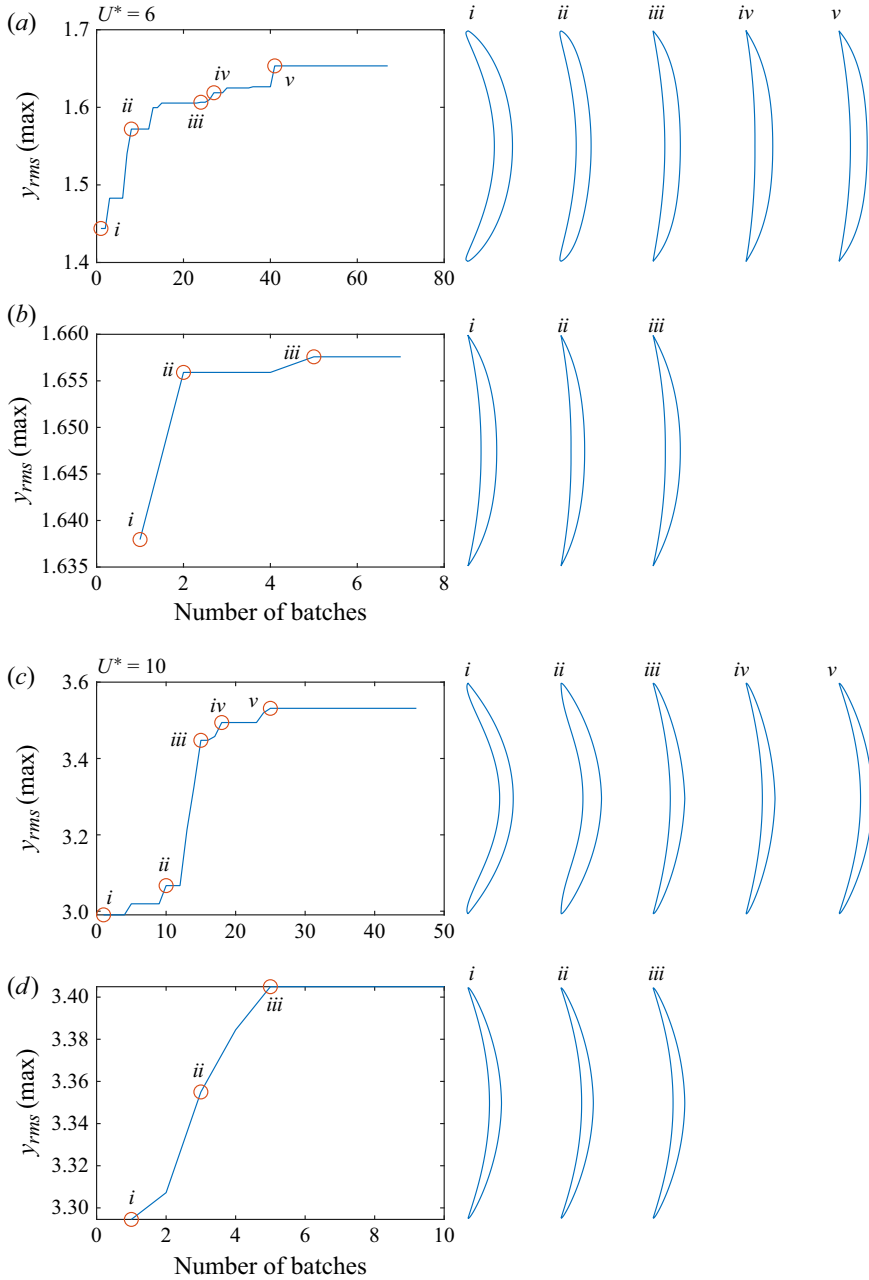


Figure 8. History of shape optimisation for (a,b) $U^* = 6$ and (c,d) $U^* = 10$ showing the best r.m.s. amplitude (y_{rms}) obtained against number of batches. The cylinder shape at various stages throughout the optimisation is also shown. At each U^* , an initial optimisation was performed using the coarse mesh (a,c), followed by a second optimisation using a fine mesh resolution (b,d). Optimisations were performed for a fixed $Re = 100$ and $m^* = 4.7124$.




Figure 9. Comparison of optimised shapes obtained for six different U^* . The label OpUX indicates the optimised shape obtained for $U^* = X$. Optimisations were performed for a fixed $Re = 100$ and $m^* = 4.7124$.

Cylinder label	U_{opt}^*	AR	α	β	m_f	n_f	m_a	n_a	x_{cam}
OpU5	5	19.13	-0.8716	0.06594	0.5763	1.802	1.894	1.655	0.089
OpU6	6	15.32	-0.2421	0.0773	1.257	1.877	0.7337	2	0.085
OpU7	7	19.85	0.3447	0.1021	1.549	1.705	1.333	0.903	0.093
OpU8	8	19.87	-0.8681	0.08598	1.762	0.4687	1.357	1.932	0.11
OpU9	9	19.65	0.09477	0.1015	1.949	1.725	0.7916	1.085	0.099
OpU10	10	19.66	-0.07547	0.1097	1.042	1.767	1.682	0.8993	0.11
OpU6_MR	6	19.73	0.7302	0.1147	1.947	1.963	1.414	1.435	0.096
OpU8_MR	8	20	0.8624	0.1554	1.966	2	2	0.4523	0.13
OpU10_MR	10	20	0.01124	0.1144	1.748	1.877	1.816	1.066	0.11
OpR60	10	18.61	-0.7136	0.1315	1.316	1.524	1.581	1.649	0.1507
OpR80	10	20	0.7669	0.1353	1.388	1.902	1.466	2	0.1161

Table 5. Shape parameters for various cylinders optimised for maximum amplitude y_{rms} , with fixed $U^* = U_{opt}^*$, $Re = 100$ and $m^* = 4.7124$. The _MR suffix indicates cylinders optimised for fixed $m_r = 15$ instead of a fixed m^* .

the optimised shapes. The general shape suggests that cylinders may behave like an airfoil at large angles of attack, with attached flow to minimise drag, and a component of lift force in line with the cylinder motion to enhance FIV. This proposal will be investigated in § 5.

Figures 10 and 11 plot two-dimensional slices of the objective function for $U^* = 5$ (figure 10) and $U^* = 10$ (figure 11). For these figures, two parameters are varied, while the remaining parameters are equal to the optimum cylinder. Colour contours indicate the expected value of the objective function, while solid lines indicate contours of the uncertainty (standard deviation of the objective function). The objective functions are shown for both the coarse-mesh and fine-mesh optimisations.

For $U^* = 5$, the objective function increases with increasing AR , while it is maximised for a moderate camber $\beta = 0.1$ (figure 10a). The parameters of the top 10 performing cylinders are indicated by blue markers. While many are clustered near the local maximum, a second peak is also observed with much lower aspect ratio ($AR \approx 8$). Note that the remaining parameters (α , m_f , n_f , m_a and n_a) are also different near this second local maximum. The value of the objective function for this second local maximum is within 2 % of the best-performing cylinder. Panel (b) shows that the objective function is maximised for $m_a = 2$, indicating a sharp upper and lower vertices.

To reduce the design space for the fine-mesh computation, a reduced domain is considered. The maximum and minimum of each parameter are determined by taking

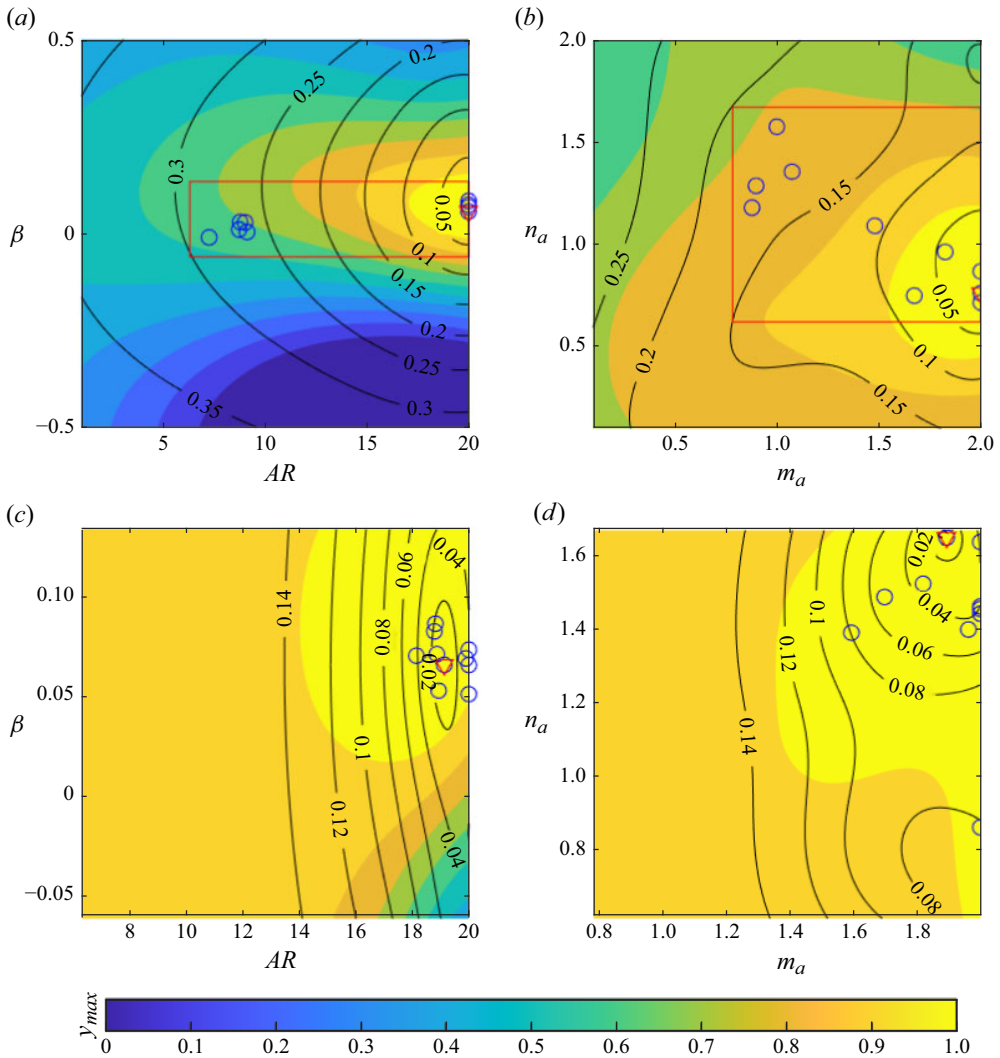


Figure 10. Two-dimensional slices of the objective function for (a,c) AR and β and (b,d) m_a and n_a , with remaining parameters equal to the optimum cylinder for $U^* = 5$ (OpU5). The objective function obtained using the coarse-mesh optimisation is shown in (a,b), while the fine-mesh optimisation is presented in (c,d). Colour contours show the expected value of y_{max} , while the solid contours show the uncertainty (standard deviation). The red triangular marker indicates the best-performing cylinder, while blue circles indicate the top 10 cylinders. Finally, the red rectangle indicates the reduced domain for the fine-mesh optimisation.

the maximum/minimum of this parameter spanned by the top 10 performing cylinders, then adding a 5 % margin (based on the overall size of the design space). The top 10 cylinders from the coarse-mesh optimisation are used as the initial batch. The reduced domain is indicated by the red rectangle in panels (a) and (b). Panel (c) shows that the local maximum identified in the coarse-mesh domain is also a local maximum on the fine mesh. Panel (d) confirms that the maximum objective function is observed for $m_a \approx 2$. The maximum value of n_a is close to the upper limit on the reduced domain. Hence, the reduced domain artificially constrains this parameter. However, this parameter has only a small effect on the objective function near the optimum, so the reduced domain is not likely to significantly affect the value of the objective function.




Figure 11. Two-dimensional slices of the objective function for (a,c) AR and β and (b,d) m_a and n_a , with remaining parameters equal to the optimum cylinder for $U^* = 10$ (OpU5). The objective function obtained using the coarse-mesh optimisation is shown in (a,b), while the fine-mesh optimisation is presented in (c,d). Colour contours show the expected value of y_{max} , while the solid contours show the uncertainty (standard deviation). The red triangular marker indicates the best-performing cylinder, while blue circles indicate the top 10 cylinders. Finally, the red rectangle indicates the reduced domain for the fine-mesh optimisation.

