

This is the accepted manuscript made available via CHORUS. The article has been published as:

Reducing the skin-friction drag of a turbulent boundary-layer flow with low-amplitude wall-normal blowing within a Bayesian optimization framework

O. A. Mahfoze, A. Moody, A. Wynn, R. D. Whalley, and S. Laizet

Phys. Rev. Fluids **4**, 094601 — Published 9 September 2019

DOI: [10.1103/PhysRevFluids.4.094601](https://doi.org/10.1103/PhysRevFluids.4.094601)

Reducing the skin-friction drag of a turbulent boundary-layer flow with low-amplitude wall-normal blowing within a Bayesian optimisation framework

O. A. Mahfoze,¹ A. Moody,² A. Wynn,¹ R. D. Whalley,² and S. Laizet¹

¹*Department of Aeronautics, Imperial College London
London, SW7 2AZ, United Kingdom*

²*School of Engineering, Newcastle University,
Newcastle NE1 7RU, United Kingdom*

Abstract

A Bayesian optimisation framework is developed to optimise low-amplitude wall-normal blowing control of a turbulent boundary-layer flow. The Bayesian optimisation framework determines the optimum blowing amplitude and blowing coverage to achieve up to a 5% net-power saving solution within 20 optimisation iterations, requiring 20 Direct Numerical Simulations (DNS). The power input required to generate the low-amplitude wall-normal blowing is measured experimentally for two different types of blowing device, and is used in the simulations to assess control performance. Wall-normal blowing with amplitudes of less than 1% of the free-stream velocity generate a skin-friction drag reduction of up to 76% over the control region, with a drag reduction which persists for up to $650\delta_0$ downstream of actuation (where δ_0 is the boundary-layer thickness at the start of the simulation domain). It is shown that it is the slow spatial recovery of the turbulent boundary-layer flow downstream of control which generates the net-power savings in this study. The downstream recovery of the skin-friction drag force is decomposed using the Fukagata-Iwamoto-Kasagi (FIK) identity, which shows that the generation of the net-power savings is due to changes in contributions to both the convection and streamwise development terms of the turbulent boundary-layer flow.

I. INTRODUCTION

Skin-friction drag reduction is a topic of great interest due to its importance in many engineering applications. Yet despite many decades of extensive research, a practical and affordable method for reducing the turbulent skin-friction drag force in air flows is yet to be found and implemented in real-world applications. Various strategies aiming to achieve this goal have been investigated with some of the most well known targeting either the near-wall turbulence structures through surface topology (García-Mayoral and Jiménez [13]) and wall-based actuation (Choi et al. [6], Quadrio [33], Whalley and Choi [46]), or the larger-scale turbulence structures further away from the wall with jets (Kang et al. [23]) and large-eddy breakup devices (Chin et al. [5]). However, passive approaches often suffer from parasitic drag effects or lose effectiveness in service (Alfredsson and Örlü [1], Spalart and McLean [40]), and the energy expenditure of typical active drag reduction strategies can be very high, often leading to net-power losses even if substantial skin-friction drag reduction is obtained (Quadrio and Ricco [34]).

In the present work, the focus is on the spatial development of a zero-pressure gradient turbulent boundary-layer flow and the resulting wall-friction after control has been applied locally using low-amplitude wall-normal blowing as a drag-reducing strategy. Mass flow injection has been well studied and can be traced back to the 1940's with surface cooling studies by Dawes and Wheeler [8] and Mickley et al. [32]. Rubesin [36] and Torii et al. [45] developed analytical formulas for the calculation of the heat transfer and skin-friction drag coefficients under wall transpiration conditions which were later confirmed experimentally by Simpson et al. [38] with measurements to analyse and model the turbulent momentum and heat transport in the presence of injection and suction. Sumitani and Kasagi [42] performed DNS of a turbulent channel flow with low-intensity (0.1% of the free-stream velocity) uniform wall injection. They observed a skin-friction drag reduction of greater than 10% and concluded that higher wall injection intensities could lead to larger skin-friction drag reduction. A DNS of a zero-pressure gradient turbulent boundary-layer flow was performed by Kim et al. [25] to examine the characteristics of wall pressure fluctuations after the sudden application of wall-blowing and wall-suction. It was found that the small-scale wall pressure fluctuations quickly recovered downstream of control, in contrast to the large-scale wall pressure fluctuations which took a large streamwise distance to return to those found in the usual

