

# Bayesian minimisation of energy consumption in turbulent pipe flow via unsteady driving

Felix Kranz<sup>1</sup> , Daniel Morón Montesdeoca<sup>1</sup>  and Marc Avila<sup>1,2</sup> 

<sup>1</sup>Center of Applied Space Technology and Microgravity (ZARM), University of Bremen, Am Fallturm 2, Bremen 28359, Germany

<sup>2</sup>MAPEX Center for Materials and Processes, University of Bremen, Am Biologischen Garten 2, Bremen 28359, Germany

**Corresponding author:** Felix Kranz, [felix.kranz@zarm.uni-bremen.de](mailto:felix.kranz@zarm.uni-bremen.de)

(Received 16 May 2025; revised 19 August 2025; accepted 18 September 2025)

---

Turbulence accounts for most of the energy losses associated with the pumping of fluids in pipes. Pulsatile drivings can reduce the drag and energy consumption required to supply a desired mass flux, when compared with steady driving. However, not all pulsation waveforms yield reductions. Here, we compute drag- and energy-optimal driving waveforms using direct numerical simulations and a gradient-free black-box optimisation framework. Specifically, we show that Bayesian optimisation is vastly superior to ordinary gradient-based methods in terms of computational efficiency and robustness, due to its ability to deal with noisy objective functions, as they naturally arise from the finite-time averaging of turbulent flows. We identify optimal waveforms for three Reynolds numbers and two Womersley numbers. At a Reynolds number of 8600 and a Womersley number of 10, optimal waveforms reduce total energy consumption by 22 % and drag by 37 %. These reductions are rooted in the suppression of turbulence prior to the acceleration phase, the resulting delay in turbulence onset, and the radial localisation of turbulent kinetic energy and production towards the pipe centre. Our results pinpoint that the predominant, steady operation mode of pumping fluids through pipes is far from optimal.

**Key words:** turbulence control, pipe flow

---

## 1. Introduction

Turbulent pipe flows are ubiquitous in engineering, ranging from large-scale oil or gas pipelines to small-scale applications like heating pipes or fresh water supply. Pumping systems are recognised as major energy consumers globally, contributing a substantial



share to the overall electrical energy demand (McKeon 2010). In certain industrial plant operations, pumping systems even account for up to 50 % of the energy usage (Frenning 2001). Compared with laminar conditions, in turbulent flows, the multi-scale eddying motion is responsible for a major part of the high friction levels and ultimately the high pumping costs (Blasius 1913). Hence, many efforts have been devoted to devising control strategies that can reduce drag or suppress turbulence. Successful examples go from passive strategies – such as the use of riblets (García-Mayoral & Jiménez 2011) – to active strategies including the application of body forces to dampen turbulent fluctuations (Kühnen *et al.* 2018; Marensi, Willis & Kerswell 2019), as well as wall-normal blowing (Mahfoze *et al.* 2019) and suction (Mallor *et al.* 2023) in turbulent boundary-layer flows. A promising active strategy involves the use of oscillatory forces to reduce wall friction, such as spanwise oscillations close to or at the wall (Quadrio & Ricco 2004; Auteri *et al.* 2010), or the use of unsteady pressure gradients to drive the flow. Both numerical (e.g. Iwamoto, Sasou & Kawamura 2007; Foggi *et al.* 2023a,b; Scarselli *et al.* 2023) and experimental (e.g. Souma, Iwamoto & Murata 2009; Kobayashi *et al.* 2021; Scarselli *et al.* 2023) studies have shown that certain unsteady (pulsatile) drivings can lead to significant drag reductions and energy savings. However, the space of possible flow parameters and pulsation waveforms capable of achieving such improvements is vast and remains only partially explored.

Pulsatile flow is governed by three factors. First, the (time-averaged) Reynolds number  $\overline{Re} = \overline{U}D/\nu$ , where  $\overline{U}$  is the time-averaged bulk velocity,  $D$  is a characteristic length and  $\nu$  is the kinematic viscosity. Second, the angular frequency  $\omega$  of the pulsation, defined using the dimensionless Womersley number  $Wo = (D/2)\sqrt{\omega/\nu}$  and third, the driving bulk velocity waveform (WF) described by the Fourier expansion

$$U(t) = \overline{U} + \sum_{k=1}^{\infty} a_k \cos(\omega kt) + \sum_{k=1}^{\infty} b_k \sin(\omega kt), \quad (1.1)$$

where  $a_k$ ,  $b_k$  are the Fourier coefficients of the pulsation.

There are two main ways to impose a pulsatile flow: either by prescribing a time-dependent streamwise pressure gradient or by imposing the desired bulk velocity. Most studies adopt the former approach, modulating the streamwise pressure gradient in time. In channel flow, Iwamoto *et al.* (2007) numerically investigated periodic square-wave pressure gradients, cyclically alternating between positive and negative values. By carefully choosing the pulsation frequency, they found that the cycle-averaged skin friction can be reduced compared with the corresponding steady flow (the one that produces the same time-averaged flow rate). Inspired by the latter, Foggi *et al.* (2023a,b) employed a simple temporal waveform for the pressure gradient, consisting of a periodic on–off pumping sequence. Using direct numerical simulation (DNS), at  $\overline{Re} \approx 4600$ , the best-performing waveform showed energy savings of up to 17 % when accounted for the achieved flow rate. Savings were obtained when the flow spent a significant fraction of the period in a transient quasi-laminar state, where there was no axial pressure gradient and the flow rate slowly decayed.

Souma *et al.* (2009) substantiated the numerical findings of Iwamoto *et al.* (2007) by doing corresponding experiments in pipe flow. More recently, Kobayashi *et al.* (2021) did a similar experimental study at low Reynolds numbers ( $\overline{Re} \approx 3600$ ). They automatically generated more than 7000 different driving pressure waveforms, confirming the reduction of cycle-averaged drag. However, drag reduction is not satisfactory to achieve a reduction of net energy consumption. For example, at a mean Reynolds number of  $\overline{Re} \approx 5900$  and high frequency regimes ( $Wo \in [39, 53]$ ), Manna, Vacca & Verzicco (2015) numerically

showed that the mean turbulent friction can be reduced by harmonically modulating the driving pressure in time. However, their method required extra energy when compared with steady driving and led to a decrease of the pumping efficiency. A more complex control strategy was investigated by Ding *et al.* (2024) in pipe flow experiments. They found that drag can be significantly reduced over a wide range of frequencies, amplitudes and Reynolds numbers by oscillating the pipe’s surface azimuthally.

Inspired by the pulsatile rhythm of the mammalian heart, Scarselli *et al.* (2023) proposed to modulate the bulk velocity – rather than adjusting the pressure gradient – with a cardiac-inspired (triangular) waveform. They found that this approach could suppress key turbulent features. For instance, in their best-performing case, at  $\overline{Re} = 8600$  and  $Wo = 14$ , they achieved drag reductions of up to 27% and energy savings of up to 9%. The authors attributed the reductions in drag to turbulence being initially ‘frozen’ during acceleration phases, which minimised the impact of mean flow velocity variations on turbulent stresses. Furthermore, the reduction of wall shear stress was linked to the delayed response of turbulence to the varying pressure gradient by Liu *et al.* (2024). In this paper, we follow up on the work of Scarselli *et al.* (2023) and aim to devise periodic (bulk velocity) waveforms that either minimise drag or energy consumption for different Reynolds numbers  $\overline{Re} \in \{4300, 5160, 8600\}$  and for Womersley numbers of  $Wo \in \{10, 10\sqrt{2}\}$ . In contrast to Scarselli *et al.* (2023) and Foggi *et al.* (2023a,b), we use optimisation algorithms to minimise an objective functional  $\mathcal{J}$  (the turbulent drag or energy consumption) while delivering a desired mean bulk  $\overline{U}_d$ , that is,

$$\min_{U(t)} \mathcal{J}(U(t)) \quad \text{subject to} \quad \overline{U} = \overline{U}_d. \quad (1.2)$$

Since there is no analytical framework to understand how different waveforms influence the objective function in (1.2) (i.e.  $\mathcal{J}$  is a black-box), evaluating  $\mathcal{J}$  requires DNS. Tackling the optimisation problem (1.2) by ordinary gradient-based algorithms demands for accurate computations of  $\mathcal{J}(U(t))$ , generally leading to long simulations of dozens of periods. For example, Scarselli *et al.* (2023) obtained cycle-average values of  $\mathcal{J}$  by considering up to 14 periods, while Foggi *et al.* (2023a,b) averaged over up to 36 periods. Therefore, to optimise  $\mathcal{J}$ , methods are needed that are robust to the statistical error arising from the averaging of turbulent flows.

The rest of the paper is structured as follows. In § 2, we present the numerical model, the DNS set-up and three finite-dimensional simplifications of (1.2) that, by design, fulfil the  $(\overline{U} = \overline{U}_d)$ -constraint from (1.2). In § 3, we describe the Bayesian optimisation method we use to find optimal waveforms (minimisers of (1.2)) and show it is vastly superior to gradient-based methods. In § 4, we present and discuss optimal driving waveforms and in § 5, we draw some conclusions.

## 2. Methods

### 2.1. Governing equations

We considered the flow of a viscous Newtonian fluid with constant properties in a straight smooth rigid pipe of circular cross-section of diameter  $D$  and, if not stated otherwise, length  $L = 5D$ . The flow is assumed to be incompressible and governed by the dimensionless Navier–Stokes equations

$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} - \frac{1}{Re} \Delta \mathbf{u} + \nabla p = \boldsymbol{\Pi}_U(t) + \mathbf{f}_b, \quad \nabla \cdot \mathbf{u} = 0, \quad (2.1)$$

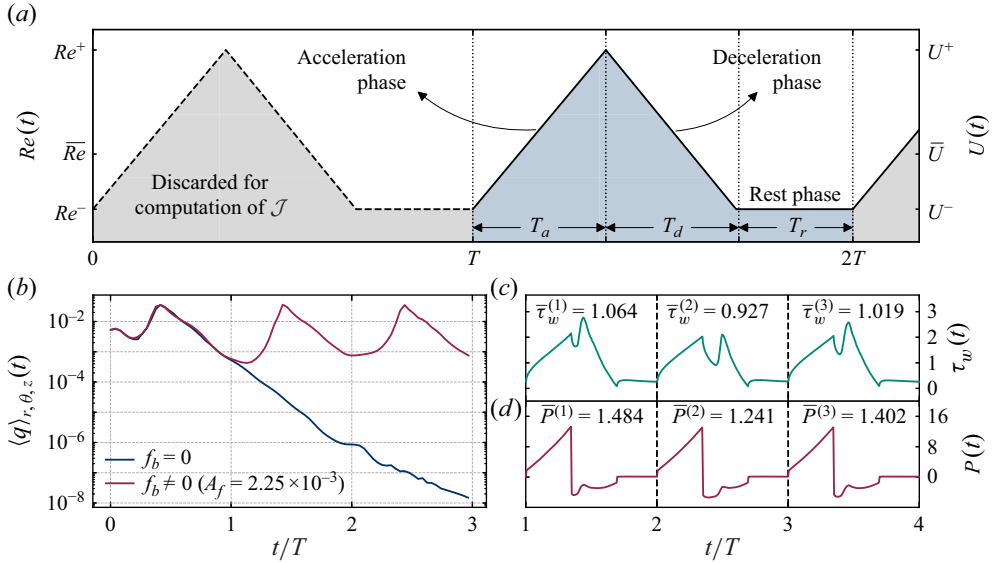


Figure 1. (a) Schematic description of the considered triangular waveforms in terms of the time variant Reynolds number  $Re(t)$  or bulk velocity  $U(t)$  (right-hand side labels). (b) Evolution of the volume-integrated cross-stream turbulent kinetic energy (in units of  $\bar{U}^2$ ) over three periods of a run driven according to panel (a) where  $\bar{Re} = 4300$ ,  $Re^+ = 9400$ ,  $Re^- = 1600$  and  $T_a = 0.345T$ . (c) Wall shear stress  $\tau_w(t)$  and power input  $P(t)$  over the last three periods of a four period run driven in the same manner as in panel (b) where  $A_f = 2.25 \times 10^{-3}$ . The wall shear stress is normalised with respect to the steady wall shear stress obtained by the Blasius friction and the power input accordingly (i.e. units of  $\bar{\tau}_{w,b}$  and  $\bar{P}_b$ , respectively).

where  $\mathbf{u} = (u_r, u_\theta, u_z)$  denotes the velocity field in a cylindrical coordinate system  $(r, \theta, z)$ ,  $p$  is the pressure,  $\mathbf{\Pi}_U(t)$  is the time-dependent axial driving force that realises a given waveform  $U(t)$  and  $\mathbf{f}_b$  is a volumetric body force (see § 2.2). All variables are rendered dimensionless using  $D$ , the time-averaged bulk velocity  $\bar{U}$  and the fluid's density  $\rho_f$ . The dimensionless period length is given by  $T = 2\pi \bar{Re} / Wo^2$ . We employ no-slip boundary conditions at the wall, and periodic boundary conditions in the axial and azimuthal direction. For the remainder of this paper,  $\overline{(\cdot)}$  indicates temporal averages, while  $\langle \cdot \rangle_\beta$  denotes spatial averaging with respect to the direction(s)  $\beta$ .