Figure 11 shows slices of the objective function for $U^* = 10$. Compared with the $U^* = 5$ optimisation, the design space is not widely explored, with the exploration concentrated near a local optimum solution. Thus, it is possible that other peaks exist, which were not adequately explored. Regardless, the optimised cylinders determined in this study significantly out-perform the geometries considered in prior studies.

Figure 12(a) plots the variation of oscillation amplitude (y_{max}) against U^* for each of the optimised cylinders. The performance of circular and elliptical cylinders with various aspect ratios are also shown for reference. Compared with the elliptical cylinders, which display only a VIV behaviour at this Re , the optimised shapes exhibit galloping-like behaviour. Specifically, there is a seemingly unbounded linearly growth in y_{max} with U^* ,

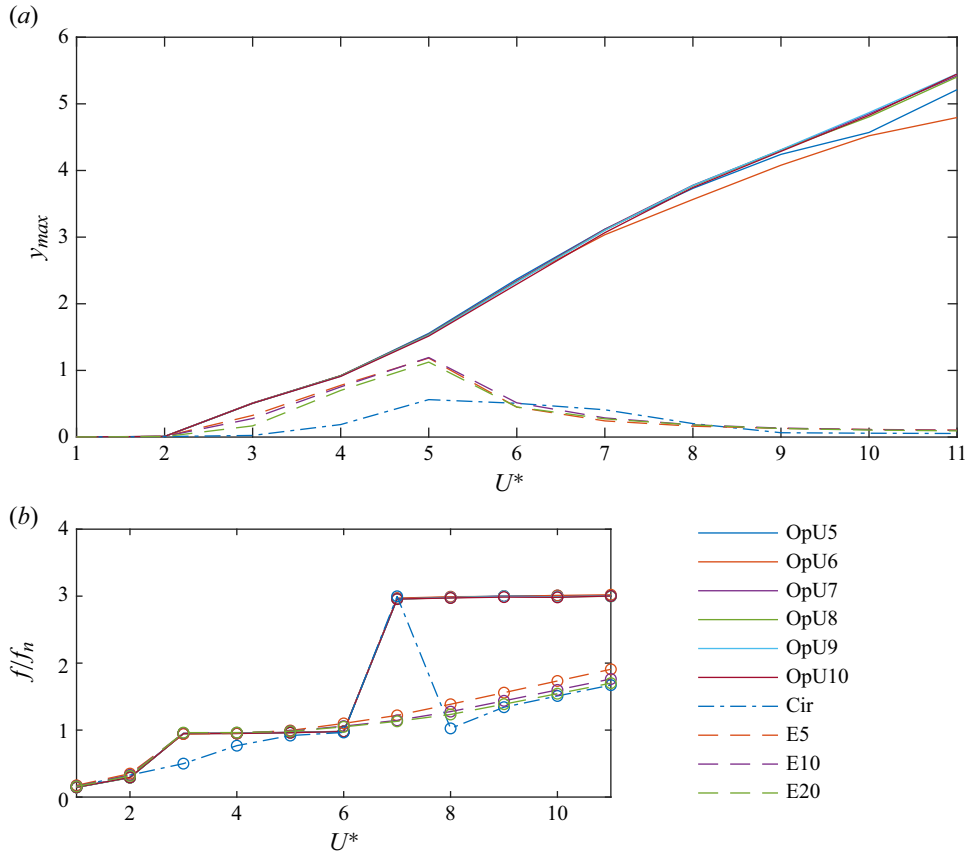


Figure 12. Variation of (a) amplitude of oscillation y_{max} and (b) frequency ratio f/f_n against U^* , where f is the dominant vortex shedding frequency and f_n is the natural frequency of the system, for the optimised cylinders. Results for the circular cylinder (Cir) and elliptical cylinders with aspect ratio 5 (E5), 10 (E10) and 20 (E20) are also shown.

which is characteristic of galloping. In § 5, we show that this behaviour is a combination of VIV and galloping. Panel (b) plots the ratio of vortex shedding frequency (determined using the maximum power spectral density of the transverse force coefficient) to the undamped natural frequency ($f_n = 1/U^*$) against U^* . For $3 \leq U^* \leq 6$, the dominant vortex shedding frequency matches the natural structural frequency, indicating lock-in behaviour. For $U^* > 6$, the vortex shedding frequency is three times the natural frequency. This still indicates a lock-in behaviour, but that the third harmonic dominates over the fundamental. The elliptical and circular cylinders become desynchronised for large U^* and the vortex shedding frequency differs from the system natural frequency.

Referring to figure 12(a), the OpU5 cylinder is the best-performing cylinder at $U^* = 5$, 6 and 7, the OpU7 cylinder is the best at $U^* = 8$, and the OpU9 cylinder is best for $U^* = 9$, 10 and 11. Several cylinders (OpU6, OpU7, OpU8 and OpU10) are not the best-performing cylinder at their design U^* . This indicates that the optimisation had not converged to a global maximum. However, each of these cylinders is within 1 % of the best-performing cylinder at their design U^* , so the optimisation is satisfactory.

The OpU6 cylinder performs satisfactorily (within 2 % of the best-performing cylinder) up to $U^* = 6$, but performs poorly for high U^* . The OpU5 cylinder is within 2 % of the




Figure 13. Comparison between the cylinder optimised for a fixed $m^* = 4.7124$ (OpU8), and the cylinder optimised for fixed $m_r = 15$ (OpU8_MR). The variation of maximum amplitude against U^* for both shapes is plotted for both $m^* = 4.7124$ and $m_r = 15$. Finally, A_c is the cross-sectional area of each cylinder.

best-performing cylinder up to $U^* = 9$, but performs poorly for $U^* \geq 10$. The remaining cylinders (OpU7, OpU8, OpU9 and OpU10) perform within 2 % of the best-performing cylinder over the entire range of U^* . Of these, the OpU7 cylinder has the lowest overall error, with less than 1 % difference from the best-performing cylinder at each U^* .

The blade-like geometry of the optimised cylinder suggests a similarity between the present study and the VIV of an airfoil. Benner *et al.* (2019) report that an airfoil aligned perpendicularly to the incoming flow achieves the highest amplitude of VIV. However, they do not observe the galloping-like behaviour seen in figure 12, with the maximum oscillation amplitude occurring at $U^* = 6$. This may be a result of either the lack of camber in their study, the fact that an airfoil has only one sharp edge, or differences in the system parameters (e.g. Reynolds number).

4.1. Optimising with a fixed mass ratio

The optimisations presented in the previous section are performed for a fixed dimensionless mass m^* . This assumes that the total oscillating mass is independent of the cross-sectional area of the cylinder. In this section, optimisations are performed with a fixed mass ratio $m_r = 15$, which assumes that the total oscillating mass is proportional to the cylinder's cross-sectional area. Three reduced velocities are considered ($U^* = 6$, $U^* = 8$ and $U^* = 10$), while the Reynolds number is held constant ($Re = 100$). Cylinders optimised for fixed m_r and $U^* = X$ are denoted OpUX_MR. The shape parameters for these cylinders are included in table 5.

Figure 13 presents a comparison between the OpU8 and OpU8_MR cylinders. For both cylinders, the aspect ratio is approximately $AR = 20$, which is the maximum AR considered in this study. While both cylinders have the same nominal AR, the OpU8_MR cylinder has a lower area than the OpU8 cylinder. As defined in this study, AR is the ratio between cylinder height and maximum cylinder width. The OpU8_MR cylinder achieves a lower area by having a sharp kink at the cylinder midpoint, so that the cylinder thickness away from the midpoint is reduced. The OpU8 cylinder has a smooth surface, with a greater overall area. For a fixed mass (m^*), the OpU8 cylinder outperforms




Figure 14. Optimised cylinders obtained for three different Reynolds numbers, at $U^* = 10$, $m^* = 4.7124$ and $\xi = 0$.

the OpU8_MR, indicating that the smooth surface leads to improved hydrodynamic performance. However, when the mass ratio (m_r) is fixed, the thinner OpU8_MR cylinder has a lower mass, and achieves a greater oscillation amplitude, indicating that the reduction in mass due to a thinner profile outweighs the hydrodynamic penalty due to the sharp kink. Overall, however, the difference in performance between the OpU8 and OpU8_MR cylinders is relatively small.

4.2. Effect of Reynolds number

In this section, optimisations are performed for two different Reynolds numbers, $Re = 60$ and $Re = 80$, to determine the influence of Re on the optimal cylinder shape within the two-dimensional laminar flow regime. The optimised cylinders are labelled OpR60 and OpR80, respectively. These optimisations are performed for $\xi = 0$, $m^* = 4.7124$ and $U^* = 10$. Figure 14 plots the optimised cylinders alongside the OpU10 cylinder, and the shape parameters are presented in table 5. Increasing Reynolds number has only a small effect on the optimal cylinder shape, with a slight decrease in the amount of camber as Re increases from $x_{cam} = 0.15$ at $Re = 60$, to $x_{cam} = 0.11$ at $Re = 100$.

5. Analysis of the optimised shape

The optimised shape obtained in this study is a thin crescent-shaped cylinder with sharp leading and trailing edges. This shape is similar to the curved blades proposed for a galloping energy harvester by Tucker Harvey *et al.* (2019, 2020). Compared with traditional bluff-body harvesting devices, the curved blade geometry was found to promote flow attachment, leading to an improved efficiency. A similar attached flow was observed by Lo *et al.* (2023) for the thin elliptical cylinder, which also results in significant improvements in power generation over traditional bluff body shapes (Lo *et al.* 2024).

Lo *et al.* (2023) propose that the observed behaviour can be explained as a combination of a vortex-induced force, and a movement-induced force (galloping). The interaction between VIV and galloping is also observed for square and rectangular cylinders (Bouclin 1979; Mannini, Marra & Bartoli 2014; He, Yang & Jiang 2018), under certain experimental conditions, and was used to enhance energy harvesting from FIV for a modified circular cylinder by Sun *et al.* (2019a).

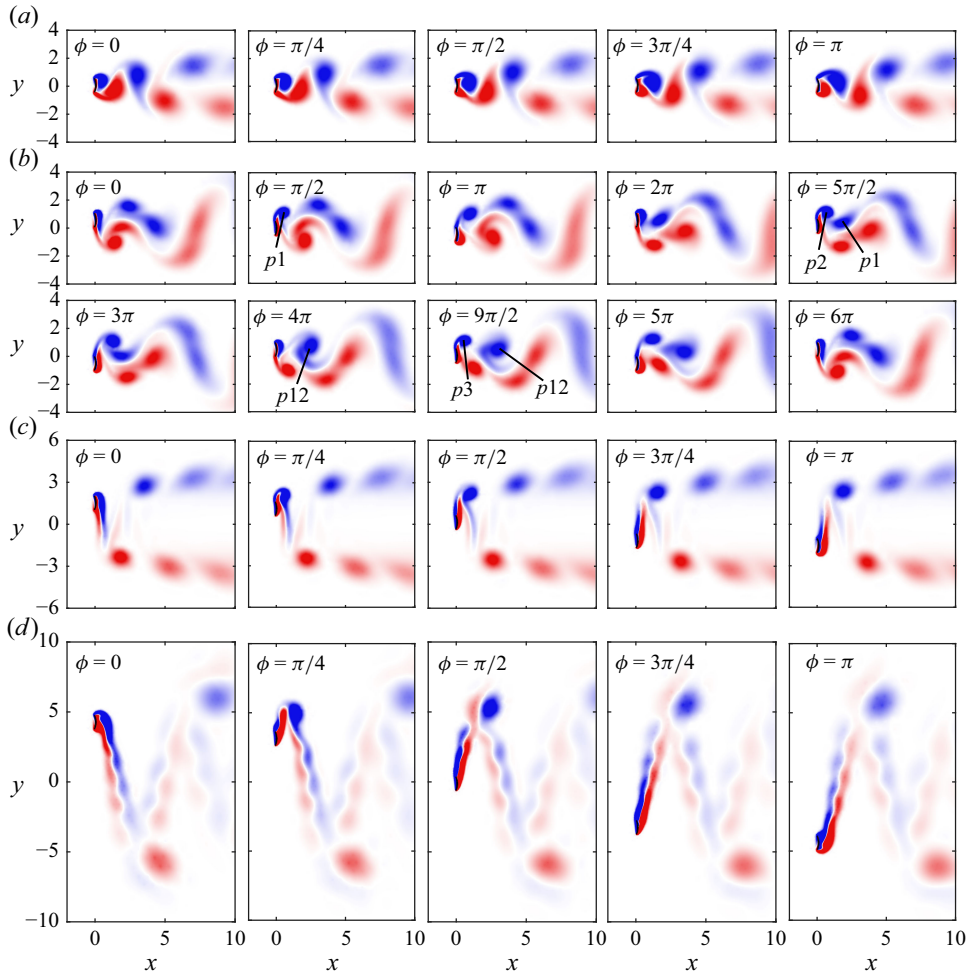


Figure 15. Visualisations of spanwise vorticity in the wake of the OpU9 cylinder, at (a) $U^* = 1$, (b) $U^* = 3$, (c) $U^* = 5$ and (d) $U^* = 10$. Here, ϕ is the phase angle, with $\phi = 0$ corresponding to the maximum displacement y of the cylinder. Red indicates positive (counter-clockwise) vorticity, while blue indicates negative (clockwise) vorticity.