canonical wall-turbulent flow. Based on the parametric study of Hwang [18], Kornilov and Boiko [26] designed a wall-blowing control system which delivered compressed air through a perforated flat plate comprised of uniformly distributed sub-millimetre diameter holes. They reported a 70% reduction in the wall-shear stress on application of wall-normal blowing with amplitudes of less than 1% of the free-stream velocity. They estimated net-power saving on the order of 5% due to the energy expenditure of the blowing system. Interestingly, they observed a slow spatial recovery of the skin-friction coefficient downstream of the control region, which hinted at the possibility of employing spatially-discontinuous blowing to achieve comparable skin-friction drag reduction with less input power (Kornilov and Boiko [27]).

The simulations of Kametani and Fukagata [20] and Kametani et al. [21] showed that uniform suction can suppress turbulence but increases the skin-friction drag, while uniform blowing can enhance turbulence but reduce wall-shear stress. These numerical studies also showed that the efficiency of the control increases with widening the streamwise length of the control section. More recently, DNS of a turbulent boundary-layer flow with a low-intensity wall-normal blowing control region have shown a local maximum skin-friction drag reduction of 60%, which persists to tens of boundary-layer thicknesses downstream of control (Stroh et al. [41]). A series of Large-Eddy Simulations (LES) of turbulent boundary-layer flows with wall-normal blowing control were performed by Kametani et al. [22] with a focus on the effect of intermittent blowing along the direction of the flow. By considering only part of the input power required to generate the wall-blowing, namely the pressure difference across the blowing wall, a very optimistic idealised net-power saving of around 18% was predicted.

Employing a reliable optimisation method to determine the optimal parameters of a wall-normal blowing control technique could potentially lead to substantial net-power savings, assuming that a low cost device could be designed to generate the wall-blowing. Bayesian optimisation is a derivative-free algorithm that works efficiently with expensive non-convex objective functions (Gelbart et al. [15]). Bayesian optimisation plays a prominent role in efficiently optimising the parameters of machine learning algorithms, such as Neural Networks, with superior performance when compared to more standard approaches (Brochu et al. [3], Snoek et al. [39]). Bayesian optimisation is yet to be fully exploited for use in fluid flow problems with very few studies combining DNS/LES and Bayesian optimisation to date. Talnikar et al. [44] developed a parallel Bayesian optimisation algorithm for LES to minimise skin-friction drag in a turbulent channel flow with travelling waves and to design

the trailing edge of a turbine blade to reduce turbulent heat transfer and pressure loss. They were able to run several simulations simultaneously, taking advantages of the concurrency offered by supercomputers.

In the present paper, DNS of zero-pressure gradient turbulent boundary layer flows were performed to investigate the potential use of Bayesian optimisation algorithms for wall-turbulence control. The focus is on achieving a skin-friction drag reduction with net-power saving using low-amplitude wall-normal blowing control strategies. For simplicity, three blowing parameters were optimised to achieve the skin-friction reduction with net-power savings. Unlike the majority of other numerical studies where idealised energy usage is assumed, energy savings were evaluated using experimental data. Two Bayesian optimisation studies are presented, the first using the pressure data from the experiments of Kornilov and Boiko [26], and the second using the power requirements from an original low-cost blowing device based on an array of miniature electromagnetic speakers.

II. NUMERICAL METHODS

The incompressible Navier-Stokes equations were solved using a recent version of the high-order flow solver **Incompact3d** (see www.incompact3d.com), adapted to parallel supercomputers using a powerful two-dimensional (2D) domain decomposition strategy (Laizet and Li [29]). This solver is based on sixth-order finite-difference schemes on a Cartesian mesh for the spatial discretization and a semi-implicit time advancement for the viscous terms. To treat the incompressibility condition, a fractional step method requires solution of a Poisson equation, fully solved in spectral space via the use of relevant 3D Fast Fourier transforms. Combined with the concept of the modified wave number (Lele [31]), this direct (i.e. non-iterative) technique allows the implementation of the divergence-free condition up to machine accuracy. A partially staggered mesh is used where the pressure mesh is shifted by a half-mesh from the velocity mesh in each direction. This type of mesh organization leads to more physically realistic pressure fields with no spurious oscillations. More details about **Incompact3d** can be found in Laizet and Lamballais [28]. Note that it has been used recently for DNS of turbulent boundary-layer flows (Diaz-Daniel et al. [10, 11]), including comparisons of wall-shear stress statistics and energy budgets with the reference data of Schlatter and Örlü [37] and Jiménez et al. [19].