## 2.2. Body force to trigger turbulence transition

Scarselli *et al.* (2023) considered waveforms consisting of a constant acceleration, followed by a constant deceleration and a constant low-velocity phase. Such a waveform, used as a baseline for our calculations, is shown in figure 1(a), where  $Re(t) = U(t)D/\nu$ ,  $Re^+ = \max_t Re(t) = 9400$ ,  $Re^- = \min_t Re(t) = 1600$  and  $\bar{Re} = 4300$ . In figure 1(b), the evolution of the volume-integrated cross-stream turbulent kinetic energy,

$$\langle q \rangle_{r,\theta,z}(t) = \frac{4}{\pi D^2 L} \int_0^L \int_0^{2\pi} \int_0^{D/2} (u_r^2 + u_\theta^2) r \, dr \, d\theta \, dz, \quad (2.2)$$

is shown as a blue line for a direct numerical simulation initialised with a turbulent flow state (for detailed information about the simulation set-up, see § 2.3). During the low-velocity phase,  $\langle q \rangle_{r,\theta,z}(t)$  decays exponentially and the flow laminarises irreversibly. This relaminarisation differs from the experimental observations reported by Scarselli *et al.* (2023), where the flow cyclically transitioned to turbulence in each period. In numerical simulations, such a cyclic transition is not observed unless residual turbulence or external



	$\overline{Re}$	$A_f$	$N_r$	$N_\theta$	$N_z$	$Re_\tau$	$\Delta r_-^+$	$\Delta r_+^+$	$\Delta(R\theta)^+$	$\Delta z^+$	$\Delta t_-$	$\Delta t_+$
Optimisation	4300	$2.25 \times 10^{-3}$	80	192	270	277	0.09	5.1	9.1	10.2	$1.3 \times 10^{-3}$	$1.6 \times 10^{-2}$
iterations	5160	$1.15 \times 10^{-3}$	96	192	300	258	0.05	4.0	8.5	8.6	$1.2 \times 10^{-3}$	$1.4 \times 10^{-2}$
	8600	$0.5 \times 10^{-3}$	128	246	420	431	0.06	5.0	11.0	11.0	$0.7 \times 10^{-3}$	$1.5 \times 10^{-2}$
WF 1 (coarse)	5160	$1.15 \times 10^{-3}$	96	192	300	215	0.05	3.3	7.0	7.1	$1.3 \times 10^{-3}$	$1.3 \times 10^{-2}$
WF 1 (fine)	5160	$1.15 \times 10^{-3}$	120	436	390	216	0.03	2.6	3.1	5.5	$1.6 \times 10^{-4}$	$1.4 \times 10^{-3}$
WF 2 (coarse)	5160	$1.15 \times 10^{-3}$	96	192	300	222	0.05	3.4	7.3	7.4	$1.9 \times 10^{-3}$	$1.2 \times 10^{-2}$
WF 2 (fine)	5160	$1.15 \times 10^{-3}$	120	436	390	222	0.03	2.7	3.2	5.7	$2.4 \times 10^{-4}$	$1.3 \times 10^{-3}$
WF 3 (coarse)	5160	$1.15 \times 10^{-3}$	96	192	300	225	0.05	3.4	7.4	7.5	$1.3 \times 10^{-3}$	$1.5 \times 10^{-2}$
WF 3 (fine)	5160	$1.15 \times 10^{-3}$	120	436	390	232	0.03	2.9	3.3	5.9	$1.2 \times 10^{-4}$	$2.8 \times 10^{-3}$
WF 4 (coarse)	8600	$0.5 \times 10^{-3}$	128	246	420	375	0.05	4.3	9.4	9.5	$7.3 \times 10^{-4}$	$1.3 \times 10^{-2}$
WF 4 (fine)	8660	$0.5 \times 10^{-3}$	120	436	390	366	0.03	3.4	3.7	5.8	$2.9 \times 10^{-4}$	$6.5 \times 10^{-3}$
WF 5 (coarse)	5160	$1.15 \times 10^{-3}$	96	192	300	247	0.06	3.7	8.1	8.2	$1.5 \times 10^{-3}$	$0.8 \times 10^{-2}$
WF 5 (fine)	5160	$1.15 \times 10^{-3}$	120	436	390	247	0.04	3.0	3.6	6.3	$2.4 \times 10^{-4}$	$1.0 \times 10^{-3}$

Table 1. In columns, from left to right: time averaged Reynolds number  $\overline{Re}$ , forcing amplitude  $A_f$ , number of physical grid points in radial, azimuthal and axial direction ( $N_r$ ,  $N_\theta$ ,  $N_z$ ), maximum friction Reynolds number ( $Re_\tau = \max_t \sqrt{\tau_w(t) / \rho_f(D/\nu)}$ ), minimum/maximum radial, azimuthal and axial resolution in inner units ( $\Delta r_-^+$ ,  $\Delta r_+^+$ ,  $\Delta(R\theta)^+$  and  $\Delta z^+$ ) and minimum and maximum time step ( $\Delta t_-$ ,  $\Delta t_+$ ) in units of  $D/\overline{U}$ . Optimisations are carried out on a coarse mesh with a large time step ( $\varepsilon \approx 10^{-9}$ ) where in row ‘Optimisation iterations’, we report the averages over all iterations. The following rows display the resolutions in the optimal waveforms (WF 1–5) as computed in the optimisation loop (coarse) where  $\varepsilon \approx 10^{-9}$  and verification cases for the optimal waveforms (fine) on a fine grid with a smaller time step ( $\varepsilon \approx 10^{-13}$ ).

perturbations are present, due to the inherent linear stability of pulsatile pipe flow at these flow parameters (Thomas *et al.* 2011). In physical experiments, turbulence is maintained (or retriggered) by unavoidable imperfections in the set-up – such as inlet disturbances – that are absent in idealised numerical simulations (Reynolds 1883; Avila, Barkley & Hof 2023). To model these experimental disturbances and enable a more faithful comparison with experiments, we introduced a small volumetric body force perturbation added to the right-hand side of the Navier–Stokes equation (2.1),

$$f_b = A_f \begin{pmatrix} 4r(r^2 - 1)^2 [(r^2 - 1) \sin^2(\theta) - (5r^2 - 1) \cos^2 \theta] - (24r^2 - 16) \sin \theta \\ [-8r(r^2 - 1)^2 \sin \theta - 72r^2 - 16] \cos \theta \\ 0 \end{pmatrix}, \tag{2.3}$$

where  $A_f \geq 0$  is the perturbation amplitude. Using an amplitude of  $A_f = 2.25 \times 10^{-3}$ , turbulence is triggered cyclically when the Reynolds number rises. Note that in the first period, the turbulent kinetic energy coincides with the undisturbed case, indicating that the perturbation – while ensuring the cyclic transition to turbulence – does not significantly alter the overall flow dynamics. This claim is later (Appendix C) verified for each of our optimal waveforms. The specific form and effect of the perturbation have been analysed in detail by Kranz (2024), where different types of forcings and their amplitudes were explored. For the remainder of this paper, we always applied the forcing (2.3) and adjusted  $A_f$  based on  $\overline{Re}$  (table 1).

### 2.3. DNS

Simulations were carried out on an NVIDIA Tesla P100 using the open-source pseudo-spectral GPU-based research code *nsPipe-CUDA* (López *et al.* 2020; Morón *et al.* 2024), which uses a Fourier–Galerkin discretisation ansatz for azimuthal ( $N_\theta$  modes) and axial directions ( $N_z$  modes) and higher-order finite differences in the radial direction ( $N_r$  points). The grid is non-uniform in the radial direction with finest resolution near

the walls. The value of the driving force  $\Pi_U(t)$  to enforce the bulk velocity waveform  $U(t)$  is adjusted in each time step. The code integrates the Navier–Stokes equations using a Crank–Nicolson scheme and iterates on the nonlinear term. The time step is adjusted dynamically so that the error between iterations on the nonlinear term is always smaller than a threshold  $\varepsilon$  (see table 1). Mostly, simulations were carried out at  $Wo = 10$ ; however, selected cases were conducted at  $Wo = 10\sqrt{2}$ . Unless stated otherwise, all simulations were initialised with a turbulent initial condition obtained from a steady run. We investigated three different average Reynolds numbers:  $\overline{Re} = 4300$ ,  $\overline{Re} = 5160$  and  $\overline{Re} = 8600$ . Optimisations were carried out with a relatively coarse spatial and temporal resolution, a fixed (turbulent) initial condition and a fixed body force (see (2.3)). As pointed out by Foggi *et al.* (2023a), small time steps may be needed if abrupt changes in the pressure gradient are present. In addition, they found that results are sensitive to the domain size. We verified the robustness of our optimal waveforms in terms of the spatial and temporal resolution, the initial condition, the pipe length and the forcing term, and refer to Appendix C. Detailed information about the spatial and temporal resolution is given in table 1. On our fine grid, we verified *a posteriori* that the worst case Courant–Friedrichs–Lewy (CFL) numbers fall below 0.1 for  $\overline{Re} = 5160$  and below 0.4 for  $\overline{Re} = 8600$ .

We considered two quantities of interest: the  $n$ -cycle-averaged wall shear stress  $\overline{\tau}_w$  and power input by  $\overline{P}$ , defined as

$$\begin{aligned} \overline{\tau}_w &= \frac{1}{n} \sum_{i=1}^n \overline{\tau}_w^{(i)}, & \overline{\tau}_w^{(i)} &= \frac{1}{T} \int_{iT}^{(i+1)T} \tau_w(t) \, dt, \\ \overline{P} &= \frac{1}{n} \sum_{i=1}^n \overline{P}^{(i)}, & \overline{P}^{(i)} &= \frac{1}{T} \int_{iT}^{(i+1)T} Q(t) \Delta p(t) \, dt, \end{aligned} \tag{2.4}$$

where  $\tau_w(t) := \mu \langle (\partial u_z / \partial r)|_{r=D/2} \rangle_{\theta,z}(t)$  is the axial and azimuthal averaged wall shear stress,  $\mu$  the dynamic viscosity,  $Q(t) = (\pi/4)U(t)$  the volumetric flow rate and  $\Delta p(t)$  the pressure drop. We refer to  $\overline{\tau}_w^{(i)}$  and  $\overline{P}^{(i)}$  as the per-period wall shear stress and power input of period  $i$ , respectively. Note that the zeroth cycle ( $i = 0$  in (2.4)) always features significantly higher turbulence levels due to the initial condition and is always discarded for the analysis. We remark that minimising the average power input is equivalent to minimising the total energy consumption  $nT \cdot \overline{P}$ . Thus, in what follows, power savings can be interpreted as energy savings and *vice versa*. The quantities  $\overline{\tau}_w$  and  $\overline{P}$  were normalised with respect to the steady-state values (indicated by  $(\cdot)_b$ ), inferred from the Blasius friction law  $2\overline{\tau}_{w,b} = 0.0791 \overline{U} \overline{Re}^{-1/4}$  (Blasius 1913). Following Scarselli *et al.* (2023), we defined the drag reduction  $D_r$  and power saving  $P_s$  as

$$D_r = \frac{\overline{\tau}_{w,b} - \overline{\tau}_w}{\overline{\tau}_{w,b}} = 1 - \overline{\tau}_w, \quad P_s = \frac{\overline{P}_b - \overline{P}}{\overline{P}_b} = 1 - \overline{P}, \tag{2.5}$$

where values larger/smaller than zero indicate reduced/increased drag or power consumption, when compared with steady driving. Because of the chaotic nature of turbulence, a single period of a pulsatile run is not representative for the overall behaviour: figure 1(c) shows the evolution of the wall shear stress and the power input over three periods for a turbulent flow periodically driven according to the waveform in figure 1(a). Analysing the per-period wall shear stress for this specific run unveils that  $\overline{\tau}_w^{(1)}$  is approximately 14.7% larger than  $\overline{\tau}_w^{(2)}$  and 4.4% larger than  $\overline{\tau}_w^{(3)}$ . For the power input,




Figure 2. (a) Relative standard error of the per-period wall shear stress ( $\zeta(\bar{\tau}_w)$ ) versus the number of averaging periods  $n$ . Red dots mark the number of periods needed to achieve values for  $\zeta^*$  of 2.5%, 1.25%, 0.625%, 0.3125% and 0.25%. The dashed line shows a  $C_1/\sqrt{n}$ -fit to the data. (b) Computational time (in hours) to achieve a given  $\zeta(\bar{\tau}_w)$ , where red dots correspond to the same  $\zeta^*$  values as in panel (a) and the dashed line shows a quadratic fit to the data.