The present section investigates the dynamics of the optimised cylinder from the perspective of combined FIV and galloping behaviour. First, in § 5.1, we investigate the vortex structures formed in the wake. Next, in § 5.2, a quasi-steady analysis is performed, to investigate galloping behaviour. Finally, in § 5.3, the quasi-steady analysis is compared with the instantaneous force coefficients, demonstrating that unsteady forces due to vortex shedding contribute substantially to the FIV behaviour.

5.1. Vorticity contours

Figure 15 presents visualisations of the vorticity field in the wake of the OpU9 cylinder, at four different values of U^* . Visualisations are provided at various phase angles $\phi = \omega t$, where ω is the oscillation frequency and $\phi = 0$ corresponds to the maximum displacement of the cylinder. Similar vorticity visualisations are obtained for all other optimised cylinders, and are not shown. Animations are provided in the online supplementary material (movies 1–4).




Figure 16. Velocity streamlines computed in a frame of reference moving with the cylinder, for (a) $U^* = 5$ and (b) $U^* = 9$, at the point of maximum vertical velocity ($\phi = \pi/2$). A visualisation of the vorticity field is also provided, with red indicating positive (counter-clockwise) vorticity and blue indicating negative (clockwise) vorticity.

For $U^* \leq 2$ (figure 15a), the oscillation amplitude is negligible. In this regime, vortices of alternating orientation are periodically shed from the cylinder, similar to the von Kármán vortex street behind a flat plate (Najjar & Vanka 1995).

At $U^* = 3$ (figure 15b), substantial FIVs begin to occur. As shown in figure 12(b), the frequency of vibration matches the natural frequency of the system. Vortex shedding is also synchronised to the cylinder's vibrations. Referring to figure 15(b), a single vortex is shed from the cylinder every half-period of the motion, after the cylinder begins to travel away from the point of maximum displacement (between $\phi = 0$ and $\phi = \pi/2$). However, there is an interesting interaction with a period of three full vortex shedding cycles. First, a vortex $p1$ is shed from the cylinder, at $\phi = \pi/2$. A second vortex $p2$ is then shed at the next cycle of motion ($\phi = 5\pi/2$). These two vortices merge to form a single vortex $p12$ ($\phi = 4\pi$). Finally, a third vortex is shed at the next cycle of motion ($\phi = 9\pi/2$). The two vortices $p12$ and $p3$ are then advected downstream of the cylinder.

For $4 \leq U^* \leq 8$ (figure 15c) the wake structure is reminiscent of the hyperbranch observed by Lo *et al.* (2023) for the thin elliptical cylinder. Specifically, a single vortex is shed from the cylinder every half-cycle of the motion (between $\phi = 0$ and $\phi = \pi/4$), after the cylinder changes direction at the point of maximum amplitude. At $\phi = \pi/2$, where the cylinder's vertical velocity is large, flow appears to be attached to the cylinder, with behaviour similar to an airfoil. Figure 16 shows velocity streamlines near the cylinder at

$\phi = \pi/2$, computed in a frame of reference attached to the cylinder. The streamlines are qualitatively similar to flow over an airfoil. Specifically, flow attaches to the cylinder at the leading edge, and separates at the trailing edge. However, there is a small recirculation region, indicating some flow separation from the rear surface. The similarity to a lifting airfoil suggests that there will be a relatively low drag force and a significant lift force with a component in line with the cylinder's velocity, which could enhance FIV. This proposal will be investigated in § 5.2. Moreover, the shedding of a single vortex at the peak cylinder displacement resembles the shedding of a starting/stopping vortex from an airfoil. Attached airfoil-like flow has been reported in previous FIV studies of thin elliptical cylinders (Lo *et al.* 2023) and curved blades (Tucker Harvey *et al.* 2019; 2020). Lo *et al.* (2023) also report a similar pattern of vortex shedding to that observed in the present study.

For $U^* \geq 9$, (figure 15*d*) the shear-layer attached to the cylinder becomes unstable, resulting in the alternate shedding of many small vortices throughout the motion. These small-scale vortex structures (SVS) were also observed by Lo *et al.* (2023) for the elliptical cylinder. In addition to the SVS, a strong primary vortex is still shed from the cylinder each half-cycle of motion, as the cylinder begins to travel away from the point of maximum displacement.

Figure 16 plots streamlines near the cylinder for $U^* = 9$, at $\phi = \pi/2$. Once again, the flow is qualitatively similar to the flow over an airfoil. Flow attaches at the leading edge, and there is a separation point at the trailing edge. However, flow separation at the rear surface is more severe at $U^* = 9$ than at $U^* = 5$, resulting in a larger recirculation bubble.

5.2. Quasi-steady analysis

Galloping is a movement-induced instability driven by changes in the transverse force due to variations in the inclination angle between the cylinder and the relative flow (Parkinson 1989). Galloping is often analysed using a quasi-steady analysis (Parkinson 1989; Barrero-Gil, Alonso & Sanz-Andres 2010; Javed, Abdelkefi & Akhtar 2016; Tucker Harvey *et al.* 2020), which assumes the fluid forces are a function only of the relative angle of attack between the cylinder and the fluid. In this section, we present a quasi-steady analysis of the optimised cylinder, showing that a hard-galloping behaviour is possible under the quasi-steady model. This is motivated by two observations. First, the variation of amplitude with U^* (figure 12) resembles galloping behaviour. Second, the proposed airfoil-like behaviour suggests a strong influence of relative angle of attack on the force coefficients, which may be predicted using quasi-steady analysis.

Note that the quasi-steady assumption is only valid at high U^* , when the oscillation frequency is much lower than the vortex shedding frequency. As shown in § 5.1, the frequency of vortex shedding for the optimised cylinder is equal to the oscillation frequency over the range of U^* considered in this study. Therefore, while the quasi-steady model indicates how the relative angle of attack affects the hydrodynamic forces, it does not capture how unsteady effects such as vortex shedding impact the motion. The influence of vortex shedding on the motion is discussed later, in § 5.3.

As shown in figure 17, the relative velocity and inclination angle between the cylinder and the oncoming flow are

$$U_{rel} = U\sqrt{1 + \dot{y}^2} = U\sqrt{1 + \tan^2 \alpha}, \quad (5.1)$$

$$\tan \alpha = \dot{y}. \quad (5.2)$$

Under the quasi-steady assumption, the lift and drag forces depend only on the inclination angle α . The quasi-steady transverse force can be computed from the lift and drag coefficients as




Figure 17. Sketch of the relative velocity experienced by the cylinder. Here, \dot{y} is the velocity of the cylinder, while U_{rel} and α are the magnitude and incidence angle of the relative velocity. Finally, C_L and C_D are the lift and drag coefficients, while C_y and C_x are the transverse and streamwise force coefficients.

$$C_{y,q}(\alpha) = (1 + \tan^2 \alpha)(C_{L,q}(\alpha) \cos \alpha - C_{D,q}(\alpha) \sin \alpha), \quad (5.3)$$

where $C_{L,q}$ and $C_{D,q}$ are the quasi-steady lift and drag coefficients. The instantaneous quasi-steady power coefficient can also be computed, as

$$C_{p,q}(\alpha) = C_{y,q} \dot{y} = \tan \alpha C_{y,q}(\alpha). \quad (5.4)$$

The quasi-steady power coefficient represents the transfer of energy from the fluid to the cylinder, due to the transverse force. When this term is positive, the quasi-steady forces enhance the vibration of the cylinder. Note that C_L and C_D are non-dimensionalised by the relative velocity U_{rel} , while C_y and $C_{p,q}$ are non-dimensionalised by the free-stream velocity U . Moreover, the effective Reynolds number for computing C_L and C_D is

$$Re_{rel} = Re \sqrt{1 + \tan^2 \alpha}, \quad (5.5)$$

where Re is the Reynolds number based on free-stream velocity.

The quasi-steady force coefficients are estimated from numerical simulations performed for a stationary cylinder, with fixed inclination angle and Reynolds number determined by (5.5). The inclination angle is varied between $\alpha = 0$ and $\alpha = (7/16)\pi$, resulting in Reynolds numbers between $Re_{rel} = 100$ and $Re_{rel} = 512$. Simulations are not performed at $\alpha = \pi/2$, as this results in an infinite relative Reynolds number.

Figure 18 presents the variation of $C_{L,q}$, $C_{D,q}$, $C_{y,q}$ and $C_{p,q}$ against α , for the OpU9 cylinder. The drag coefficient $C_{D,q}$ is largest at $\alpha = 0$, when the cylinder is perpendicular to the oncoming flow. As α is increased, the profile is more streamlined, resulting in a substantial decrease in drag coefficient. The lift coefficient $C_{L,q}$ is zero at $\alpha = 0$, due to symmetry. At large α , the cylinder profile is similar to a cambered airfoil, generating a substantial positive lift coefficient.

From (5.3), the lift force has a contribution in line with the cylinder's direction of motion, while the drag force has a component opposed to the cylinder's motion. For small α , the lift coefficient is small, and therefore the drag term dominates. This produces a negative power coefficient, opposing the cylinder motion. For moderate inclination angles ($0.84 < \alpha < 1.27$) the lift force is large, and dominates the drag term in (5.3), leading to a positive power coefficient, increasing the cylinder motion. Although the lift force remains large for $\alpha > 1.27$, $\cos \alpha$ becomes small (i.e. the lift force is not sufficiently aligned in the transverse direction), so that the contribution of the lift force to the power coefficient decreases, and the power coefficient becomes negative.