Throughout the paper x , y and z denote the streamwise, wall-normal and spanwise directions, and the symbols U and V indicate the instantaneous streamwise and wall-normal velocities, respectively. Fluctuating velocities are represented by lower case symbols (e.g. u) and primed symbols (e.g. u') denote r.m.s values of these fluctuating quantities. An overbar (e.g. \overline{U}) indicates a quantity which is averaged in space and/or in time.

The present simulations were performed for a domain size $L_x \times L_y \times L_z = 750\delta_0 \times 80\delta_0 \times 15\delta_0$ discretized with $n_x \times n_y \times n_z = 3073 \times 513 \times 128$ mesh nodes in the streamwise, wall-normal, and the spanwise directions, respectively. Here, δ_0 is the boundary-layer thickness at the inlet ($x = 0$). A laminar Blasius boundary layer was prescribed at the inlet boundary condition in the streamwise direction, with a Reynolds number of $Re_\theta = U_\infty \theta / \nu = 170$ based on the momentum thickness θ , free-stream velocity U_∞ and kinematic viscosity ν . A 1D convection equation was solved for the outlet boundary condition, where the Reynolds number reaches $Re_\theta \approx 2100$. In the spanwise direction, the boundary conditions were periodic while a homogeneous Neumann condition was imposed on the three velocity components at the top of the domain. The mesh was uniformly spaced in the streamwise and the spanwise directions, and increased in size further from the wall. The resolution in wall viscous units for $Re_\theta = 365$ was $\Delta x^+ = 16.6$, $0.53 \leq \Delta y^+ \leq 135.5$ and $\Delta z^+ = 8$. The simulation time step was $\Delta T = 0.004\delta_0/U_\infty$. Figure 1 illustrates the computational domain and the control region. The random volume forcing designed by Schlatter and Örlü [37] was used to trip the boundary layer and trigger transition-to-turbulence. The forcing volume has a Gaussian distribution with length scales of $1.4\delta_0$, $0.35\delta_0$ and $1.8\delta_0$ in the x , y and z directions, respectively, and a decay time of $1.4\delta_0/U_\infty$. Those parameters were chosen to ensure a smooth and rapid transition to turbulence. The tripping region is located at $x = 3.5\delta_0$ and occupies the full spanwise extent of the domain. The steady low-amplitude wall-normal blowing, v_w , was applied throughout the control region using an inhomogeneous wall-boundary condition.

The numerical set-up and numerical methods were validated by comparing data with the canonical turbulent boundary-layer flow of Schlatter and Örlü [37] and the low-amplitude wall-normal blowing controlled turbulent boundary-layer flow of Stroh et al. [41]. The control region was located at a distance of $x_{Bs} = 68\delta_0$ from the inlet and had a streamwise extent of $L_B = 77\delta_0$, corresponding to $470 \leq Re_\theta \leq 700$ in the canonical flow. For the validation cases, the blowing coefficient $C_B = v_w/U_\infty = 0.005$ and the wall-normal blowing is applied uniformly across the span of the domain and was steady in time. Figure 2(a) shows the




FIG. 1: Schematic of the computational domain. The control region is highlighted in pink.

streamwise evolution of the friction coefficient c_f with and without low-amplitude wall-normal blowing control. Data are averaged in the spanwise direction and in time over $T \sim 8000\delta_0/U_\infty$. The friction coefficient is evaluated using

$$c_f(Re_\theta) = \frac{\overline{\tau_w}(Re_\theta)}{0.5\rho U_\infty^2}, \quad (1)$$

where $\overline{\tau_w}$ is the averaged wall-shear stress and ρ is the density of the fluid. Good agreement with past studies is observed for the streamwise evolution of the skin-friction coefficient with and without the low-amplitude wall-normal blowing control. The current wall-normal blowing simulation (red solid line) captures the same qualitative and quantitative trend as the reference data (Stroh et al. [41], red dashed line), showing a maximum skin-friction drag reduction of 56% at $Re_\theta \approx 650$. Downstream of control there is a sharp recovery of the skin-friction coefficient towards its canonical counterpart; however, a significant level of skin-friction drag reduction persists far downstream. A noticeable difference between the present data and the reference data is the slightly smaller values for the skin-friction coefficient. These small differences may be attributed to the use of different flow solvers, domain sizes and number of mesh nodes across the data sets. Despite the use of different flow solvers, figure 2 (b) shows there is excellent agreement in the canonical turbulent boundary-layer profiles at $Re_\theta = 1000$ (black lines) and $Re_\theta = 1410$ (red lines) when compared with data from Schlatter and Örlü [37].