$\bar{P}^{(1)}$  is approximately 19.5% larger than  $\bar{P}^{(2)}$  and 5.8% larger than  $\bar{P}^{(3)}$  (figure 1d). This motivates the question of finding the minimum number of averaging periods that approaches the statistically steady values of  $\bar{\tau}_w$  and  $\bar{P}$ . To address this question, we considered the relative standard error of the vector of per-period wall shear stress  $\bar{\tau}_w = (\bar{\tau}_w^{(i)})_{i=1,\dots,n}$  and power input  $\bar{P} = (\bar{P}^{(i)})_{i=1,\dots,n}$ , which are

$$\zeta(\bar{\tau}_w) = \frac{\sigma(\bar{\tau}_w)}{\bar{\tau}_w\sqrt{n}}, \quad \zeta(\bar{P}) = \frac{\sigma(\bar{P})}{\bar{P}\sqrt{n}}, \quad (2.6)$$

where  $\sigma(\cdot)$  denotes the sample standard deviation and is required to fall under a threshold  $\zeta^*$ . Initially, simulations were conducted and post-processed averaging over three periods, and were automatically extended by another period until  $\zeta(\cdot) \leq \zeta^*$ . Reducing  $\zeta^*$  comes with extensive computational efforts, based on the simulation disseminated in figure 1(c). In figure 2(a), we illustrate the number of averaging periods needed to achieve a given  $\zeta^*$  (computed on  $\bar{\tau}_w$ ). Realising  $\zeta^* = 2.5\%$  requires three averaging periods, corresponding to a computational time of approximately 48 min (see figure 2b). Halving  $\zeta^*$  requires nine periods (2.5 h of computing time), while reducing it by a factor of ten requires almost 100 periods (27.2 h).

#### 2.4. Finite-dimensional optimisation problem and waveform design

In the optimisation problem, (1.2), we sought bulk velocity waveforms  $U(t)$  that deliver a desired averaged velocity bulk  $\bar{U}_d$ , while minimising either the mean wall shear stress ( $\mathcal{J} = \bar{\tau}_w$ ) or the mean power input ( $\mathcal{J} = \bar{P}$ ) (see (2.4)). Equation (1.2) is a complex optimisation problem: it is constrained by a partial differential equation ( $\mathbf{u}$  has to fulfil the Navier–Stokes equation (2.1)), it features state constraints (the desired flux  $\bar{U}_d = 1$  needs to be satisfied) and it is of infinite dimensions (seeking functions  $U(t)$ ). To simplify the optimisation problem (1.2) and overcome the last two complexity features, we considered two different waveforms  $U_\eta(t)$  that can be defined by a finite number of parameters  $\eta \in \mathbb{R}^d$  and (by design) deliver  $\bar{U}_d = 1$ .

First, we considered triangular waveforms as was done by Scarselli *et al.* (2023). The waveforms consist of an acceleration phase of length  $T_a$ , accelerating the flow from a minimum Reynolds number ( $Re^-$ ) to a maximum Reynolds number ( $Re^+$ ), a deceleration phase of length  $T_d$  slowing it down to  $Re^-$  and a rest phase of length  $T_r$  of steady driving, see figure 1(a). We parametrised the waveform by the acceleration time, and maximum

and minimum Reynolds number ( $\eta = (T_a, Re^+, Re^-)$ ), obtaining the shape

$$U_{(T_a, Re^+, Re^-)}(t) = \frac{\bar{U}}{Re} \cdot \begin{cases} Re^- + \frac{Re^+ - Re^-}{T_a} t, & 0 \leq t < T_a, \\ Re^+ + \frac{Re^- - Re^+}{T_d} (t - T_a), & T_a \leq t < T_a + T_d, \\ Re^-, & T_a + T_d \leq t \leq T. \end{cases} \quad (2.7)$$

By restricting the waveform in terms of fixing  $Re^+$  and  $Re^-$ , this triangular waveform can be simplified into a uni-variant one, only dependent on the acceleration time  $\eta = T_a$ . Note that  $T_d$  is adjusted based on  $T_a = T_a(T_d)$ , so that the desired mass flux is satisfied, while  $T_r$  is fixed by  $Re^+$  and  $Re^-$ .

Second, we used a truncated Fourier series to express the waveform, allowing for a wide range of continuous waveform shapes. We also incorporate the desired bulk velocity  $\bar{U}_d = 1$ , where we fix the phase of the first mode ( $b_1 = 0$ ) to avoid waveforms that are identical up to a time shift. The waveform is given by

$$U_\eta(t) = 1 + \sum_{k=1}^N a_k \cos\left(\frac{2\pi}{T} nt\right) + \sum_{k=2}^N b_k \sin\left(\frac{2\pi}{T} nt\right), \quad \eta := (a_1, \dots, a_N, b_2, \dots, b_N). \quad (2.8)$$

In other words, we aimed to identify the  $2N - 1$  Fourier coefficients, such that the bulk velocity waveform minimises the mean wall shear stress or mean power input.

Both waveform designs automatically satisfy the desired flux constraint in (1.2) and the optimisation problem can be reduced to finite dimensions as

$$\min_{\eta \in \mathcal{Q} \subset \mathbb{R}^d} \mathcal{J}(U_\eta(t)), \quad (2.9)$$

where  $\mathcal{Q} \subset \mathbb{R}^d$  is the admissible set realising bounds for the parameter vector, specified later.

### 3. Optimisation method

A challenge in solving the optimisation problem (2.9) is that the functional  $\mathcal{J}$  is noisy. The noise amplitude is inversely proportional to the square root of the number of periods considered in the DNS, see figure 2(a) and (2.6). As shown in Appendix A, gradient-based methods are not well suited for this optimisation problem owing to the large number of periods needed to reduce the noise level, so the robust evaluation of gradients of  $\mathcal{J}$  is ensured. Bayesian optimisation (BO) naturally embraces the noisy nature of (2.9) and, as shown later, produces results that are consistent regardless of the number of periods chosen to evaluate  $\mathcal{J}$ .

BO is a global data-driven optimisation technique for black-box functions that feature stochastic noise, are costly to evaluate, and possibly non-convex and non-differentiable (Kushner 1964; Mockus 1972). After initialising a dataset by observing the objective at initial points,  $\mathcal{D}^{(0)} = \bigcup_i \{(\eta^{(0)_i}, \mathcal{J}(\eta^{(0)_i}))\}$ , where  $\eta^{(0)_i}$  denote the initial points, a probabilistic surrogate model,  $\mathcal{J}_s^{(0)}$ , usually in the form of Gaussian processes (GPs), is fitted to  $\mathcal{D}^{(0)}$ . An acquisition function  $\alpha(\eta)$  is then used to decide where to evaluate the function next by balancing the trade-off between exploration (sampling where the uncertainty of the GP is high) and exploitation (sampling where the surrogate model predicts low values). A common choice for  $\alpha$  (chosen in this study) is the expected improvement,  $\alpha(\eta) = \mathbb{E}[\max(0, \mathcal{J}^* - \mathcal{J}(\eta))]$ , where  $\mathcal{J}^*$  is the current best observed value. The next point to sample  $\mathcal{J}$  is given by  $\eta^{(j+1)} = \arg \max_{\eta \in \mathcal{Q}} \alpha(\eta)$  and we set

$\mathcal{D}^{(j+1)} = \mathcal{D}^{(j)} \cup (\boldsymbol{\eta}^{(j+1)}, \mathcal{J}(\boldsymbol{\eta}^{(j+1)}))$ . Iteratively, an updated surrogate  $\mathcal{J}_S^{(j+1)}$  is fitted to the updated dataset, a new point maximising  $\alpha$  is sampled and the dataset is extended. Detailed descriptions of BO are given by Mockus (1994) and Jones, Schonlau & Welch (1998) and references therein.

For the BO algorithm, we constructed the initial dataset  $\mathcal{D}^{(0)}$  by randomly sampling  $5d$  initial data points, in other words,  $\mathcal{D}^{(0)} = \bigcup_{i=1}^{5d} \{(\boldsymbol{\eta}^{(0)_i}, \mathcal{J}(\boldsymbol{\eta}^{(0)_i}))\}$ , where  $\boldsymbol{\eta}^{(0)_i}$  is drawn from a continuous uniform distribution. Convergence criteria in BO are not as straightforward as in gradient-based methods, as, due to the exploring–exploiting approach, the iterative function values are generally not monotonically converging towards the optimum. As the dataset grows, the surrogate  $\mathcal{J}_S$  is expected to improve continuously and so the expected minimum (the minimum of  $\mathcal{J}_S$ ) is gradually refined. However, if the acquisition function favours exploring new areas and the observed value falls out of line with the existing surrogate, the expected minima may change. To capture that the algorithm converges to an exploration-unbiased minimum, we demanded the relative difference between expected minima to fall beneath a threshold  $\epsilon$ ,

$$\left| \frac{\min_{\boldsymbol{\eta}} \mathcal{J}_S^{(j+1)}(\boldsymbol{\eta}) - \min_{\boldsymbol{\eta}} \mathcal{J}_S^{(j)}(\boldsymbol{\eta})}{\min_{\boldsymbol{\eta}} \mathcal{J}_S^{(j)}(\boldsymbol{\eta})} \right| \leq \epsilon \quad (3.1)$$

for a patience of at least five iterations. Note that this way, an optimisation is at least  $5(d + 1)$  iterations long; however, the risk of aborting prematurely is reduced. Convergence criteria of this kind have been reported in the literature and have been demonstrated to yield reasonable timings to stop the BO algorithm (Ishibashi *et al.* 2023).

We assessed the performance in terms of feasibility, computational effort and robustness of BO using the restricted (uni-variant) triangular waveform described in § 2.4. Here, choosing  $Wo = 10$ ,  $Re = 4300$ ,  $Re^+ = 9400$  and  $Re^- = 1600$  (see figure 1a), we aimed to identify the  $\bar{\tau}_w$ -optimal acceleration time. The acceleration time is bounded by  $0.01T$  and  $0.68T$  ( $T_a \in \mathcal{Q} = [0.01, 0.68]T$ ) to obtain continuous waveforms and to ensure that  $Re$  is realised.

In figure 3(a–c), we show the best surrogates (surrogates based on the last dataset) for the wall shear stress ( $\mathcal{J}_{\tau_w}^*$ ) for three different values of  $\zeta^*$ . In the Bayesian optimisation, regardless of the admissible standard error, expected optima (red markers) are obtained at the lower bound of  $\mathcal{Q}$  ( $T_a = 0.01T$ ) and surrogate models visually coincide. Using a 95 % of confidence interval (CI), the differences in the expected function values are statistically insignificant (0.950 and 0.949 are within the 95 % CI of 0.954 with a standard error of  $\zeta = 0.625$  %). We tested the performance of BO by further reducing the number of initial averaging periods to  $n = 2$ . We refrained from decreasing it to one period only as the  $\zeta^*$ -criterion does not apply in that case and, thus, no upper bound for the noise can be given. As shown in figure 3(d), the surrogate for  $\bar{\tau}_w$  as well as the expected minimum of  $\bar{\tau}_w = 0.946$  are almost identical to cases with a higher number of averaging periods, see figure 3(a–c). In summary, unlike gradient-based methods (see Appendix A, figure 10), BO produces results that are consistent as  $\zeta^*$  is reduced.

In figure 3(d), we also show the best surrogate for the mean power input ( $\mathcal{J}_{P,S}^*$ ). Here, the surrogate close to the expected minimum of  $\bar{P} \approx 1.3$  at  $T_a = 0.16$  is flat, which intensifies the problem of accurately approximating gradients.