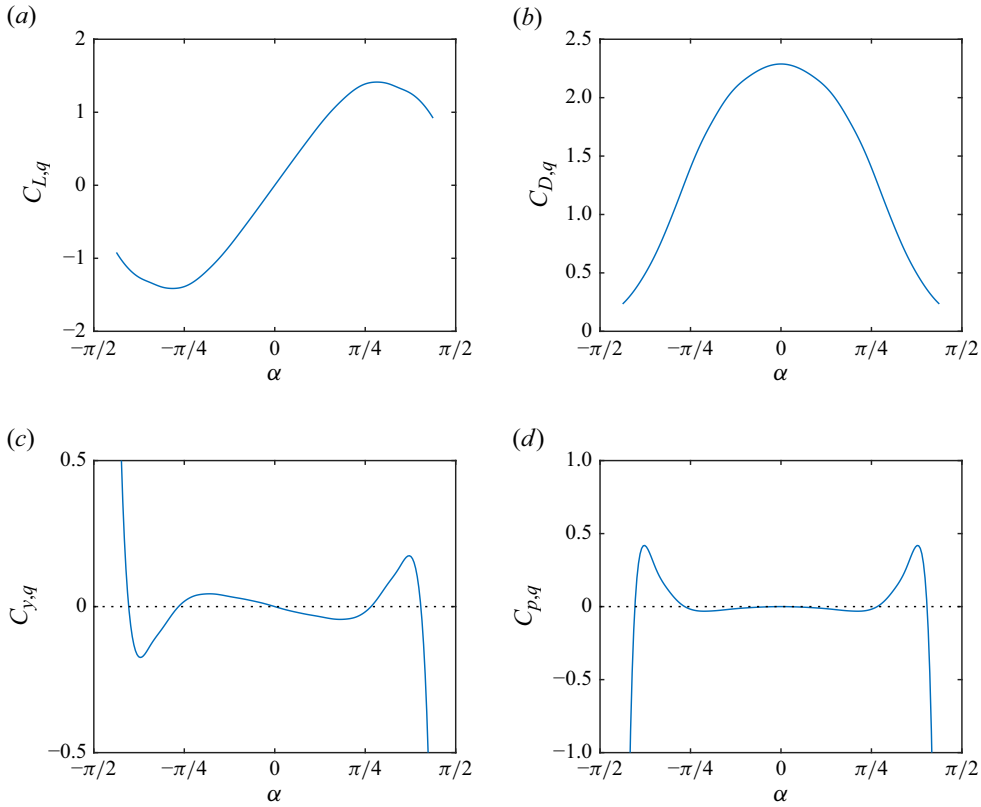


Figure 18. Variation of the (a) lift $C_{L,q}$, (b) drag $C_{D,q}$, (c) transverse $C_{y,q}$ and (d) power $C_{p,q}$ coefficients against inclination angle α under the quasi-steady assumption, obtained from simulations of a stationary OpU9 cylinder. The inclination angle was varied from $0 \leq \alpha \leq (7/16)\pi$, and symmetry used to compute the force coefficients for negative α . Simulations were not performed at $\alpha = \pi/2$, since the Reynolds number becomes infinite.

Therefore, according to the quasi-steady analysis, there is a range of relative inclination angles ($0.84 < \alpha < 1.27$) where the component of the lift force in the transverse direction exceeds the drag force, leading to a positive instantaneous power coefficient, increasing the energy of the cylinder. This is due to the large lift force and small drag force occurring at large inclination angles, due to the airfoil-like shape.

To determine if galloping is possible, the power coefficient is integrated over a single oscillation cycle (Parkinson & Smith 1964; Tucker Harvey *et al.* 2020). We assume the cylinder displacement is sinusoidal

$$y = A \sin \omega t, \quad \dot{y} = A\omega \cos \omega t, \quad (5.6)$$

so that the instantaneous power coefficient is expressed as

$$C_{p,q} = C_{y,q}(\tan^{-1}(A\omega \cos(\omega t))) A\omega \cos(\omega t). \quad (5.7)$$

The cycle-averaged power coefficient is a function only of the combined parameter $A\omega$

$$\bar{C}_{p,q}(A\omega) = \frac{A\omega}{2\pi} \int_0^{2\pi} C_{y,q}(\tan^{-1}(A\omega \cos u)) \cos u \, du. \quad (5.8)$$

Figure 19(a) presents the variation of $\bar{C}_{p,q}(A\omega)$ against $A\omega$ for the OpU9 cylinder. For small-amplitude oscillations ($A\omega < 1.29$), the net power added to the cylinder during

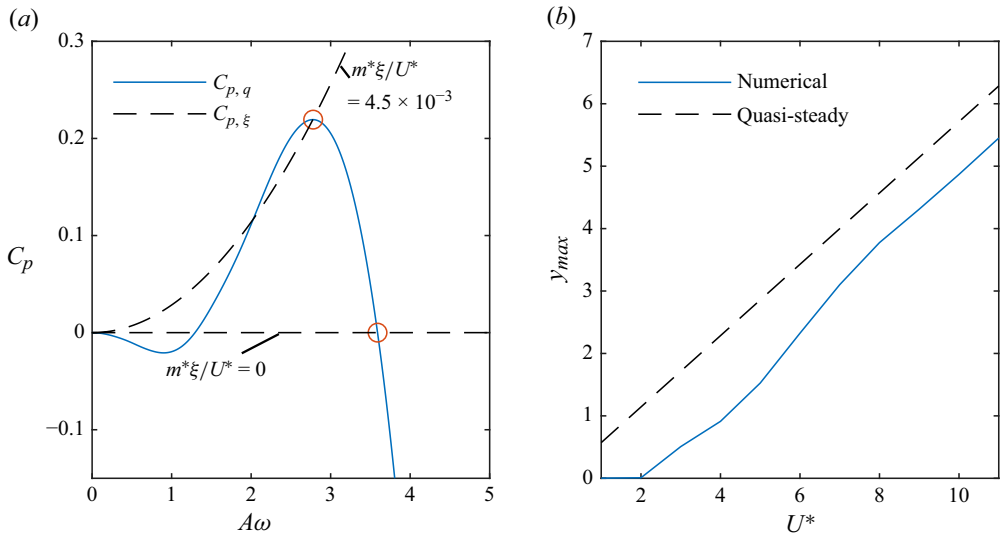


Figure 19. (a) Variation of the cycle-averaged power coefficient for the quasi-steady model ($\bar{C}_{p,q}$), against $A\omega$, for the OpU9 cylinder. Dashed lines show the power extracted by mechanical damping ($\bar{C}_{p,\xi}$), for either $m^*\xi/U^* = 0$ (undamped), or $m^*\xi/U^* = 4.5 \times 10^{-3}$ (optimally damped). Red circles indicate stable limit cycles. (b) Comparison between the predicted maximum amplitude obtained from the quasi-steady model and numerical simulations.

one cycle is negative, and therefore the amplitude of oscillation will decrease to zero. For $1.29 < A\omega < 3.59$, the cycle-averaged power is positive, so that the amplitude will increase until it reaches a value of $A\omega = 3.59$. For $A\omega > 3.59$, the cycle-averaged power is again negative, and the amplitude of oscillation will decrease.

For an undamped system, this indicates that there are two stable equilibria: a stationary cylinder ($A\omega = 0$), and a stable periodic limit cycle with $A\omega = 3.59$. There is also an unstable equilibrium at $A\omega = 1.29$. This indicates hard-galloping behaviour, in which small-amplitude oscillations are stable, and will decay to zero. However, given a sufficient initial displacement ($A\omega > 1.29$), the body will gallop with a large amplitude $A\omega = 3.59$.

Assuming the system vibrates at the natural frequency $\omega = \sqrt{k/m} = 2\pi/U^*$, the amplitude of oscillation predicted by the quasi-steady model is $A = 0.5714U^*$. Figure 19(b) compares this prediction with the numerical simulations, showing reasonable agreement, given the limitations of the quasi-steady model. However, the amplitude obtained from numerical simulations is lower than the predictions obtained from the quasi-static approximation.

5.3. Instantaneous force coefficients

The quasi-steady analysis presented in § 5.2 assumes that the hydrodynamic force applied to the cylinder depends only on the relative angle of attack α . However, the shedding of a primary vortex each half-cycle of motion (figure 15) will also affect the forces, which will not be captured by the quasi-steady model. In this section, we compare the force coefficients obtained from numerical simulations with the predictions of the quasi-static model, to deduce the relative significance of the quasi-steady forces and the vortex-induced forces to the motion.

Bouclin (1979) models the VIV–galloping interaction by including an unsteady fluctuating fluid force, representing the vortex-induced forces, in addition to the




Figure 20. Comparison between the instantaneous (a) lift, (b) drag, (c) transverse force and (d) power coefficients obtained from numerical simulations at various U^* and the quasi-steady model. Solid lines indicate $y > 0$, while dashed lines indicate $y < 0$.

quasi-steady forces. Therefore, we interpret the difference between the instantaneous and quasi-steady force coefficients as representing the effects of vortex shedding on the instantaneous force and moment coefficients.

Figure 20 presents a comparison between the instantaneous lift (C_L), drag (C_D) transverse force (C_y) and power (C_p) coefficients for $U^* = 4, 6, 8$ and 10 , to the quasi-static prediction. Here, the transverse and streamwise force coefficients, C_y and C_x , were obtained directly from the numerical simulations. Then, the added mass force was subtracted from the transverse force, to obtain the vortex force ($C'_y = C_y + 2m_{a.m.}^* \ddot{y}$, where $m_{a.m.}^*$ is the dimensionless added mass). To estimate the added mass, a numerically simulated pluck test in still fluid is performed using the same numerical set-up described in § 2.3, but with $U = 0$ at the inlet. The cylinder is given an initial displacement, and the frequency of the subsequent oscillations is measured. The added mass is then estimated from the relationship $\omega = \sqrt{k/(m + m_{a.m.})}$, where ω is the angular frequency of oscillations.

The lift and drag coefficients were then computed according to

$$C_L = (C'_y \cos \alpha + C_x \sin \alpha) / (1 + \tan^2 \alpha), \quad (5.9)$$

$$C_D = (C'_y \sin \alpha - C_x \cos \alpha) / (1 + \tan^2 \alpha). \quad (5.10)$$

While the quasi-steady model captures the general behaviour of C_L and C_D , particularly for large inclination angles, substantial deviations are observed. This indicates that, while the quasi-steady model provides useful information about the effect of α on the lift and drag coefficients, other factors, such as vortex shedding, are also significant.

In [figure 20](#), solid lines indicate $y > 0$, while dashed lines indicate $y < 0$. For $y > 0$ and $\alpha > 0$ (or $y < 0$ and $\alpha < 0$), the cylinder is in decelerating motion as it approaches the point of maximum displacement. In this region, the lift and drag coefficients are both slightly lower than the quasi-static prediction. However, the power coefficient is negative, indicating that the drag term dominates over the lift. For $y > 0$ and $\alpha < 0$ (or $y < 0$ and $\alpha > 0$), the cylinder is accelerating as it travels away from the point of maximum displacement. In this region, both C_L and C_D are substantially larger than the quasi-static prediction. The power coefficient is positive, indicating that the lift term dominates over the drag.

Note that C_y , and therefore also C_p , contains opposite-signed contributions from C_L and C_D (5.3). For the quasi-steady model, these contributions are close to equal for small to moderate α , so that the quasi-steady transverse force contribution $C_{y,q}$ is relatively small. Although the instantaneous lift and drag coefficients approximately follow the quasi-steady prediction, the instantaneous transverse force and power coefficient do not obviously follow the quasi-steady prediction, and are instead dominated by other unsteady contributions, such as vortex shedding.

[Figure 20\(b\)](#) shows a sharp decrease in the power coefficient when the inclination angle exceeds $\alpha = 1.27$, for both quasi-steady and instantaneous power coefficients. This is due to the poor alignment between the lift force and the transverse direction at high inclination angles. Therefore, the quasi-steady force will limit the maximum inclination angle, which limits the maximum value of $A\omega$.

[Figure 21](#) presents the time history of the instantaneous force and power coefficients (C_y and C_p), as well as the quasi-steady predictions ($C_{y,q}$ and $C_{p,q}$), against the phase angle $\phi = \omega t$, where ω is the natural frequency. In general, there is little correlation between the instantaneous force coefficients and the quasi-steady predictions, indicating that the transverse force is dominated by vortex shedding or other unsteady effects.

For moderate and high U^* ($U^* \geq 6$), the motion appears to be driven by a large peak in both transverse force and power coefficient, which occurs immediately after the point of maximum displacement (between $0 \leq \phi \leq \pi/4$ and $\pi \leq \phi \leq 5\pi/4$). As shown in [figure 15](#), this corresponds to the shedding of a primary vortex from the cylinder. Therefore, there is a substantial contribution from VIV to the observed motion, which enhances FIVs.

Thus, we have the following interpretation of the motion at high U^* . A primary vortex is shed each half-cycle, as the cylinder accelerates away from the point of maximum displacement, similar to the shedding of a starting vortex. This shed vortex results in a large transverse force in line with the cylinder velocity, adding energy to the cylinder which enhances FIV motion. As the cylinder travels downwards, there is a large inclination angle relative to the flow. This produces both a drag force, which opposes the motion, and a lift force, which enhances FIV motion. For $\alpha \lesssim 1.27$, the quasi-steady lift and drag terms are approximately equal, so the quasi-steady effect provides a negligible contribution to the transverse force. However, for $\alpha \gtrsim 1.27$, the component of lift in the transverse direction becomes small, so that the quasi-steady drag dominates. This effectively limits the maximum amplitude for a given frequency, so that the amplitude increases approximately linearly with U^* .