FIG. 2: (a) Streamwise evolution of the friction coefficient as a function of the Reynolds number for the canonical and controlled turbulent boundary-layer flows. (b) Mean and fluctuating streamwise velocity profiles of canonical turbulent boundary-layer flows at $Re_\theta = 1000$ (black lines) and $Re_\theta = 1410$ (blue lines). Solid lines correspond to the present results and the dashed lines to the reference data of Schlatter and Örlü [37].

III. BAYESIAN OPTIMISATION ALGORITHM

In the present study, a Bayesian optimisation algorithm was used to optimise low-amplitude wall-normal blowing control of a turbulent boundary-layer flow to achieve a skin-friction drag reduction with net-power saving. Generically, Bayesian optimisation algorithms seek to minimise a chosen objective function over a given set of parameter values. Bayesian optimisation algorithms have two stages. First, given knowledge of the objective at a known set of parameters, a probability density function (PDF) for the objective function is computed. This encapsulates a *best guess* of the objective and quantifies the uncertainty of the approximation. Second, an *acquisition function* is minimised to determine the next set of parameter values to be sampled. This typically involves a trade-off between minimising the objective and reducing uncertainty of its approximation.

Specifically, consider an experiment with m input parameters, denoted $\mathbf{x} \in \mathbb{R}^m$, and a scalar-valued objective function $f(\mathbf{x}) \in \mathbb{R}$ that is to be minimised. Suppose that n experiments have been conducted at input values $(\mathbf{x}_i)_{i=1}^n$ and that the objective function

values $f(\mathbf{x}_i)$ are known. Collecting these values as

$$X := \begin{bmatrix} \uparrow & & \uparrow \\ \mathbf{x}_1 & \cdots & \mathbf{x}_n \\ \downarrow & & \downarrow \end{bmatrix} \in \mathbb{R}^{m \times n}, \quad \mathbf{f} := \begin{bmatrix} f(\mathbf{x}_1) \\ \vdots \\ f(\mathbf{x}_n) \end{bmatrix} \in \mathbb{R}^n, \quad (2)$$

a *training set* is defined as $\mathcal{D} := \{X, \mathbf{f}\}$.

The aim is to approximate the value of the objective function at a new *test set* of input values $(\mathbf{x}_i^*)_{i=1}^q \subset \mathbb{R}^m$. To achieve this, a Bayesian optimisation methodology assumes a particular form of Gaussian uncertainty in the relation between the input parameters \mathbf{x} and the objective function value $f(\mathbf{x})$. Under this assumption, the value of the objective at each test point \mathbf{x}_i^* is itself a random variable, which is denoted f_i^* . Letting $\mathbf{f}^* = (f_i^*)_{i=1}^q$ and $X^* := [\mathbf{x}_1^*, \dots, \mathbf{x}_q^*] \in \mathbb{R}^{m \times q}$, it can be shown (Rasmussen [35]) that, given knowledge of the training set \mathcal{D} and chosen test inputs X^* ,

$$(\mathbf{f}^* | X^*, \mathcal{D}) \sim \mathcal{N}(\mu(X^*, \mathcal{D}), \sigma(X^*, \mathcal{D})). \quad (3)$$