To put computational efforts into perspective: the optimisations shown in figure 3(d) took approximately 9 h, whereas to arrive at a similar minimum with the gradient-based SLSQP method described in Appendix A took 4.5 days. These results show that BO is a powerful tool to approach the optimisation problem at hand: (i) it can deal by design

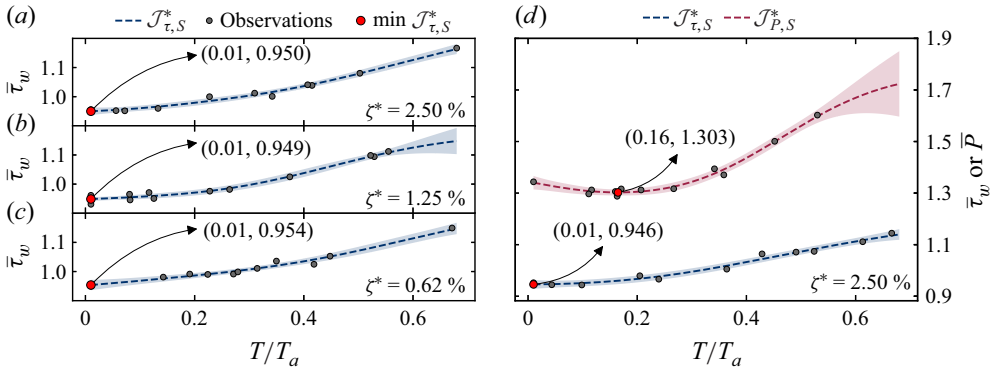


Figure 3. (a)–(c) Best surrogate for the mean wall shear stress ( $\mathcal{J}_{\tau,S}^*$ ) and expected optima for different allowable standard errors  $\zeta^* \in \{2.5, 1.25, 0.625\}$  %. (d) Best surrogates for the wall shear stress and power input,  $\mathcal{J}_{\tau,S}^*$  and  $\mathcal{J}_{P,S}^*$ , respectively, as well as expected optima, when reducing the number of initial averaging periods to two. In all cases, the flow was driven according to the waveform from figure 1(a), where  $\overline{Re} = 4300$ ,  $Re^- = 1600$  and  $Re^+ = 9400$ . Shaded areas indicate the uncertainty of the surrogate model. Wall shear stresses are given in units of  $\bar{\tau}_{w,b}$  and power inputs in  $\bar{P}_b$ .

with the noisy setting; (ii) it is efficient in terms of the number of function evaluations; (iii) already made observations are reusable. Hence, for all optimisations in this paper, BO was the tool of choice. Specifically, we used the implementation of BO available in the python library [scikit-optimize](https://scikit-optimize.github.io/).

## 4. Results

### 4.1. Uni-variant approach

At  $Wo = 10$ , we optimised the uni-variant waveform at Reynolds numbers,  $\overline{Re} = 5160$ ,  $Re^+ = 11\,280$ ,  $Re^- = 1920$ , averaging initially over two periods and enforcing  $\zeta^* = 2.5\%$ . In figure 4(a), we compare the (best) surrogates for the mean wall shear stress ( $\mathcal{J}_{\tau,S}^*$ ) and the mean power input ( $\mathcal{J}_{P,S}^*$ ) at  $\overline{Re} = 4300$  and  $\overline{Re} = 5160$ . Again, minima of  $\bar{\tau}_w$  are obtained at the smallest admissible acceleration time,  $T_a = 0.01T$ , whereas the drag reduction improves from  $D_r = 5\%$  at  $\overline{Re} = 4300$  to  $D_r = 14\%$  at  $\overline{Re} = 5160$ . The  $\bar{P}$ -optima deviate slightly in terms of the acceleration time ( $T_a = 0.16T$  at  $\overline{Re} = 4300$  and  $T_a = 0.13T$  at  $\overline{Re} = 5160$ ), whereas the power loss is decreased from  $P_s = -30\%$  to  $P_s = -14\%$ . Both surrogates for  $\overline{Re} = 5160$  show increased noise levels compared with the lower  $\overline{Re}$  case and thus the uncertainty of each surrogate (shaded areas) also increases. The average standard deviation of single observations regarding the surrogate increases from 0.7% to 1.3% for the  $\bar{\tau}_w$  surrogate and from 0.8% to 2.1% for the  $\bar{P}$  surrogate. Considering the noise, the difference in expected optimal points for  $\bar{P}$  is insignificant: evaluation of the  $\bar{P}$ -surrogate for  $\overline{Re} = 5160$  at  $T_a = 0.16$  yields a power loss of 15.0% (versus 14.4% at  $T_a = 0.13$ ). The increased noise levels underscore that gradient-based methods are inappropriate for the given problem.

In figure 4(b), we show the  $\bar{\tau}_w$ -optimal driving waveform ( $T_a = 0.01T$ ), as well as the evolution of the resulting wall shear stress and power input, averaged over five periods. At  $\zeta = 2.1\%$ , the drag reduction that is actually achieved is insignificantly larger than that projected by the surrogate ( $D_r = 15\%$  compared with  $D_r = 14\%$ ). Due to abrupt acceleration, peak values for the wall shear stress overshoot steady values by up to seven times in the beginning of the period. During deceleration, the wall shear stress quickly declines, eventually rising again until hitting a second peak at approximately



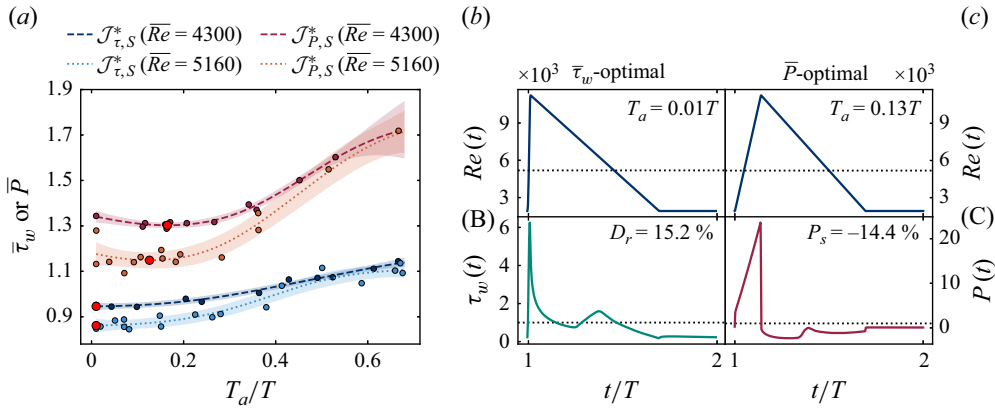


Figure 4. (a) Best surrogates and expected minima for the mean wall shear stress  $\mathcal{J}_{\tau,S}^*$  and the mean power input  $\mathcal{J}_{P,S}^*$  at Reynolds numbers of  $\overline{Re} = 4300$  and  $\overline{Re} = 5160$ . Shaded areas indicate the uncertainty of the surrogate model. (b,c) The  $\bar{\tau}_w$ - and  $\bar{P}$ -optimal and sub-optimal waveforms ( $\overline{Re} = 5160$ ), and the resulting evolutions of the wall shear stress and power over the time span of two periods (where capital letters are associated with the lower case letters). The average Reynolds number as well as steady values for the wall shear stress and power input, obtained by Blasius’ friction law, are indicated by dotted lines. Wall shear stresses are given in units of  $\bar{\tau}_{w,b}$  and power inputs in  $\bar{P}_b$ .

$0.5T$ . This behaviour may be explained as follows. In low velocity phases ( $\overline{Re} < 5160$ ), turbulence levels decay, setting a favourable initial condition for the following acceleration. Opposed to slow accelerations, where turbulence has sufficient time to rise until being fully developed at peak velocities and realising accordingly large wall shear stresses, in fast accelerations, peak turbulence levels are attained during deceleration, where flow velocities are already comparably small.

A reduction of the mean wall shear stress does not ensure that the mean power input is lower compared with steady conditions. In fact, the negative power saving of  $P_s = -14\%$  achieved by the *a posteriori* simulation of the  $\bar{P}$ -optimal waveform ( $T_a = 0.13T$ ) in figure 4(b) shows that, at this Reynolds and Womersley number, any pulsatile driving of this shape results in power losses. The power loss compared with the steady case is caused by the additional energy input required to accelerate the flow. Accelerating the flow within short time frames requires a large pressure gradient and hence the required instantaneous power input increases up to 24 times the steady value.

#### 4.2. Tri-variant approach

We kept  $(\overline{Re}, Wo) = (5160, 10)$  and now allowed maximum and minimum Reynolds numbers  $Re^+ \in [9400, 11280]$  and  $Re^- \in [1600, 1920]$  to change (tri-variant approach) – a combination of the uni-variant case at  $\overline{Re} = 4300$  and  $\overline{Re} = 5160$ . In figure 5, the optimisation outcome is displayed using partial dependence plots. The partial dependence is calculated by averaging the objective value for a number of random samples in the search-space, while keeping the remaining dimensions fixed (Goldstein *et al.* 2014). This averages out the effect of varying the other dimensions and shows the influence of one or two dimensions on the objective function.

Figure 5(a) shows the dependence of the mean wall shear stress on  $T_a$ ,  $Re^+$  and  $Re^-$ . The function value is mostly ruled by the acceleration time as nearly no dependence on the maximum and minimum Reynolds number can be observed (equal values up to the fifth decimal place at lower and upper bounds for  $Re^+$  and  $Re^-$ ). Thus, the same drag reductions of approximately 15% as in the uni-variant case are obtained at the upper

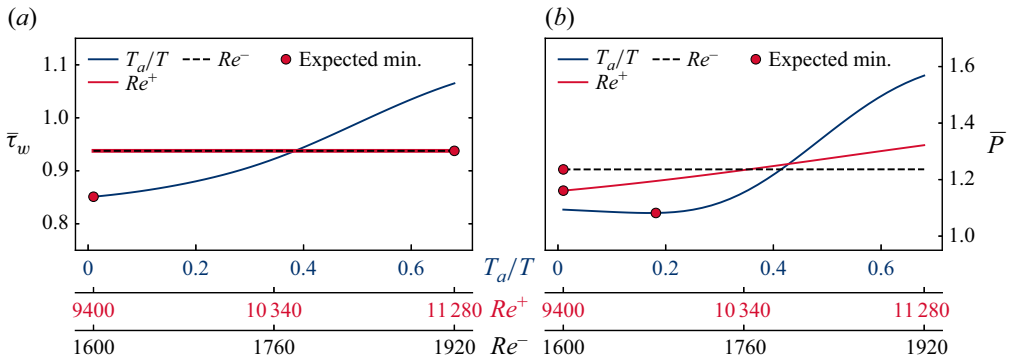


Figure 5. Partial dependence plots for (a) the wall shear stress in units of  $\bar{\tau}_w$  and (b) the power input in units of  $\bar{P}$  in the tri-variant optimisation. Circles show the expected minimum.

bounds of  $Re^+$  and  $Re^-$  and the lower bound of  $T_a$ . As shown in figure 5(b), the mean power input shows dependencies on the maximum Reynolds number and the acceleration time. The  $T_a$ -dependence is more substantial, spanning power losses from  $-8\%$  to  $-58\%$ . Varying  $Re^+$  also has an impact on  $\bar{P}$ . Specifically, modifying  $Re^+$  for a fixed acceleration time can influence the power saving by up to 16% ( $\bar{P} = 1.14$  for  $Re^+ = 9400$  and  $\bar{P} = 1.30$  for  $Re^+ = 11280$  (red curve in figure 5b)). The largest power saving of ( $P_s = -2\%$ , power loss of 2%) is found at an acceleration time of  $T_a = 0.18T$ , and at lower bounds for maximum and minimum Reynolds numbers ( $Re^+ = 9400$ ,  $Re^- = 1600$ ), corresponding to the shortest possible rest phase. As argued before, the slight difference in optimal acceleration times ( $T_a = 0.16T$  for the uni-variant and  $T_a = 0.18T$  for the tri-variant case) is insignificant considering the flatness of the  $T_a$ -dependence around the minimum in combination with the noisy setting. The expected minimum of  $\bar{P} = 1.02$  falls in line with the *a posteriori* evaluation at the optimal parameters ( $\bar{P} = 1.015$ , not shown here).