This section has demonstrated that, while the quasi-steady model captures some effect of the relative angle of attack, unsteady forces are also substantial. Given the similarities

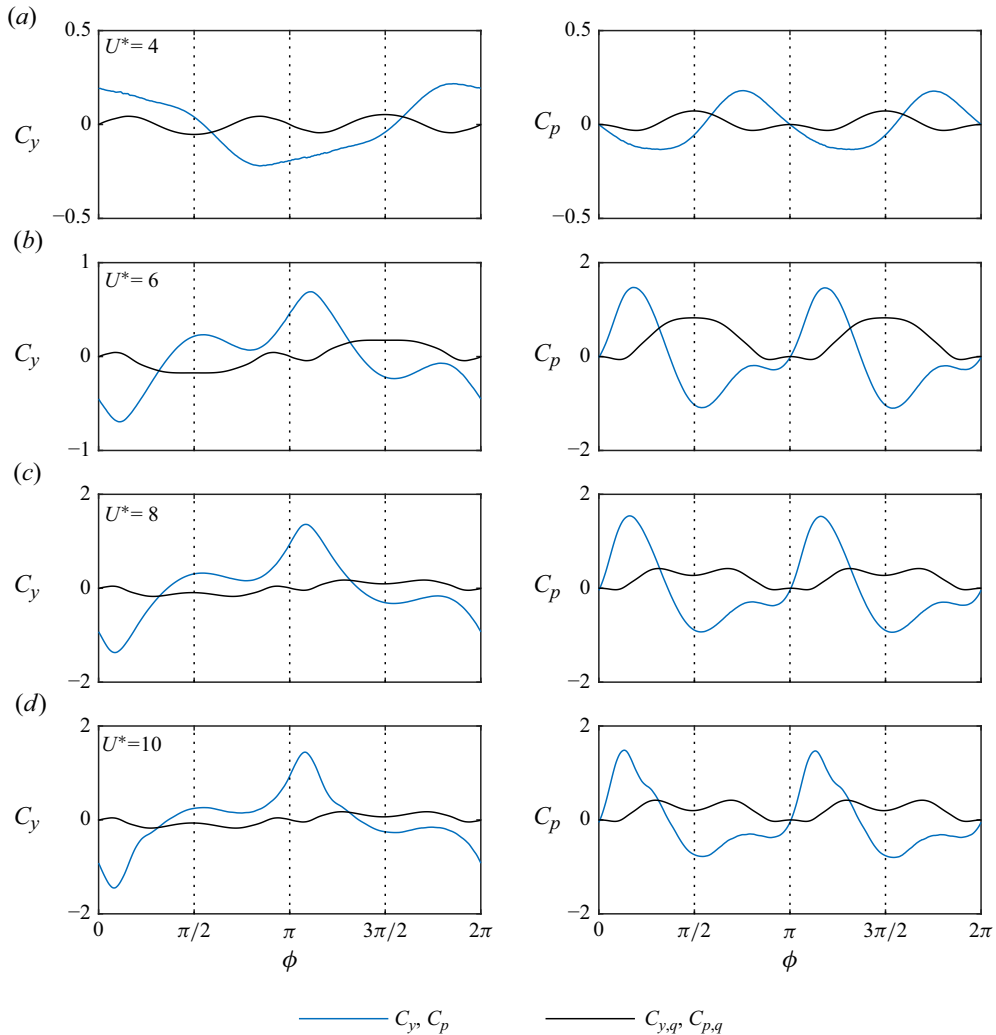


Figure 21. Comparison between the instantaneous transverse force coefficient (left) and power coefficient (right) and the predictions of the quasi-steady model, for (a) $U^* = 4$, (b) $U^* = 6$, (c) $U^* = 8$ and (d) $U^* = 10$.

between the crescent cylinders and airfoil/turbine blades, modelling approaches developed for dynamic stall (Bangga, Lutz & Arnold 2020; Mohamed *et al.* 2024) may prove useful in providing further insight into the observed motion. These models include contributions from attached flow, separated flow and vortex shedding all of which are observed in the present flow.

6. Optimising power generation

We have thus far considered undamped vibrations, with the optimisation performed to maximise the vibration amplitude. For such cases, the power generation is zero. Attaching a transducer to the system will extract mechanical energy, which is modelled here as a linear damper. The power coefficient (2.8) obtained under this assumption represents the




Figure 22. Comparison of the (a) power coefficient \bar{C}_p and (b) amplitude parameter $A\omega$ predicted from the quasi-steady analysis, and numerical simulations performed at various U^* , for the OpU9 cylinder at $Re = 100$ and $m^* = 4.7124$.

maximum power that can be extracted for the given damping coefficient ξ , assuming no other sources of damping.

The structure of this section is as follows. First, in § 6.1, we estimate the power coefficient using the quasi-steady model, and compare this with the numerical simulations. Then, in § 6.2, a shape optimisation is performed to maximise the power coefficient.

6.1. Power generation according to the quasi-steady model

Assuming a sinusoidal displacement profile $y = A \sin \omega t$, the power extracted from the cylinder via mechanical damping (2.8) is

$$\bar{C}_{p,\xi} = \frac{2\pi m^*\xi}{U^*} (A\omega)^2, \quad (6.1)$$

which depends on two parameters: the amplitude parameter $A\omega$, and the mass-damping parameter $m^*\xi/U^*$. Figure 19 plots $\bar{C}_{p,\xi}$ against $A\omega$, for $m^*\xi/U^* = 4.5 \times 10^{-3}$. When $\bar{C}_{p,\xi} > \bar{C}_{p,q}$, the energy added to the cylinder by quasi-steady forces over one cycle is less than the energy extracted by mechanical damping, and so the amplitude will decrease. Likewise, the amplitude will increase for $\bar{C}_{p,\xi} < \bar{C}_{p,q}$. Therefore, there is a stable equilibrium when $C_{p,\xi} = \bar{C}_{p,q}$ and $\partial \bar{C}_{p,\xi} / \partial (A\omega) > \partial \bar{C}_{p,q} / \partial (A\omega)$, indicated by the red circle in figure 19. The value of $m^*\xi/U^* = 4.5 \times 10^{-3}$ corresponds to the maximum value of $\bar{C}_{p,\xi}$ at this stable equilibrium, and therefore the maximum power generation under the quasi-static assumption.

Figure 22 plots a comparison between the predicted power coefficient and vibration amplitude according to the quasi-steady model, and numerical simulations of the OpU10 cylinder, performed for various values of U^* . The quasi-steady model predicts a maximum power coefficient of $\bar{C}_{p,\xi} = 0.219$, occurring at $m^*\xi/U^* = 4.5 \times 10^{-3}$, and that galloping cannot be sustained for $m^*\xi/U^* > 5.2 \times 10^{-3}$.

However, the numerical simulations demonstrate that significant FIVs are sustained for damping coefficients much larger than this, and yield significantly larger maximum power coefficients (up to $\bar{C}_{q,\xi} = 0.364$). This shows that the energy added to the cylinder

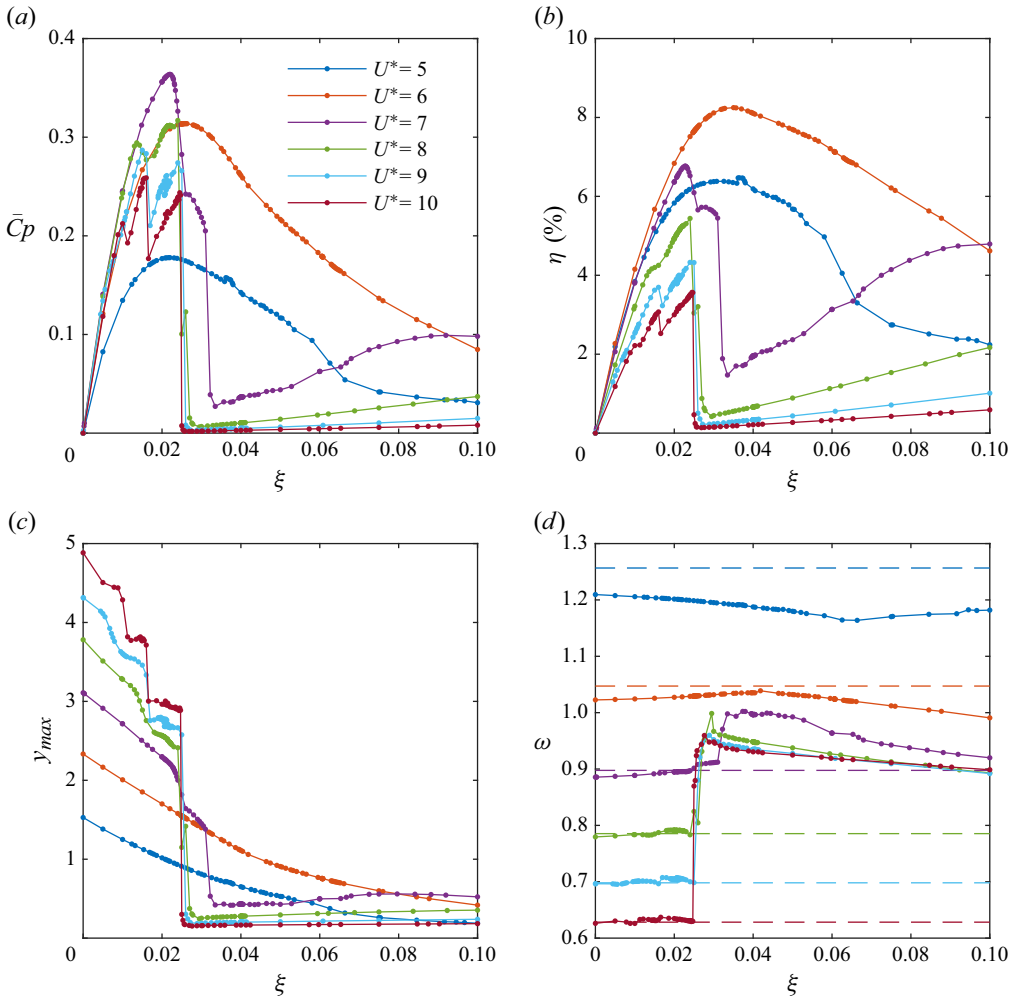


Figure 23. Variation of (a) power coefficient \bar{C}_p , (b) efficiency η , (c) oscillation amplitude y_{max} and (d) angular frequency ω , against damping coefficient ξ , for numerical simulations of the OpU9 cylinder performed at $Re = 100$ and $m^* = 4.7124$. Dashed lines in (d) indicate the undamped natural frequency $\omega_n = 2\pi/U^*$.

by vortex-induced forces provide a significant contribution to power generation, allowing more power to be extracted via mechanical damping per cycle of motion.

Figure 23 plots the variation of power coefficient \bar{C}_p , efficiency η , amplitude y_{max} and angular frequency ω against ξ , for the numerical simulations. Increasing ξ initially increases power extraction, as more power is extracted by mechanical damping. However, damping also decreases the amplitude of vibration, leading to a reduction in the power coefficient for large ξ . For $U^* \geq 7$, the galloping-like behaviour cannot be sustained for $\xi \gtrsim 0.025$, leading to a sudden decrease in amplitude and power coefficient. Figure 23(d) shows that the oscillation frequency ω (solid lines) shifts away from the natural frequency (dashed lines) when this occurs, indicating a desynchronised behaviour. This suggests that the synchronised behaviour that occurs at high U^* is due to the large oscillation amplitude, which significantly alters the vortex shedding. When the damping is sufficiently large to prevent the galloping-like behaviour, vortex shedding instead occurs at a different frequency (presumably the natural frequency of vortex shedding).