### 4.3. Truncated Fourier series

In the Fourier series expansion of arbitrary waveforms, we only optimised for the power input ( $\mathcal{J} = \bar{P}$ ) for computational reasons. However, we note that all found power-reducing waveforms also reduce the mean wall shear stress substantially. We first considered  $Wo = 10$  and later extended the analysis to  $Wo = 10\sqrt{2}$ . At  $\bar{Re} = 5160$ , we considered five coefficients  $a_1, a_2, a_3$  and  $b_2, b_3$  in (2.8) and to restrict the maximum Reynolds number, we limited the magnitude of each by  $1/6$  ( $|a_k|, |b_k| \leq 1/6, k = 1, \dots, N$ ), because in the BO toolkit used, it is not trivial to realise non-box constraints (i.e. to restrict  $Re^+$  directly).

Figure 6(a) shows one period of the  $\bar{P}$ -optimal waveforms for three independent optimisation runs (labelled 1–3). On average, four periods were necessary to achieve  $\zeta^* = 2.5\%$  and 46 iterations of the BO algorithm are run to satisfy the convergence criterion (3.1). Therefore, one optimisation at this parameter regime took approximately 87 h of computing time. All optimisations yield different optima: WF 1 ( $\eta = (-1/6, 0.08, 0.08, -1/6, -0.02)$ ) realises the greatest power saving ( $P_s = 11.8\%$ ), while WF 2 ( $\eta = (0.014, -0.16, -0.14, -1/6, -0.02)$ ) and WF 3 ( $\eta = (1/6, -0.07, 0.13, -1/6, 0.06)$ ) yield power savings of  $P_s = 10.8\%$  and  $P_s = 10.4\%$ , respectively. Statistically, using a 95% CI and  $\zeta^* = 2.5\%$ , the difference in power savings is insignificant (a CI of approximately 10.8% is given by [6.4%, 15.2%]). All power-optimal waveforms also reduce drag substantially: WF 1 reduces drag by  $D_r = 21.1\%$  when compared with steady driving, while WF 2 and 3 realise drag reductions of  $D_r = 20.6\%$  and  $D_r = 21.7\%$ .

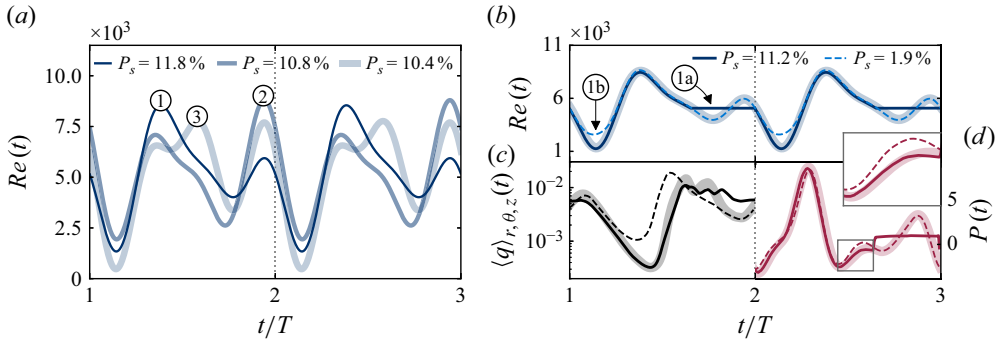


Figure 6. (a) Power-optimal waveforms obtained from three independent runs of the truncated Fourier approach at  $\overline{Re} = 5160$  and  $Wo = 10$ , where  $N = 3$  (see (2.8)) and  $|a_k|, |b_k| \leq 1/6, k = 1, 2, 3$  (phase-adjusted for the minimum velocity). (b) Best-performing waveform from panel (a) (WF 1) and two modifications thereof (1a and 1b). (c) and (d) Evolution of the cross-sectional kinetic energy (units of  $\overline{U}^2$ ) and of the power input (units of  $\overline{P}_b$ ; line styles according to panel b).

Foggi *et al.* (2023a) proposed a different metric than (2.5) to evaluate power savings. They accounted for the impossibility of reaching below the laminar state, defining the saving  $P_e = (1 - \overline{P}) / (1 - \overline{P}_\ell)$ , where  $\overline{P}_\ell$  is the power input of the laminar flow state. Using that metric, waveforms 1–3 realise increased savings of  $P_e = 17.6\%$ ,  $P_e = 16.1\%$  and  $P_e = 15.6\%$ , respectively.

The finding of multiple minima may have various sources. First, the initial dataset is randomly sampled across the parameter space, which can initially guide the BO loop towards different regions of the parameter space. Since the computational resources are limited, the loop cannot gather sufficiently many data points to ensure that the surrogate functions are independent of the initialisation. Second, the noisiness can lead to multiple noise-induced minima, especially in flat  $\mathcal{J}(\eta)$ -landscapes, where noise may be larger than actual changes in the functional.

Even if found at different parameters, all waveforms exhibit similar features. After the minimum flow rate, they follow a steep acceleration phase. As in the power-optimal triangular waveforms, this fast acceleration delays the onset of turbulence (Scarselli *et al.* 2023). Indeed, the acceleration length of approximately  $0.21T$  (measured from the global minimum to the first local maximum in figure 6a) is similar to that found in the tri-variant approach ( $T_a = 0.18T$ ). Also, as in the triangular waveforms, the maximum velocity is followed by a deceleration phase. Additionally, in contrast to the triangular waveforms, no distinct rest phase exists, and, instead, a second valley–peak combination occurs.

We analysed the role of the second peak by modifying the best-performing waveform (WF 1 in figure 6a). Specifically, we replaced the second peak with a rest phase while keeping  $\overline{Re} = 5160$  (see WF 1a figure 6b). The power saving for that waveform remains (statistically) identical (11.2%). A different modification was applied to WF 1, namely the minimum Reynolds number was increased from 1500 to 3400 (WF 1b in figure 6b). Note that WF 1b realised a slightly larger average Reynolds number of  $\overline{Re} = 5330$ ; however, we also normalised the power with the steady power input associated with  $\overline{Re} = 5330$ . This modification of the waveform results in a substantial decrease of power saving from 11.8% to 1.9%. The additional power spending in WF 1b is linked to the turbulent kinetic energy  $\langle q \rangle_{r,\theta,z}(t)$ , shown in figure 6(c). As a result of the flatter minimum of WF 1b, compared with 1 and 1a, turbulence does not sufficiently decay during the pre-peak phase. Consequently, the onset of turbulence is quicker and peak intensities occur in earlier phases of the deceleration, where the velocity is comparably large. In other words, compared with

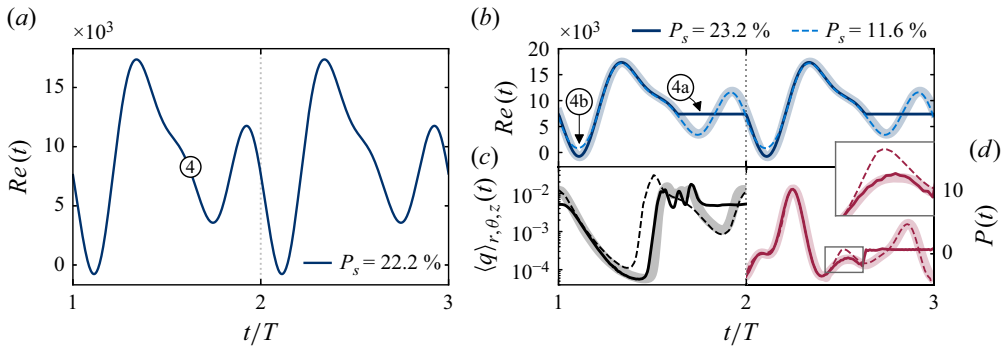


Figure 7. (a) Power-optimal waveform obtained by the truncated Fourier approach ( $\overline{Re} = 8600$ ,  $Wo = 10$ ) where  $N = 3$  (see (2.8)) and  $|a_k|, |b_k| \leq 1/4$ ,  $k = 1, 2, 3$  (phase-adjusted). (b) WF 4 and two modifications thereof (4a and 4b). (c) and (d) Evolution of the cross-sectional kinetic energy (units of  $\overline{U}^2$ ) and the power input (units of  $\overline{P}$ ); line styles according to panel b).

1 and 1a, in 1b, a larger amount of fluid has to be displaced at higher turbulence levels, resulting in a larger power consumption during that phase. In figure 6(d), we demonstrate that differences in the evolution of the power input of WF 1 and 1b mainly occur within the phase of large turbulence intensities.

As all optimal waveforms realise the bounds for at least one coefficient, for WF 1, we allowed for larger magnitudes of  $1/4$ . Using all observations for WF 1, the optimisation was continued until criterion (3.1) was fulfilled (31 iterations). At  $\eta = (-0.17, 0.06, 0.1, -0.17, -0.09)$  (found within the interior of the admissible set), the obtained waveform (not shown here) is nearly identical to WF 1 and realises a similar power saving ( $P_s = 12.3\%$ ).

Additionally, we investigated  $\overline{Re} = 8600$ , drastically increasing the computational effort. Even with our rather coarse computational mesh (see table 1), single periods took approximately 4 h of computational time, and, additionally, the noise level increases, requiring up to 52 periods to fulfil  $\zeta(\overline{P}) \leq 2.5\%$ . In total, the 63 iterations of the BO loop needed to satisfy criterion (3.1) approximately took 65 days of computing time.

Figure 7(a) shows the power-optimal WF 4 ( $\eta = (0.23, 0.25, -0.16, -0.15, -0.05)$ ) obtained for  $N = 3$  and  $|a_k|, |b_k| \leq 1/4$ . Compared with the best-performing waveform at  $\overline{Re} = 5160$ , the power saving doubles ( $P_s = 22.2\%$ , or, if preferred,  $P_e = 28.7\%$ ), whereas the drag is reduced by 36.5%. WF 4 exhibits similar characteristics to waveforms 1–3; however, the minimum flow rate becomes negative, corresponding to backflow ( $U < 0$ ).

In figure 7(b), we show two modifications of WF 4, namely 4a and 4b as done earlier for  $\overline{Re} = 5160$ . The findings are shown in figure 7(c,d) and are analogous: the second peak does not play an important role and the critical aspect of the profile is the deep minimum in the bulk velocity that contributes to sinking turbulence intensities drastically before acceleration. Lastly, the effect of the Womersley number was investigated. At  $\overline{Re} = 5160$  and  $Wo = 10\sqrt{2}$ , we carried out an additional optimisation, where  $N = 3$  and  $|a_k|, |b_k| \leq 1/6$ ,  $k = 1, 2, 3$ . The optimal WF 5 ( $\eta = (1/6, -1/6, 0.03, 0.01, 0.10)$ ) realises a similar power saving of 10.0% ( $P_e = 14.9\%$ ), while reducing drag by approximately 20.6% (evaluated on our fine grid, see table 1). The waveform exhibits the same characteristics as WF 1–4, featuring lowest bulk velocities prior to the main acceleration phase. Again, the onset of turbulence is delayed into phases of low Reynolds numbers. In Appendix B, we show the optimal waveform for  $Wo = 10\sqrt{2}$  as well as the evolution of the power input and turbulent kinetic energy.

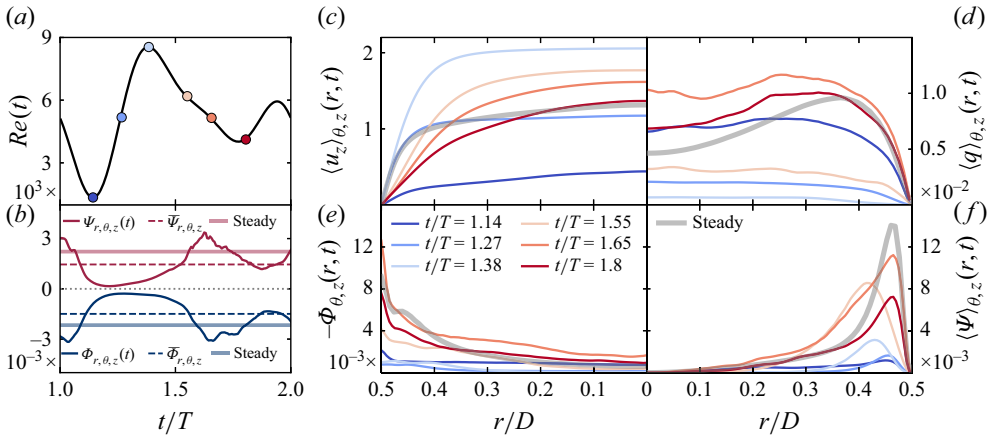


Figure 8. (a) Optimal WF 1 where colours encode the times in panels (c)–(f). (b) Time evolution of production ( $\Psi$ ) and dissipation ( $\Phi$ ) in the optimal waveform and steady pipe flow ( $Re = 5160$ ) in units of  $(\bar{U}^3/D)$ . (c–f) Axial and azimuthal averaged axial velocity (units of  $\bar{U}$ ), turbulent kinetic energy (units of  $\bar{U}^2$ ), dissipation and production (units of  $(\bar{U}^3/D)$ ), averaged over eight periods.

#### 4.4. Spatio-temporal intra-cycle mechanisms

We investigated the physical mechanisms that yield energy/drag reductions by computing the production ( $\Psi$ ) and dissipation ( $\Phi$ ) of the velocity fluctuations  $\mathbf{u}'$  (Feldmann, Morón & Avila 2020),

$$\Psi_\beta(t) = - \left( \langle u'_r u'_z \rangle_\beta \frac{\partial \langle u_z \rangle_\beta}{\partial r} \right) (t), \quad \Phi_\beta(t) = - \left( \frac{1}{Re} \langle \nabla \mathbf{u}' : \nabla \mathbf{u}' \rangle_\beta \right) (t), \quad (4.1)$$

as well as  $\langle u_z \rangle_{\theta,z}(r, t)$  and  $\langle q \rangle_{\theta,z}(r, t)$ .

In figure 8, we show these quantities for different times within WF 1. Additionally, in figure 9, we show instantaneous snapshots of the wall shear stress  $(\partial u_z / \partial r)(D/2, \theta, z, t)$  in a  $\theta$ – $z$ -plane for a steady case (left panel) and our optimal WF 1 (other panels).

The minimum of turbulent production occurs right after the minimum bulk velocity (figure 8b). At this time, turbulent fluctuations are almost completely damped. During the acceleration phase, turbulent production and dissipation are still low ( $t/T = 1.27$  in figures 8b, 8e and 8f). Therefore, although during the last half of the acceleration the mean profile has a larger gradient at the wall than the mean profile of the steady case (see figure 8c), the magnitude of the spatially averaged wall shear stresses (see figure 9) is almost equal between the two. At peak bulk velocity ( $t/T = 1.38$ ), the strong acceleration has elongated all velocity streaks close to the wall, resulting in an homogeneous wall shear stress distribution ( $t/T = 1.38$  in figure 9). During deceleration, turbulent production (and dissipation) increases. Interestingly, first ( $t/T = 1.55$ ), the peak of turbulent production is radially localised farther from the wall than the peak production of the steady case (figure 8f). Thus, turbulent fluctuations are mostly found towards the pipe centre at later pulsation phases ( $t/T = 1.65$  and  $t/T = 1.8$  in figure 8d). At  $t/T = 1.65$ , even though the magnitude of the turbulent fluctuations peaks, the wall shear stress is mostly unaffected (the spatial average amounts to approximately  $0.3\bar{\tau}_{w,b}$ ). Only towards the end of the deceleration is the wall shear stress distribution visually similar to the steady case ( $t/T = 1.8$ ). At these phases of the period however, production and (therefore) dissipation decrease, which explains the almost complete dampening of turbulence before the next acceleration phase.

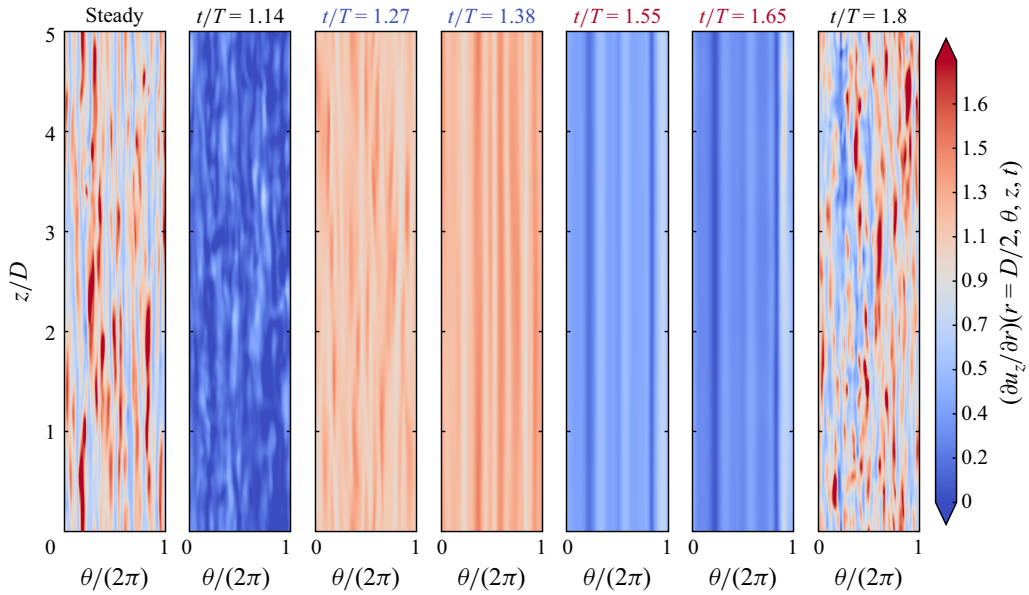


Figure 9. Instantaneous snapshots of the wall shear stress in units of  $\bar{\tau}_{w,b}$  in steady pipe flow ( $\bar{Re} = 5160$ , left panel) and in WF 1 for the same phases in the period as in figure 8(a). Blue/red titles are associated with acceleration/deceleration.

To complement our analysis, in Appendix D, we analogously show the quantities as in figures 8 and 9, but for a single harmonic (sine) waveform, with the same maximum and mean Reynolds number as WF 1. The waveform results in an extra energy consumption of approximately 15 % and neither drag reduction nor drag increase ( $D_r \approx 0$ ). Due to slower acceleration than in WF 1, we observe smaller wall shear stress in this phase of the period ( $t/T = 1.4$  and  $t/T = 1.65$  in figure 14). However, the elongated wall shear stress distribution remains. During deceleration, both production and dissipation rise faster (and are substantially larger) than in WF 1 (figure 13b). Further, here, the radial localisation of the production is similar to the steady case (figure 13f,  $t/T = 1.72, 1.82, 1.92$ ). Therefore, peak turbulent fluctuations appear in earlier phases of deceleration (at larger instantaneous Reynolds numbers) and radial localisation resembles the steady case. At late phases of the deceleration ( $t/T = 1.92$  in figure 14), the wall shear stress distribution is, like in WF 1, similar to the steady case; however, with a larger magnitude.

## 5. Conclusions

We show that by modulating the bulk velocity in time, while realising a fixed time-averaged volume flux, the average wall shear stress and net energy consumption can be greatly reduced when compared with steady driving. This is possible due to the nonlinear response of turbulence to unsteady driving: if and to what extent energy is saved or drag is reduced is highly dependent on the specific driving waveform – motivating optimisation for the latter.

Solving the optimisation problem using DNS is computationally demanding and results in a noisy estimation of the power input and drag due to the finite-time sampling of the turbulent attractor. Uncertainty in  $\mathcal{J}$  can be reduced by running longer simulations (i.e. averaging over more periods) at the cost of increasing computing times. In gradient-based methods, including adjoint optimisation, noise must be small in each iteration to enable



reasonable gradient approximations. BO deals by design with noisy objective functionals by modelling noise with Gaussian processes. Hence, single functional evaluations can be noisy and thus cheap. As more observations are added, the surrogate model is iteratively refined and the uncertainty in the surrogate is much smaller than the noise in the functional. We note that BO had already been used in fluid problems, e.g. for flow control by Morita *et al.* (2022) and Mallor *et al.* (2023), who showed that a few function evaluations are needed to converge to global optima. Our work unveils the ability of BO to deal with (large) noise levels and highlights its prowess to optimise average turbulent flow quantities. Generally, we expect BO to significantly outperform gradient-based algorithms in the optimisation of ergodic properties of chaotic systems such as turbulence.

We find that, to save energy, it is mandatory to decrease the turbulence intensity prior to the acceleration phase, thereby delaying the onset of turbulence, in agreement with Scarselli *et al.* (2023). Independently of the parametrisation method considered (uni-/tri-variant or Fourier series), all found optimal waveforms have a steep acceleration phase in common. As noted by Morón *et al.* (2022) in the context of transitional pulsatile pipe flow and Scarselli *et al.* (2023) in the case of fully turbulent flow, at these Womersley numbers, a fast acceleration is related to a rapid turbulence collapse. This is ultimately rooted in the time delay between the driving pressure gradient and the turbulence response, which for these flow parameters was shown to be approximately  $T/4$  (Morón & Avila 2024). After modifying our optimal waveforms with different techniques, we show that changing the acceleration phase slope (and thereby not delaying the onset of turbulence) has the largest detrimental impact on energy consumptions. Compared with a single harmonic waveform, the optimal waveform exhibits a substantially larger time shift between the peak bulk velocity, and the peaks of dissipation and production. Further, in optimal waveforms, both turbulent kinetic energy and production are more radially localised towards the pipe centre, resulting in lower spatially and temporally averaged wall shear stresses.

In conclusion, thanks to the use of Bayesian optimisation, coarse meshes and a state-of-the-art GPU DNS code, we solve the optimisation problem in a broad parametric regime and computationally efficiently. We consider a large space of continuous waveforms by optimising a small number of Fourier coefficients of the driving waveform (see (2.8)). Here, at  $Re = 8600$ , the optimal waveform saves up to 23 % of energy and reduces drag by 37 %, clearly outperforming triangular waveforms at the same Reynolds number (Scarselli *et al.* 2023 reported a power saving of 9 % and a drag reduction of 27 % at this  $Re$ ). We only considered two Womersley numbers ( $Wo \in \{10, 10\sqrt{2}\}$ ) and investigated three average Reynolds numbers ( $Re \in \{4300, 5160, 8600\}$ ). Hence, a large parameter space remains to be explored in the future. In addition, for computational reasons, we restricted the number of Fourier coefficients and their magnitude, imposing a strong locality to the optima. However, the presented savings are already higher than many flow control techniques documented in the literature (e.g. Foggi *et al.* 2023a,b; Scarselli *et al.* 2023). Finally, we note that in practical applications, further steps are needed to realise the power savings. For example, power gains occur in part of the period where the fluid is decelerated and storing this energy will unavoidably result in losses.

**Acknowledgements.** We appreciate fruitful discussions with Dr José M. Lopez (University of Málaga) and the members of the fluid modelling and simulation group at ZARM.

**Funding.** Support from the Deutsche Forschungsgemeinschaft (DFG, German Science Foundation) under grant number 540652448 is gratefully acknowledged.

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

**Data availability statement.** The data that support the findings of this study will be made openly available at <https://doi.pangaea.de/10.1594/PANGAEA.986097>.

### Appendix A. Gradient-based optimisation

In this appendix, we describe the results of using a gradient-based optimisation method. In particular, we consider an ordinary gradient-based sequential least-squares programme (SLSQP, Schittkowski 1982; Kraft 1988) and study its performance with respect to the desired level of standard error  $\zeta^*$ . We implement the version from *scipy*'s *optimise* package.

The SLSQP method is a quasi-Newton method, suited for nonlinear constrained optimisation problems. It follows the basic update strategy  $\boldsymbol{\eta}^{(j+1)} = \boldsymbol{\eta}^{(j)} + \eta^{(j)} \mathbf{s}^{(j)}$ , where  $\eta^{(j)}$  and  $\mathbf{s}^{(j)}$  denote the step size and search direction at iterate  $j$ . In each iteration, the search direction is computed by solving the least squares sub-problems (Schittkowski 1982) arising from a factorisation of a Broyden–Fletcher–Goldfarb–Shanno (BFGS) approximation to the Hessian (Broyden 1970; Fletcher 1970; Goldfarb 1970; Shanno 1970).

The Armijio condition (Armijio 1966) is used in the line search to guarantee an appropriate step size  $\eta^{(j)}$ . A detailed description of the SLSQP method can be found from Schittkowski (1982), Kraft (1988) or Ma *et al.* (2024) and the references therein.