Cylinder label	$U_{opt.}^*$	AR	α	β	m_f	n_f	m_a	n_a	x_{cam}	ξ	\bar{C}_p
OpU6_P	6	19.36	-0.14	0.2782	1.356	0.6099	1.03	1.973	0.2818	0.0681	0.402
OpU8_P	8	18.56	-0.3822	0.2828	1.982	0.6445	1.486	1.742	0.2931	0.0802	0.4329
OpU10_P	10	19.67	0.3212	0.196	1.789	1.206	0.948	1.083	0.1879	0.0495	0.5130
OpUM_P	6,8,10	19.64	-0.6262	0.2915	1.538	1.099	1.871	1.079	0.3074	0.0715	0.4055

Table 6. Shape parameters for various cylinders optimised for maximum power coefficient \bar{C}_p , with fixed $U^* = U_{opt.}^*$, $Re = 100$ and $m^* = 4.7124$. The OpUM_P cylinder was optimised to maximise the average power coefficient across three different conditions ($U^* = 6, 8$ and 10).




Figure 24. Comparison between cylinders optimised for maximum power coefficient (OpU6_P, OpU8_P, OpU10_P, OpUM_P) and cylinders optimised for maximum amplitude (OpU6, OpU8, OpU10).

At high U^* , the power coefficient profiles display evidence of multiple different flow regimes (sudden jumps to different trend lines), depending on the precise value of ξ . This results in sharp peaks in the power coefficient, meaning that ξ must be precisely tuned to maximise the power coefficient. The overall maximum power coefficient, $\bar{C}_p = 0.364$ is obtained for $U^* = 7$ and $\xi = 0.0220$. At $U^* = 6$, there is a much broader peak in \bar{C}_p with a maximum value of $\bar{C}_p = 0.314$ and $\xi = 0.0259$. The broad peak means that \bar{C}_p is less sensitive to ξ , which will produce more consistent power generation.

The efficiency η (2.9) represents the ratio of power extracted from power available in the swept area of the device. This is plotted in figure 23(b), against ξ , for the OpU9 cylinder. The maximum efficiency of 8.25 % is obtained for $U^* = 6$ and $\xi = 0.0350$, and the power coefficient at this condition is $\bar{C}_p = 0.289$. The efficiency decreases as U^* is increased. At $U^* = 7$, when the power coefficient is maximised, the efficiency is $\eta = 6.71$ %. The reduction in efficiency is due to an increased swept area with increasing U^* .

6.2. Optimising the power coefficient

We now consider shape optimisation to maximise the power coefficient, using the BO-GPR method. For this case, the objective function to be maximised is the mean power coefficient ($r(\mathbf{q}) = \bar{C}_p$), and the design parameter vector \mathbf{q} contains all shape parameters and the damping coefficient ξ : $\mathbf{q} = (AR, \alpha, \beta, m_f, n_f, m_a, n_a, \xi)$. The remaining details of the implementation are the same as in § 4. For each optimisation, U^* is held constant, and the Reynolds number and dimensionless mass are $Re = 100$ and $m^* = 4.7124$, respectively. Finally, the damping coefficient is restricted to the range $0 \leq \xi \leq 0.2$.

Table 6 lists the parameters of the optimised shapes for $U^* = 6$, $U^* = 8$ and $U^* = 10$, as well as the mean power coefficient. We use the label OpUX_P to indicate shapes optimised to maximise power coefficient at $U^* = X$. The optimised shapes are plotted in figure 24, as well as the OpU6, OpU8 and OpU10 cylinders. The power-optimised cylinder have

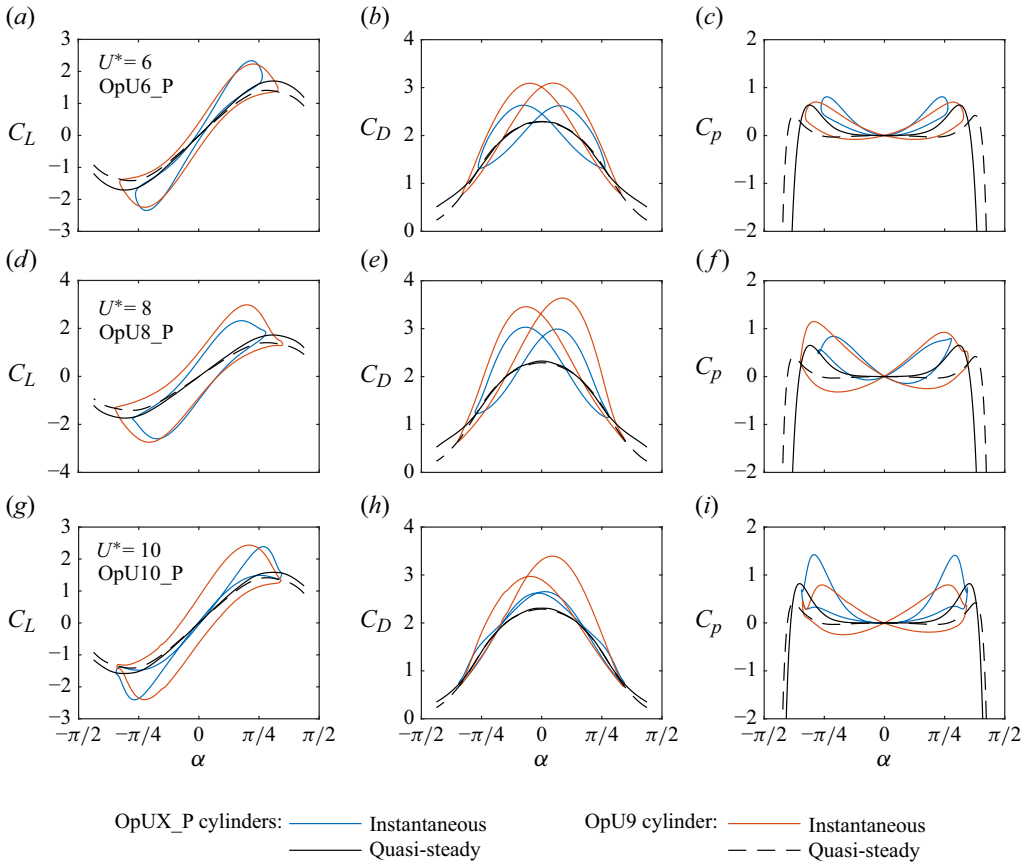


Figure 25. Variation of the instantaneous and quasi-steady lift (*a,d,g*), drag (*b,e,h*) and power (*c,f,i*) coefficients at $Re = 100$ and $m^* = 4.7124$: (*a,b,c*) OpU6_P at $U^* = 6$ and $\xi = 0.0681$; (*d,e,f*) OpU8_P at $U^* = 8$ and $\xi = 0.0802$; (*g,h,i*) OpU10_P at $U^* = 10$ and $\xi = 0.0495$. Plots of the OpU9 cylinder with $\xi = 0.0220$ are also provided for comparison.

the same general features as the amplitude-optimised cylinders; namely a thin crescent blade with sharp vertices. However, compared with the amplitude-optimised cylinders, the power-optimised cylinders have a greater curvature. The OpU6_P and OpU8_P cylinders have a mean camber of $x_{cam} \approx 0.29$, approximately three times greater than the camber of the OpU6 and OpU8 cylinders. The OpU10_P cylinder has a mean camber of $x_{cam} \approx 0.1897$, which is approximately twice the camber of the OpU10 cylinder.

Optimising at a single U^* does not guarantee that the cylinder will perform well over a range of different operating conditions. An additional optimisation was performed using the average power coefficient across three different reduced velocities ($U^* = 6, 8$ and 10) as the objective function. This cylinder, labelled OpUM_P, is plotted in figure 24, and its parameters are provided in table 6. This cylinder has a similar shape to the OpU6_P and OpU8_P cylinders, but with a slightly higher camber.

To investigate the influence of the increased camber on the power generation, figure 25 compares the instantaneous and quasi-steady lift, drag and power coefficients for the OpUX_P to the OpU10 cylinder, at $U^* = X$. Increasing camber results in an increased maximum power coefficient according to the quasi-steady approximation, due to an increased lift coefficient. However, the increased drag coefficient reduces the maximum




Figure 26. Vorticity visualisations for the optimised cylinders (a) OpU6_P at $U^* = 6$ and $\xi = 0.0681$; (b) OpU8_P at $U^* = 8$ and $\xi = 0.0802$; (c) OpU10_P at $U^* = 10$ and $\xi = 0.0495$. Vorticity visualisations for the OpU9 cylinder with $\xi = 0.0220$ at (d) $U^* = 6$, (e) $U^* = 8$ and (f) $U^* = 10$. Red indicates positive (counter-clockwise) vorticity, while blue indicates negative (clockwise) vorticity.

angle of attack for which the quasi-steady power coefficient is positive. Therefore, increasing camber reduces the maximum amplitude, but increases the power generation at smaller amplitudes.

Increasing camber may also affect the vortex-induced forces. Figure 26 plots vorticity visualisations for the power-optimised cylinders OpU6_P, OpU8_P and OpU10_P. Vorticity visualisations for the OpU9 cylinder with $\xi = 0.0220$ are also provided for comparison. Animations are provided in the online supplementary material (movies 5–7). For the OpU6_P and OpU8_P cylinders, the increased camber alters the structure of vortices as they are shed from the cylinder, compared with the OpU9 cylinder. The resulting power coefficients (figure 25c,f) are positive for all inclination angles, while the power coefficients for the OpU9 cylinder are sometimes negative. Therefore, changes to the vortex-induced forces as a result of cylinder camber also appear to enhance power generation.

For the OpU10_P cylinder, the vortex structure is similar to the OpU9 cylinder (figure 26c,f). However, figure 25(g–f) shows that the peak in power coefficient associated with vortex shedding occurs at a larger angle of attack (i.e. delayed later in the vortex

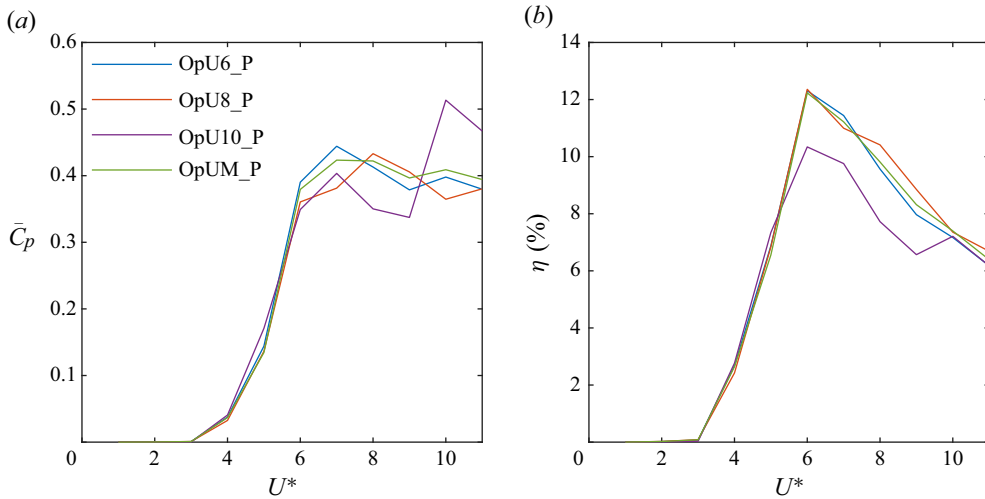


Figure 27. Variation of (a) power coefficient \bar{C}_p and (b) efficiency η against reduced velocity U^* , for various optimised cylinders, at $Re = 100$ and $m^* = 4.7124$.

shedding cycle), resulting in a much larger peak power coefficient, which enhances the power extracted from FIV.

Figure 27 plots the variation of power coefficient and efficiency against U^* for the optimised cylinders. For the OpU6_P, OpU8_P and OpU10 cylinders, there is a broad region of high power generation for reduced velocities $U^* \geq 6$, with $\bar{C}_p \approx 0.4$. Within this region, each cylinder achieves maximum power coefficient at its design U^* , with a decrease in performance away from the design point. The maximum overall power coefficient of $\bar{C}_p = 0.513$ is achieved with the OpU10 cylinder, however, this cylinder performs worse than the OpU6_P and OpU8_P cylinders between $U^* = 6$ and $U^* = 9$. Maximum efficiency of $\eta \approx 12.2\%$ is achieved for $U^* = 6$, and the efficiency decreases as U^* increases beyond this value.