We assessed the performance in terms of feasibility, computational effort and robustness of the SLSQP method using the restricted (uni-variant) triangular waveform described in § 2.4. Here, choosing  $\bar{Re} = 4300$ ,  $Re^+ = 9400$  and  $Re^- = 1600$  (see figure 1a), we aim to identify the  $\bar{\tau}_w$ -optimal acceleration time. The acceleration time is bounded by  $0.01T$  and  $0.68T$  ( $T_a \in \mathcal{Q} = [0.01, 0.68]T$ ) to obtain continuous waveforms and to ensure that  $\bar{Re}$  is realised. The gradient  $\partial\mathcal{J}/\partial T_a$  (used for the computation of the search direction) is approximated using first-order finite differences with a step size of  $2.5 \times 10^{-2}T$ . The optimisation is started with  $\boldsymbol{\eta}^{(0)} = 0.345$  and the optimisation loop is terminated if either  $|\mathcal{J}(\boldsymbol{\eta}^{(j)}) - \mathcal{J}(\boldsymbol{\eta}^{(j-1)})| \leq \epsilon$  or  $\|\mathbf{s}^{(j)}\|_2 \leq \epsilon$ , where the tolerance is chosen as  $\epsilon = 5 \times 10^{-3}$ .

For different values of  $\zeta^*$  and initially  $n = 3$ , the evolution of the function value  $\bar{\tau}_w$  (blue lines) and the corresponding acceleration time (grey lines), as the SLSQP-optimisation loop progresses, are shown in figure 10(a–e). For  $\zeta^* = 2.5\%$  (figure 10a), the gradient at the initial guess points towards larger acceleration times ( $T_a$  is increased and  $\bar{\tau}_w$  decreases). In fact, the gradient is so large that the first line search iteration moves towards the upper bound of  $\mathcal{Q}$  ( $T_a = 0.68T$ ), increasing the function value to  $\bar{\tau}_w = 1.17$ . Subsequent line search iterations find an acceleration time where, compared with the initial value of  $\bar{\tau}_w = 1.025$ , the mean wall shear stress is decreased to  $\bar{\tau}_w = 1.004$ . Even though being close to the initial guess, at  $T_a = 0.355$ , the now gradient points towards shorter accelerations (second  $\nabla$ -marker in figure 10a), which would indicate a minimum between 0.345 and 0.355. Subsequently, the line search fails to find a step that decreases the function value for 11 consecutive function calls and the optimisation loop is terminated. These line searches unveil that the underlying optimisation problem exhibits features of ill-posedness: minor changes in the acceleration time yield large changes in the function value. Thus, gradient approximation and line searches are prone to following those noise-induced changes rather than capturing the overall trend of the function. This motivates the question of how many periods to run to capture the underlying function trend.

For  $\zeta^* = 1.25\%$ , the gradient at  $\boldsymbol{\eta}^{(0)}$  also points towards larger accelerations (if not so strongly), and the line search finds such a small step size that  $\mathcal{J}(\boldsymbol{\eta}^{(0)})$  and  $\mathcal{J}(\boldsymbol{\eta}^{(5)})$  almost coincide and the loop terminates. Reducing  $\zeta^*$  to 0.625% in figure 10(c) leads to an almost zero gradient at the initial guess and the loop is terminated by the  $\|\mathbf{s}^{(j)}\|_2 \leq \epsilon$

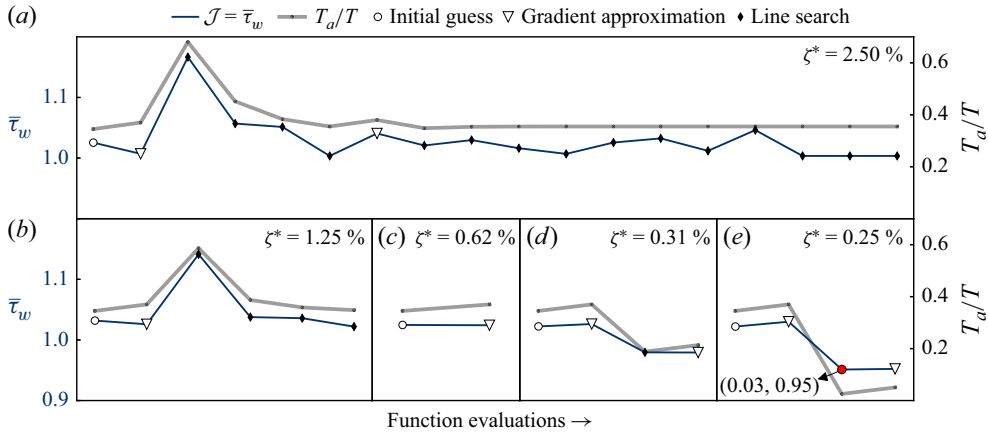


Figure 10. Evolution of the optimisation process for different choices of the admissible standard error  $\zeta^* \in \{2.5, 1.25, 0.625, 0.3125, 0.25\}$  %. Blue lines indicate the wall shear stress obtained when the acceleration time is chosen according to the grey lines (uses right-hand side labels), where gradient-approximations are indicated by downward triangles and line searches are indicated by diamonds.

criterion. For  $\zeta^* = 0.3125\%$  (figure 10d), the gradient points towards shorter acceleration times and the line search finds a suitable step size after only one iteration, indicating the gradient is now capturing the trend of the function. Lastly, for  $\zeta^* = 0.25\%$ , the slope towards smaller acceleration times is even larger, so that the first line search iteration (red marker in figure 10e) finds  $\mathcal{J}(U_{0.025}(t)) = 0.951$ . Even at  $\zeta^* = 0.25\%$ , which in our case is equivalent to running up to 98 periods, the SLSQP method evidently fails to find minima: at  $T_a = 0.01T$ , using  $\zeta^* = 0.25\%$ , a mean wall shear stress of  $\bar{\tau}_w = 0.945$  is obtained, which is statistically significantly smaller than the obtained minimum of  $\bar{\tau}_w = 0.951$  (the 95 % CI of 0.951 is given as  $[0.946, 0.956]$ ). Clearly, further reducing  $\zeta^*$  is not a viable strategy to find ‘true’ minima, as it comes with extensive computational effort. In fact, the four function evaluations from figure 10(e) took 4.5 days. Assuming that the optimisation would converge for a smaller standard error of, for example, 0.125 %, according to (2.6), four times as many periods (approximately 400) are necessary, which would lead to a computational time of 4.5 days per function evaluation. We conclude that, for the multidimensional optimisation problems proposed here, where the number of function evaluations in general scales with  $d^2$ , gradient-based methods like the SLSQP method are computationally infeasible.

Increasing the step size for the finite difference gradient approximation can be a way to overcome noise-induced local minima. However, the chosen value of  $2.5 \times 10^{-2}$  is already relatively large and increasing it further reduces the accuracy substantially (i.e. we would be unable to capture short term trends in  $\mathcal{J}_S$ ).

## Appendix B. Effect of the Womersley number

For  $\overline{Re} = 5160$  and  $Wo = 10\sqrt{2}$ , we show the optimal waveform in figure 11(a). Here, a power saving of 10.0 % is realised, while drag is reduced by 20.7 %. The waveform shows the same characteristics as WF 1–4. The pre-peak velocity is reduced substantially, so that the onset of turbulence is delayed (see figure 11b). As before, quick acceleration needs large instantaneous pressure gradients (figure 11c) and therefore, the instantaneous power input is up to 14 times larger than in steady conditions (figure 11e).

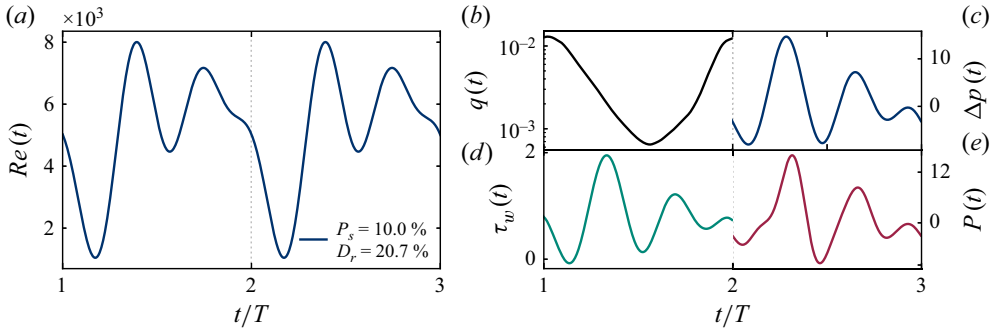


Figure 11. (a) Power-optimal waveform obtained by the truncated Fourier approach ( $\overline{Re} = 5160$ ,  $Wo = 10\sqrt{2}$ ), where  $N = 3$  (see (2.8)) and  $|a_k|, |b_k| \leq 1/6$ ,  $k = 1, 2, 3$  (phase-adjusted). (b) Evolution of the cross-sectional kinetic energy (units of  $\overline{U}^2$ ) and the power input (units of  $\overline{P}_b$ ). (c–e) Evolution of the pressure drop, the wall shear stress and the power input (in units of the steady state counterparts), respectively.

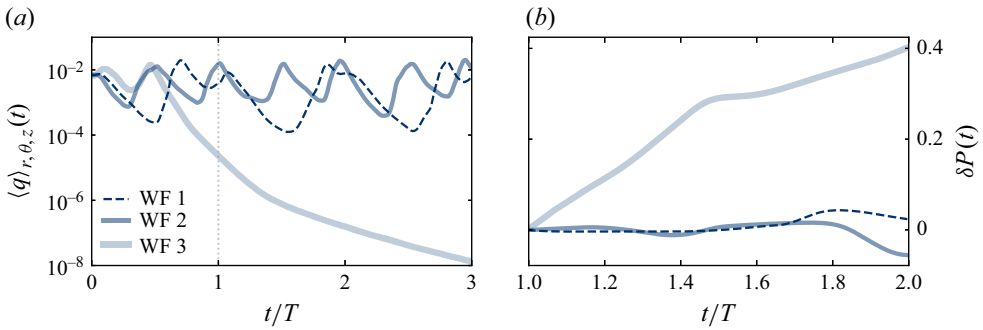


Figure 12. (a) Evolution of the turbulent kinetic energy (in units of  $\overline{U}^2$ ) for waveforms 1–3 ( $\overline{Re} = 5160$ ) without forcing ( $A_f = 0$ ). (b) Time integrated difference of the power inputs  $P(t)$ ,  $\delta P(t)$ , with ( $A_f = 1.15 \times 10^{-3}$ ) and without forcing.

### Appendix C. Robustness of optimal waveforms

First, we analyse the sensitivity with respect to the spatial and temporal resolution. As shown in table 1, for waveform 1–3 ( $\overline{Re} = 5160$ ), we conducted additional simulations with significantly finer resolutions while reducing the time step approximately tenfold (see table 1). We note that our code does not adjust the time step based on the Courant–Friedrichs–Lewy (CFL) number, but based on an error threshold of the predictor–corrector loop for the nonlinear term. The observed changes in power savings are insignificant for any waveform when considering the 95 % CI with the acceptable standard error of 2.5 %. In WF 1, power saving is increased to 14.5 % (originally 11.8 %), while in WF 2, the power saving is nearly identical (10.1 % versus 10.8 % originally) and in WF 3, the saving is slightly reduced to 8.5 % (10.4 % originally). In WF 4, power savings are slightly increased to 23.4 % (originally 22.2 %). We stress that satisfying such strict spatial and temporal resolutions during the whole optimisation process is infeasible: for the  $\overline{Re} = 5160$  case, the single function evaluations take approximately 2.5 days, while they take 9 days for the  $\overline{Re} = 8600$  case.

Second, we examine the effect of the body force (2.3). On our finest grid (see table 1), we have conducted additional simulations of the optimal waveforms 1–3 without any forcing term ( $A_f = 0$ ). In figure 12(a), we show the evolution of the turbulent kinetic energy during the first 3 periods for waveforms 1–3. While in waveforms 1 and 2, turbulence is

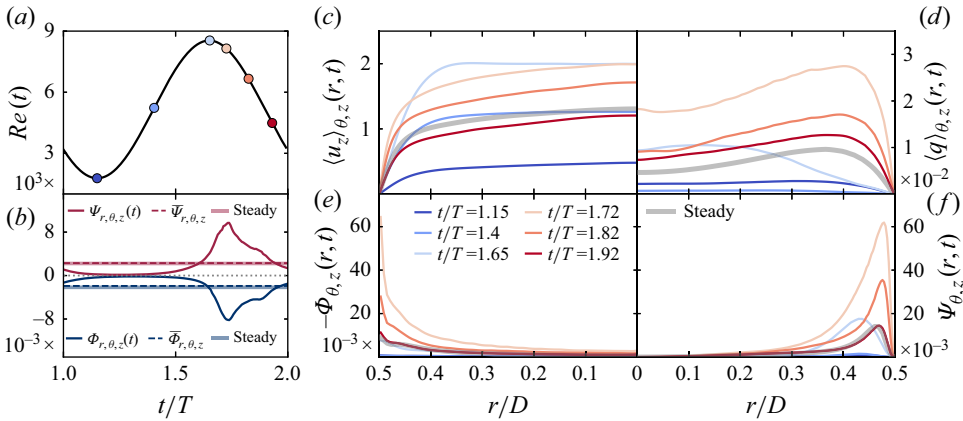


Figure 13. (a) Single harmonic waveform ( $\overline{Re} = 5160$ ,  $Re^+ = 8550$ ,  $Re^- = 1779$ ) where colours encode the times in panels (c)–(f). (b) Time evolution of production ( $\Psi$ ) and dissipation ( $\Phi$ ) in this waveform and steady pipe flow ( $\overline{Re} = 5160$ ) in units of  $(\overline{U}^3/D)$ . (c–f) Axial and azimuthal averaged axial velocity (units of  $\overline{U}$ ), turbulent kinetic energy (units of  $\overline{U}^2$ ), dissipation and production (units of  $(\overline{U}^3/D)$ ).

sustained, WF 3 shows exponential decay of turbulence. Thus, in WF 3, laminar conditions are approached, resulting in an unrealistic power saving of approximately 40%. Drag is reduced by 44%. In the optimisation loop, waveforms that relaminarise would always be preferable over those that do not – even if in reality, they might perform worse. In waveforms 1 and 2, the forcing has minimal impact on power savings: figure 12(b) shows the time integrated difference between the power input  $P(t)$  with and without forcing,  $\delta P(t)$ . In the majority of the period, power inputs with and without forcing are almost equal ( $\delta P(t) \approx 0$ ). Towards the end of the period, they slightly diverge; however,  $\delta P(T)$  (which is the difference  $P_s$  with and w/o forcing) is insignificant ( $\approx 4\%$ ). In WF4 ( $Re = 8600$ ), turbulence is sustained. The change in the power saving is insignificant ( $P_s = 22.1\%$  versus  $P_s = 22.2\%$  originally).

The optimal waveforms are robust with respect to the pipe length. For waveforms 1–3, we conducted simulations on a longer pipe ( $L = 10D$ ). Power savings fall in line with those obtained in the short pipe ( $P_s = 12.6\%$  for waveform 1,  $P_s = 10.9\%$  for WF 2 and  $P_s = 10\%$  for WF 3).

Lastly, we have also confirmed the robustness with respect to different initial conditions. For waveforms 1–3, instead of using an already turbulent initial condition, obtained from a steady run, we initialise the flow laminar ( $u_z = 1 - r^2$ ) and introduce a (strong) localised perturbation of the form

$$\begin{aligned} u_r &= 0.2(1 - r^2)^2 \exp(-10 \sin^2(\pi z/L)) \sin(\theta), \\ u_\theta &= 0.2((1 - r^2)^2 - 4r^2(1 - r^2)) \exp(-10 \sin^2(\pi z/L)) \cos(\theta). \end{aligned} \quad (C1)$$

This unveils that the waveforms are memoryless, in the sense that even strong initial conditions have no influence on the flow behaviour after the first period. Precisely, for waveforms 1, 2 and 3 (see figure 6), we now obtain power savings of 13%, 9.6% and 9.2%, respectively. These differences are insignificant considering the standard error of 2.5% that was used during optimisation.




Figure 14. Instantaneous snapshots of the wall shear stress in units of  $\bar{\tau}_{w,b}$  in steady pipe flow ( $\overline{Re} = 5160$ , left panel) and in the single harmonic waveform for the same phases in the period as in figure 13(a).

#### Appendix D. Spatio-temporal intra-cycle mechanics in sub-optimal waveforms

Here, we show the figures for the intra-cycle mechanism of a single harmonic waveform, namely figures 13 and 14.

#### REFERENCES

- ARMIJO, L. 1966 Minimization of functions having Lipschitz continuous first partial derivatives. *Pac. J. Maths* **16** (1), 1–3.
- AUTERI, F., BARON, A., BELAN, M., CAMPANARDI, G. & QUADRIO, M. 2010 Experimental assessment of drag reduction by traveling waves in a turbulent pipe flow. *Phys. Fluids* **22** (11), 115103–115114.
- AVILA, M., BARKLEY, D. & HOF, B. 2023 Transition to turbulence in pipe flow. *Annu. Rev. Fluid Mech.* **55** (1), 575–602.
- BLASIUS, H. 1913 *Mitteilungen über Forschungsarbeiten Auf Dem Gebiete Des Ingenieurwesens – Das Ähnlichkeitsgesetz Bei Reibungsvorgängen in Flüssigkeiten*. Springer Berlin Heidelberg.
- BROYDEN, C.G. 1970 The convergence of a class of double-rank minimization algorithms 1. General considerations. *IMA J. Appl. Maths* **6** (1), 76–90.
- DING, L., SABIDUSSI, L.F., HOLLOWAY, B.C., HULTMARK, M. & SMITS, A.J. 2024 Acceleration is the key to drag reduction in turbulent flow. *Proc. Natl Acad. Sci. USA* **121** (43), 1–5.
- FELDMANN, D., MORÓN, D. & AVILA, M. 2020 Spatiotemporal intermittency in pulsatile pipe flow. *Entropy* **23** (1), 46.
- FLETCHER, R. 1970 A new approach to variable metric algorithms. *Comput. J.* **13** (3), 317–322.
- FOGGI, R., GIULIO, M., ALESSANDRO, R., MARCO, E. & QUADRIO, M. 2023a On–off pumping for drag reduction in a turbulent channel flow. *J. Fluid Mech.* **966**, 1–12.
- FOGGI, R., GIULIO, M., ALESSANDRO, R., MARCO, E. & QUADRIO, M. 2023b Saving energy in turbulent flows with unsteady pumping. *Sci. Rep.-UK* **13** (1), 1299.
- FRENNING, L. 2001 *Pump Life Cycle Costs — a Guide to LCC Analysis for Pumping Systems*. Hydraulic Institute & Europum.
- GARCÍA-MAYORAL, R. & JIMÉNEZ, J. 2011 Drag reduction by riblets. *Phil. Trans. R. Soc. A Math. Phys. Engng Sci.* **369** (1940), 1412–1427.
- GOLDFARB, D. 1970 A family of variable-metric methods derived by variational means. *Maths Comput.* **24** (109), 23–26.



- GOLDSTEIN, A., KAPELNER, A., BLEICH, J. & PITKIN, E. 2014 Peeking inside the black box: visualizing statistical learning with plots of individual conditional expectation, [arXiv:1309.6392](https://arxiv.org/abs/1309.6392).
- ISHIBASHI, H., KARASUYAMA, M., TAKEUCHI, I. & HINO, H. 2023 A stopping criterion for bayesian optimization by the gap of expected minimum simple regrets. In *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, Proceedings of Machine Learning Research, vol. 206, pp. 6463–6497. PMLR.
- IWAMOTO, K., SASOU, N. & KAWAMURA, H. 2007 *Direct Numerical Simulation of Pulsating Turbulent Channel Flow for Drag Reduction*. Springer Berlin Heidelberg.
- JONES, D.R., SCHONLAU, M. & WELCH, W.J. 1998 Efficient global optimization of expensive black-box functions. *J. Glob. Optim.* **13** (4), 455–492.
- KOBAYASHI, W., SHIMURA, T., MITSUISHI, A., IWAMOTO, K. & MURATA, A. 2021 Prediction of the drag reduction effect of pulsating pipe flow based on machine learning. *Intl J. Heat Fluid Flow* **88**, 108783.
- KRAFT, D. 1988 A software package for sequential quadratic programming. Technical report. <https://books.google.de/books?id=4rKaGwAACAAJ>.
- KRANZ, F. 2024 Simulation-based optimization of driving waveforms in turbulent pipe flow. *Master's thesis*, University of Bremen, Bremen, Germany.
- KUSHNER, H.J. 1964 A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise. *J. Basic Engng* **86** (1), 97–106.
- KÜHNEN, J., SONG, B., SCARSELLI, D., BUDANUR, N.B., RIEDL, M., WILLIS, A.P., AVILA, M. & HOF, B. 2018 Destabilizing turbulence in pipe flow. *Nat. Phys.* **14** (4), 386–390.
- LIU, X., ZHU, H., BAO, Y., SRINIL, N., ZHOU, D. & HAN, Z. 2024 Time-delayed characteristics of turbulence in pulsatile pipe flow. *J. Fluid Mech.* **979**, 1–2.
- LÓPEZ, J.M., FELDMANN, D., RAMP, M., VELA-MARTÍN, A., SHI, L. & AVILA, M. 2020 nsCouette – A high-performance code for direct numerical simulations of turbulent Taylor–Couette flow. *SoftwareX* **11**, 100395.
- MA, Y., GAO, X., LIU, C. & LI, J. 2024 Improved SQP and SLSQP algorithms for feasible path-based process optimisation. *Comput. Chem. Engng* **188**, 108751.
- 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. 2023 Bayesian optimization of wall-normal blowing and suction-based flow control of a NACA 4412 wing profile. *Flow Turbul. Combust.* **113** (1), 93–118.
- MANNA, M., VACCA, A. & VERZICCO, R. 2015 Pulsating pipe flow with large-amplitude oscillations in the very high frequency regime. Part 2. Phase-averaged analysis. *J. Fluid Mech.* **766**, 272–296.
- MARENSE, E., WILLIS, A.P. & KERSEWELL, R.R. 2019 Stabilisation and drag reduction of pipe flows by flattening the base profile. *J. Fluid Mech.* **863**, 850–875.
- MCKEON, B.J. 2010 Controlling turbulence. *Science* **327** (5972), 1462–1463.
- MOCKUS, J. 1972 On bayes methods for seeking an extremum. *Avtomatika i Vychislitel'naja Technika* **3**, 53–62.
- MOCKUS, J. 1994 Application of Bayesian approach to numerical methods of global and stochastic optimization. *J. Glob. Optim.* **4** (4), 347–365.
- 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.
- MORÓN, D. & AVILA, M. 2024 Turbulent puffs in transitional pulsatile pipe flow at moderate pulsation amplitudes. *Phys. Rev. Fluids* **9** (2), 024601.
- MORÓN, D., FELDMANN, D. & AVILA, M. 2022 Effect of waveform on turbulence transition in pulsatile pipe flow. *J. Fluid Mech.* **948**, 1–14.
- MORÓN, D., VELA-MARTÍN, A. & AVILA, M. 2024 Predictability of decay events in transitional wall-bounded flows. *J. Phys.: Conf. Ser.* **2753** (1), 012009.
- QUADRIO, M. & RICCO, P. 2004 Critical assessment of turbulent drag reduction through spanwise wall oscillations. *J. Fluid Mech.* **521**, 251–271.
- REYNOLDS, O. 1883 An experimental investigation of the circumstances which determine whether the motion of water shall be direct or sinuous, and of the law of resistance in parallel channels. *Proc. R. Soc. Lond.* **35**, 84–99.
- SCARSELLI, D., LOPEZ, J.M., VARSHNEY, A. & HOF, B. 2023 Turbulence suppression by cardiac-cycle-inspired driving of pipe flow. *Nature* **621** (7977), 71–74.
- SCHITTKOWSKI, K. 1982 The nonlinear programming method of Wilson, Han, and Powell with an augmented Lagrangian type line search function. *Numer. Math.* **38** (1), 83–114.

- SHANNO, D.F. 1970 Conditioning of quasi-Newton methods for function minimization. *Maths Comput.* **24** (111), 647–656.
- SOUMA, A., IWAMOTO, K. & MURATA, A. 2009 Experimental investigation of pump control for drag reduction in pulsating turbulent pipe flow. In *Proceeding of Sixth International Symposium on Turbulence and Shear Flow Phenomena*, pp. 761–765. Begellhouse.
- THOMAS, C., BASSOM, A.P., BLENNERHASSETT, P.J. & DAVIES, C. 2011 The linear stability of oscillatory Poiseuille flow in channels and pipes. *Proc. R. Soc. A Math. Phys. Engng Sci.* **467** (2133), 2643–2662.