In practice, devices must perform well over a range of operating conditions. The OpUM_P cylinder was optimised using the average power coefficient obtained at three different reduced velocities ($U^* = 6, 8$ and 10) as the objective function. As expected, the OpUM_P cylinder achieves a more consistent power coefficient for different reduced velocities, but has a lower maximum power coefficient than the single U^* cylinders at their design condition.

As discussed in the introduction (table 1), previous studies have considered a range of different cylinder geometries. Since these studies are at much higher Reynolds numbers than the present study, we cannot compare with their findings directly. Instead, a series of optimisations were performed for several geometries listed in table 7. For these optimisations, the free parameters to be optimised are either the damping coefficient ξ (circle, square, rectangle-triangle) or both AR and ξ (triangle, rectangle, ellipse, semi-ellipse). Of these shapes, the triangular and semi-elliptical cylinders achieve the best performance. This is consistent with prior studies (Sun *et al.* 2019b; Tamimi *et al.* 2019; Huang *et al.* 2024) which find that triangular and semicircular afterbodies generate the most power.

Figure 28 compares the power coefficient and efficiency obtained using the circular, triangular and semicircular cylinders with the OpU10_P and OpUM_P cylinders. The optimised cylinders (OpU10_P and OpUM_P) achieve a substantially greater power

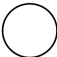






			\bar{C}_P	U^*	ξ	AR
C5	Circle		8.24×10^{-2}	5	8.95×10^{-2}	1
S5	Square		1.51×10^{-2}	5	5.73×10^{-2}	1
T7	Triangle		0.254	7	0.193	1.35
R5	Rectangle		8.21×10^{-2}	5	3.17×10^{-2}	5.97
E5	Ellipse		0.149	5	3.80×10^{-2}	2.40
SE6	Semi-ellipse		0.264	6	2.93×10^{-2}	4.95
RT7	Rectangle-triangle		0.147	7	0.100	1

Table 7. Optimised parameters ξ and AR for various elementary cylinder geometries at $Re = 100$ and $m^* = 4.7124$. A sketch of each shape is provided, with the flow direction from left to right.

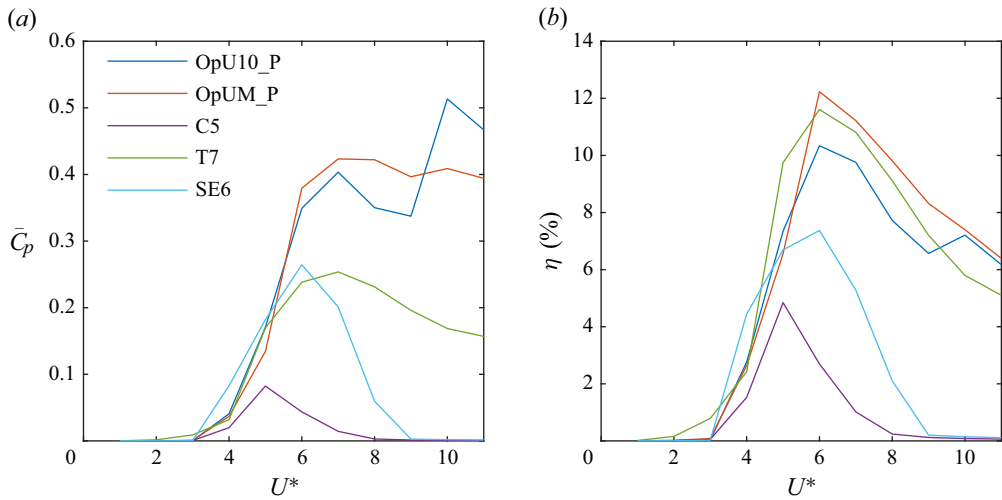


Figure 28. Variation of (a) power coefficient \bar{C}_p and (b) efficiency η against reduced velocity U^* , for various cylinders at $Re = 100$ and $m^* = 4.7124$.

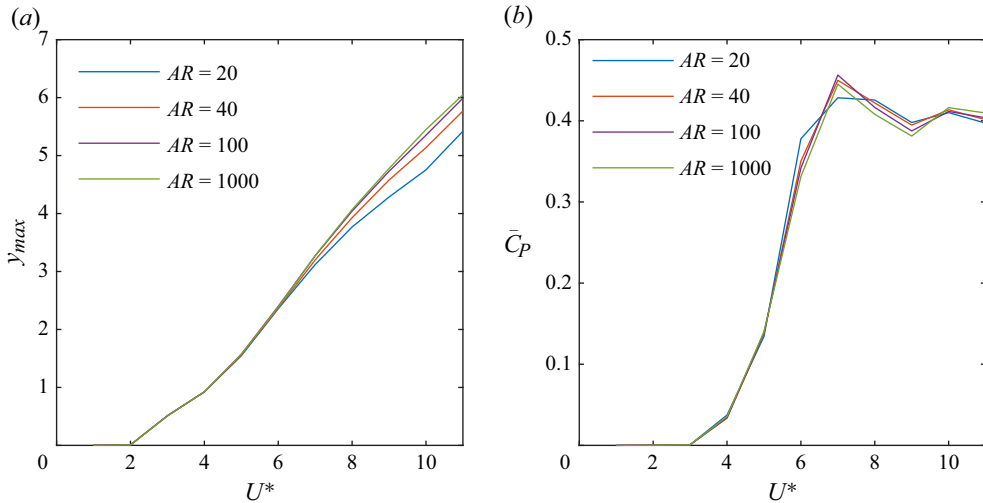


Figure 29. (a) Variation of oscillation amplitude (y_{max}) against U^* for modified OpU7 cylinders with various aspect ratios. (b) Variation of mean power coefficient \bar{C}_p against U^* for modified OpUM_P cylinders with various aspect ratios.

coefficient over a much larger range of U^* compared with the circular, triangular and semi-elliptical cylinders. The maximum power coefficient $\bar{C}_p = 0.513$ is obtained using the OpU10_P cylinder at $U^* = 10$. This is 6.23 times larger than the maximum power coefficient obtained using the circular cylinder ($\bar{C}_p = 0.0824$), and 1.94 times greater than that obtained using the semi-elliptical cylinder ($\bar{C}_p = 0.264$). The OpUM_P cylinder has an average power coefficient of $\bar{C}_p = 0.404$ over the range $6 \leq U^* \leq 11$, which is 1.95 times the average power coefficient of the triangular cylinder over this same range ($\bar{C}_p = 0.208$). The efficiency is maximised at $U^* = 6$, with a maximum efficiency of $\eta = 12.2\%$ obtained using the OpUM_P cylinder, which is 2.52 times the efficiency of the circular cylinder ($\eta = 4.84\%$), but only slightly greater than the efficiency of the triangular cylinder ($\eta = 11.6\%$).

Note that we have not optimised to maximise the efficiency. Despite this, the OpU6_P, OpU8_P and OpUM_P cylinders achieve a better efficiency than other cylinders. For future studies, it may be worth performing a combined optimisation to balance maximising both power coefficient and efficiency.

6.3. Effect of maximum aspect ratio

Given that the optimised cylinders have aspect ratio close to the upper bound of $AR = 20$, this suggests that cylinders with a higher AR may achieve even greater performance. Figure 29(a) presents the amplitude response for cylinders with parameters equal to those of the OpU7 cylinder, but with various aspect ratios. This confirms that the oscillation amplitude increases as AR increases, with a zero-thickness blade being the ideal cylinder to maximise the amplitude response. In practice the maximum aspect ratio will be limited by engineering constraints, such as the strength and stiffness of the cylinder.

Figure 29(b) also plots the power coefficient for cylinders of different aspect ratios, with remaining parameters equal to the OpUM_P cylinder. While the maximum power coefficient increases with increasing AR , for some reduced velocities the power coefficient decreases with increasing AR . Given the optimised cylinders all have an AR close to the

upper bound of 20, it is likely that a zero-thickness blade will achieve the best overall power generation, however, this needs further investigation.

7. Conclusions

Bayesian optimisation with GP regression was performed to optimise the shape of an elastically mounted cylinder undergoing transverse FIV. First, optimisation was performed to maximise the amplitude of undamped oscillations, for several different reduced velocities (U^*) at $Re = 100$ and $m^* = 4.7124$. The optimised shape was found to be a thin curved blade with sharp leading and trailing edges, similar to the curved blade geometry studied by Tucker Harvey *et al.* (2019, 2020).

The optimised cylinder exhibits a combined VIV–galloping response. A quasi-steady analysis reveals that hard galloping is possible – the cylinder cannot gallop from rest, but galloping may occur given a sufficiently large initial displacement. At large relative inclination angles, the curved blade acts like an airfoil, with a relatively low drag coefficient, and a large lift coefficient. For certain angles of attack, the component of lift in the transverse direction exceeds the drag component, producing a net positive contribution to the cylinder motion. The FIV is further enhanced by vortex-induced forces. A single primary vortex is shed from the cylinder each half-cycle of motion, resulting in a large transverse force in line with the cylinder’s velocity.

Finally, shape optimisation was performed to enhance the power coefficient for damped vibrations. These optimised cylinders were also found to be curved blades with sharp leading and trailing edges, however, with a larger curvature than the cylinders optimised for maximum amplitude. The increased camber allows more power to be extracted from the cylinder during a single cycle of motion, compared with the amplitude-optimised cylinder.

The shape optimisation produced remarkable improvements in both power coefficient and efficiency. The maximum power coefficient was $\bar{C}_p = 0.513$, obtained using the OpU10_P cylinder at $U^* = 10$. This is an increase of 523 % compared with the circular cylinder ($\bar{C}_p = 0.0824$) and of 94 % compared with the semi-elliptical cylinder ($\bar{C}_p = 0.264$). The OpUM_P achieved a relatively consistent power coefficient of $\bar{C}_p = 0.404$ over the range $6 \leq U^* \leq 11$, which is an improvement of 95 % compared with the triangular cylinder ($\bar{C}_p = 0.208$) over this same range. The maximum efficiency of $\eta = 12.2$ % was obtained using the OpUM_P cylinder at $U^* = 6$, which is an improvement of 152 % compared with a circular cylinder ($\eta = 4.84$ %), but only 5.2 % better than the triangular cylinder ($\eta = 11.6$ %).

The present study was limited to low Reynolds numbers ($Re = 100$), and considered only numerical simulations. For future work, the optimisation should be performed at larger Reynolds numbers, using an appropriate turbulence model validated against experimental measurements. The performance of the optimised shape should then be validated experimentally. Increasing the Reynolds number is expected to yield even greater power coefficients and efficiencies, due to reduced viscous damping. For comparison, Lo *et al.* (2024) obtained a maximum power coefficient of $\bar{C}_p = 3.31$ and maximum efficiency of $\eta = 22.7$ % for a thin elliptical cylinder with Reynolds numbers in the range $990 \leq Re \leq 4,390$, which are substantially greater than the efficiency and power coefficient obtained in the present study.

Finally, given the present study obtained a thin curved blade as the optimised shape, future optimisation studies may focus on shape parameters for a thin blade, rather than attempting to parametrise a wide range of different geometries as was done in the present study. Geometry parametrisation techniques developed for airfoil and compressor-blade sections (Kulfan 2010; Shi 2021) will be useful for this purpose.

Supplementary movies. Supplementary movies are available at <https://doi.org/10.1017/jfm.2025.10901>.

Funding. This work was supported by the Australian Government through the Australian Research Council's Discovery Projects funding scheme (projects DP200100704, DP210100990 and DP190103388), and by computational resources provided by the National Computational Infrastructure and Pawsey Supercomputer Centre (Merit Grant d71) under the National Computational Merit Allocation Scheme.

Declaration of interests. The authors report no conflict of interest.

REFERENCES

- ALHUSSEIN, H., DALAQ, A.S. & DAQAQ, M. 2024 AI-based shape optimization of galloping micro-power generators: exploring the benefits of curved surfaces. *Sci. Rep.-UK* **14** (1), 1552.
- BANGGA, G., LUTZ, T. & ARNOLD, M. 2020 An improved second order dynamic stall model for wind turbine airfoils. *Wind Energy Sci. Discuss.* **2020**, 1–36.
- BARRERO-GIL, A., ALONSO, G. & SANZ-ANDRES, A. 2010 Energy harvesting from transverse galloping. *J. Sound Vib.* **329** (14), 2873–2883.
- BEGALL, M.J., SCHWEIDTMANN, A.M., MHAMDI, A. & MITSOS, A. 2023 Geometry optimization of a continuous millireactor via CFD and Bayesian optimization. *Comput. Chem. Engng* **171**, 108140.
- BENNER, B.M., CARLSON, D.W., SEYED-AGHAZADEH, B. & MODARRES-SADEGHI, Y. 2019 Vortex-induced vibration of symmetric airfoils used in vertical-axis wind turbines. *J. Fluids Struct.* **91**, 102577.
- BERNITSAS, M.M., RAGHAVAN, K., BEN-SIMON, Y. & GARCIA, E.M.H. 2008 VIVACE (Vortex Induced Vibration Aquatic Clean Energy): a new concept in generation of clean and renewable energy from fluid flow. *J. Offshore Mech. Arctic Engng* **130** (4), 041101.
- BOUCLIN, D.N. 1979 Hydroelastic oscillations of square cylinders. PhD thesis, University of British Columbia, Canada.
- CHEN, W., LI, X. & YANG, W. 2024 Shape optimization to enhance energy harvesting from vortex-induced vibration of a circular cylinder. *Phys. Fluids* **36**, 023612.
- CHEN, W. & LI, Y. 2023 Evidence and physical mechanism for vortex-induced vibration of a bluff body without an afterbody. *Phys. Fluids* **35**, 065143.
- CHEN, W. & LI, Y. 2024 Energy harvesting performance of an elastically mounted semi-circular cylinder. *Renew. Energy* **229**, 120692.
- DIESSNER, M., O'CONNOR, J., WYNN, A., LAIZET, S., GUAN, Y., WILSON, K. & WHALLEY, R.D. 2022 Investigating Bayesian optimization for expensive-to-evaluate black box functions: application in fluid dynamics. *Frontiers Appl. Maths Stat.* **8**, 1076296.
- DING, L., ZHANG, L., WU, C., MAO, X. & JIANG, D. 2015 Flow induced motion and energy harvesting of bluff bodies with different cross sections. *Energy Convers. Manage.* **91**, 416–426.
- GEUZAIN, C., REMACLE, J.-F. & DULAR, P. 2009 Gmsh: a three-dimensional finite element mesh generator. *Int'l J. Numer. Meth. Engng* **79** (11), 1309–1331.
- GONZALEZ, J., DAI, Z., HENNIG, P. & LAWRENCE, N. 2016 Batch Bayesian optimization via local penalization. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics* (ed. A. Gretton & C. C. Robert), Proceedings of Machine Learning Research, vol. 51, pp. 648–657. PMLR.
- GUPTA, S., LO, J.C.C., ZHAO, J., THOMPSON, M.C. & HOURIGAN, K. 2025 The role of an afterbody in flow-induced vibration of cylinders at low to moderate Reynolds numbers. *J. Fluid Mech.* **1007**, A6.
- HE, X., YANG, X. & JIANG, S. 2018 Enhancement of wind energy harvesting by interaction between vortex-induced vibration and galloping. *Appl. Phys. Lett.* **112**, 033901.
- HUANG, J., LUO, W., ZHOU, M., WANG, B., XU, Z. & QIN, H. 2024 Influences of bluff body geometries upon performances of piezoelectric wind energy harvesters. *Phys. Fluids* **36**, 097117.
- JAVED, U., ABDELKEFI, A. & AKHTAR, I. 2016 An improved stability characterization for aeroelastic energy harvesting applications. *Commun. Nonlinear Sci. Numer. Simul.* **36**, 252–265.
- KULFAN, B.M. 2010 Recent extensions and applications of the 'CST' universal parametric geometry representation method. *Aeronaut. J.* **114** (1153), 157–176.
- KUMAR, A., PATIL, S., KOVACEVIC, A. & PONNUSAMI, S.A. 2024 Performance prediction and Bayesian optimization of screw compressors using Gaussian process regression. *Engng Applies Artif. Intell.* **133**, 108270.
- LEONTINI, J.S., THOMPSON, M.C. & HOURIGAN, K. 2006 The beginning of branching behaviour of vortex-induced vibration during two-dimensional flow. *J. Fluids Struct.* **22** (6–7), 857–864.
- LIAN, J., YAN, X., LIU, F. & ZHANG, J. 2017 Analysis on flow induced motion of cylinders with different cross sections and the potential capacity of energy transference from the flow. *Shock Vib.* **2017** (1), 4356367.

- LO, J.C.C., HOURIGAN, K., THOMPSON, M.C. & ZHAO, J. 2023 The effect of structural damping on flow-induced vibration of a thin elliptical cylinder. *J. Fluid Mech.* **974**, A5.
- LO, J.C.C., THOMPSON, M.C., HOURIGAN, K. & ZHAO, J. 2024 Order of magnitude increase in power from flow-induced vibrations. *Renew. Sustainable Energy Rev.* **205**, 114843.
- LV, Y., SUN, L., BERNITSAS, M.M. & SUN, H. 2021 A comprehensive review of nonlinear oscillators in hydrokinetic energy harnessing using flow-induced vibrations. *Renew. Sustainable Energy Rev.* **150**, 111388.
- MAHFOZE, O.A., MOODY, A., WYNN, A., WHALLEY, R.D. & LAIZET, S. 2019 Reducing the skin-friction drag of a turbulent boundary-layer flow with low-amplitude wall-normal blowing within a Bayesian optimization framework. *Phys. Rev. Fluids* **4** (9), 094601.
- MALLOR, F., SEMPRINI-CESARI, G., MUKHA, T., REZAEIRAVESH, S. & SCHLATTER, P. 2024 Bayesian optimization of wall-normal blowing and suction-based flow control of a NACA 4412 wing profile. *Flow Turbul. Combust.* **113** (1), 93–118.
- MANNINI, C., MARRA, A.M. & BARTOLI, G. 2014 VIV–galloping instability of rectangular cylinders: review and new experiments. *J. Wind Engng Ind. Aerodyn.* **132**, 109–124.
- MOHAMED, O.S., MELANI, P.F., BANGGA, G., ARYAN, N., GRECO, L. & BIANCHINI, A. 2024 Accuracy assessment of beddoes-leishman and iag dynamic stall models for wind turbine applications. *J. Phys.: Conf. Ser.* **2767**, 052053.
- MORITA, Y., REZAEIRAVESH, S., TABATABAEI, N., VINUESA, R., FUKAGATA, K. & SCHLATTER, P. 2022 Applying Bayesian optimization with Gaussian process regression to computational fluid dynamics problems. *J. Comput. Phys.* **449**, 110788.
- NABAE, Y. & FUKAGATA, K. 2021 Bayesian optimization of traveling wave-like wall deformation for friction drag reduction in turbulent channel flow. *J. Fluid Sci. Technol.* **16** (4), JFST0024.
- NAJJAR, F.M. & VANKA, S.P. 1995 Simulations of the unsteady separated flow past a normal flat plate. *Intl J. Numer. Meth. Fluids* **21** (7), 525–547.
- NOVAK, M. & TANAKA, H. 1974 Effect of turbulence on galloping instability. *J. Engng Mech. Div.* **100** (1), 27–47.
- PARKINSON, G. 1989 Phenomena and modelling of flow-induced vibrations of bluff bodies. *Prog. Aerosp. Sci.* **26** (2), 169–224.
- PARKINSON, G.V. & SMITH, J.D. 1964 The square prism as an aeroelastic non-linear oscillator. *Q. J. Mech. Appl. Maths* **17** (2), 225–239.
- RAJAMUNI, M.M., THOMPSON, M.C. & HOURIGAN, K. 2020 Efficient FSI solvers for multiple-degrees-of-freedom flow-induced vibration of a rigid body. *Comput. Fluids* **196**, 104340.
- SARPKAYA, T. 2004 A critical review of the intrinsic nature of vortex-induced vibrations. *J. Fluids Struct.* **19** (4), 389–447.
- SHI, H. 2021 A parametric blade design method for high-speed axial compressor. *Aerospace* **8** (9), 271.
- SIROHI, J. & MAHADIK, R. 2011 Piezoelectric wind energy harvester for low-power sensors. *J. Intell. Mater. Syst. Struct.* **22** (18), 2215–2228.
- SUN, W., JO, S. & SEOK, J. 2019a Development of the optimal bluff body for wind energy harvesting using the synergetic effect of coupled vortex induced vibration and galloping phenomena. *Intl J. Mech. Sci.* **156**, 435–445.
- SUN, W., WANG, Y., LIU, Y., SU, B., GUO, T., CHENG, G., ZHANG, Z., DING, J. & SEOK, J. 2024 Navigating the future of flow-induced vibration-based piezoelectric energy harvesting. *Renew. Sustainable Energy Rev.* **201**, 114624.
- SUN, W., ZHAO, D., TAN, T., YAN, Z., GUO, P. & LUO, X. 2019b Low velocity water flow energy harvesting using vortex induced vibration and galloping. *Appl. Energy* **251**, 113392.
- TALNIKAR, C., BLONIGAN, P., BODART, J. & WANG, Q. 2014 Parallel optimization for LES. In Proceedings of the Summer Program, pp. 315–324. Center for Turbulence Research.
- TAMIMI, V., ARMIN, M., SHAHVAGHAR-ASL, S., NAEENI, S.T.O. & ZEINODDINI, M. 2019 Fiv energy harvesting from sharp-edge oscillators. In *International conference on offshore mechanics and arctic engineering*, vol. 58899, pp. V010T09A001. American Society of Mechanical Engineers.
- TUCKER HARVEY, S., KHOVANOV, I.A. & DENISSENKO, P. 2019 A galloping energy harvester with flow attachment. *Appl. Phys. Lett.* **114** (10), 104103.
- TUCKER HARVEY, S., KHOVANOV, I.A., MURAI, Y. & DENISSENKO, P. 2020 Characterisation of aeroelastic harvester efficiency by measuring transient growth of oscillations. *Appl. Energy* **268**, 115014.
- WANG, J., GENG, L., DING, L., ZHU, H. & YURCHENKO, D. 2020a The state-of-the-art review on energy harvesting from flow-induced vibrations. *Appl. Energy* **267**, 114902.
- WANG, J.-S., FAN, D. & LIN, K. 2020b A review on flow-induced vibration of offshore circular cylinders. *J. Hydrodyn.* **32** (3), 415–440.

- WELLER, H.G., TABOR, G., JASAK, H. & FUREBY, C. 1998 A tensorial approach to computational continuum mechanics using object-oriented techniques. *Comput. Phys.* **12** (6), 620–631.
- YANG, Y., ZHAO, L. & TANG, L. 2013 Comparative study of tip cross-sections for efficient galloping energy harvesting. *Appl. Phys. Lett.* **102**, 064105.
- ZHANG, B., SONG, B., MAO, Z., TIAN, W. & LI, B. 2017 Numerical investigation on VIV energy harvesting of bluff bodies with different cross sections in tandem arrangement. *Energy* **133**, 723–736.
- ZHANG, B., MAO, Z., SONG, B., DING, W. & TIAN, W. 2018*a* Numerical investigation on effect of damping-ratio and mass-ratio on energy harnessing of a square cylinder in fim. *Energy* **144**, 218–231.
- ZHANG, B., WANG, K.-H., SONG, B., MAO, Z. & TIAN, W. 2018*b* Numerical investigation on the effect of the cross-sectional aspect ratio of a rectangular cylinder in FIM on hydrokinetic energy conversion. *Energy* **165**, 949–964.
- ZHANG, J., LIU, F., LIAN, J., YAN, X. & REN, Q. 2016 Flow induced vibration and energy extraction of an equilateral triangle prism at different system damping ratios. *Energies* **9** (11), 938.
- ZHAO, J., HOURIGAN, K. & THOMPSON, M.C. 2018 Flow-induced vibration of D-section cylinders: an afterbody is not essential for vortex-induced vibration. *J. Fluid Mech.* **851**, 317–343.