That is, the unknown values of the objective function \mathbf{f}^* at testing points \mathbf{x}_i^* have a multivariate normal distribution with mean $\mu = \mu(X^*, \mathcal{D})$ and covariance matrix $\sigma = \sigma(X^*, \mathcal{D})$. This distribution is commonly referred to as the *posterior*. Its mean and covariance are given by

$$\mu = K(X^*, X)K(X, X)^{-1}\mathbf{f}, \quad (4)$$

$$\sigma = K(X^*, X^*) - K(X^*, X)K(X, X)^{-1}K(X, X^*), \quad (5)$$

where $K(A, B) \in \mathbb{R}^{\ell \times p}$ is a kernel matrix calculated from inputs $A \in \mathbb{R}^{m \times \ell}, B \in \mathbb{R}^{m \times p}$. In this study, the elements of $K(A, B)$ are chosen to be the *Matérn 5/2* kernels

$$K(A, B)_{ij} = k(\mathbf{a}_i, \mathbf{b}_j) = \theta_0 \left(1 + \sqrt{\frac{5\|\mathbf{a}_i - \mathbf{b}_j\|_2^2}{l^2}} + \frac{5}{3} \frac{\|\mathbf{a}_i - \mathbf{b}_j\|_2^2}{l^2} \right) \exp \left(-\sqrt{\frac{5\|\mathbf{a}_i - \mathbf{b}_j\|_2^2}{l^2}} \right) \quad (6)$$

for $i = 1, \dots, \ell, j = 1, \dots, p$ ($\mathbf{a}_i, \mathbf{b}_i$ are columns of A and B). Here, θ_0 is the covariance amplitude, and l is a length scale which determines the smoothness of the posterior. Note that different length scales can be assigned to each parameter independently.

The second stage of the Bayesian optimisation algorithm is to select the next sample point \mathbf{x}_{n+1} , given the training set \mathcal{D} and the computed posterior distribution \mathbf{f}^* . This is performed by considering an acquisition function $a(\mathbf{x})$, which trades off *exploitation* (to

select the sample of the lowest mean) and *exploration* (to sample from a region of high uncertainty), by computing

$$\mathbf{x}_{n+1} = \operatorname{argmin}\{a(\mathbf{x}) : \mathbf{x} \in X^*\}. \quad (7)$$

Many acquisition functions have been proposed in the literature (Brochu et al. [3], Snoek et al. [39]), although there is little consensus regarding the optimal choice of such a function. The acquisition function used in this study is the *expected improvement* (EI) function. It takes into account the probability of improvement, and the magnitude of the expected improvement, with respect to the best known value of the objective function $f_{\text{best}} := \min_{\mathbf{x} \in X} f(\mathbf{x})$, given knowledge of \mathcal{D} . In particular,

$$a(\mathbf{x}) = \sigma(\mathbf{x}|\mathcal{D}) [\gamma(\mathbf{x})\Phi(\gamma(\mathbf{x})) + \phi(\gamma(\mathbf{x}))], \quad \mathbf{x} \in X^*, \quad (8)$$

where

$$\gamma(\mathbf{x}) = \frac{\mu(\mathbf{x}|\mathcal{D}) - f_{\text{best}}}{\sigma(\mathbf{x}|\mathcal{D})}, \quad \mathbf{x} \in X^*, \quad (9)$$

and $\phi(\cdot)$ and $\Phi(\cdot)$ are the PDF and the cumulative distribution function (CDF), of a standard $\mathcal{N}(0, 1)$ distribution, respectively.

For a simple illustrative example of locating a local minimum with a Bayesian optimisation algorithm, consider a 1D problem that has a noise-free objective function, $f(x) = \sin(x)/x$, for $-2 \leq x \leq 1$. In this example, the fitting parameters are fixed with $\theta_0 = 1$ and $l = 0.8$. Starting with an arbitrary input parameter $x_1 = -1.5$ and corresponding training set $\mathcal{D} = \{-1.5, f(-1.5)\}$, figure 3 shows the development of the posterior distribution over five iterations of the Bayesian optimisation algorithm. Near the training points (black solid circles), the posterior mean μ (dark blue curve) and the true objective function (black curve) match and the posterior covariance σ (blue shaded area) vanishes; conversely, the uncertainty of the predictive model increases with distance from the observation points: see figure 3 (a). Using the EI acquisition function, the first new input parameter x_2 shown in figure 3 (b) corresponds to the lowest value of the posterior mean. Subsequently, since x_2 is close to the lowest value of the new posterior mean (almost unchanged from the previous figure), the acquisition function selects a testing point x_3 in a region of high uncertainty in figure 3 (c). This behaviour is repeated in figure 3 (d), after which the minimum of the posterior mean approximately coincides with the true optimum in figure 3 (e).

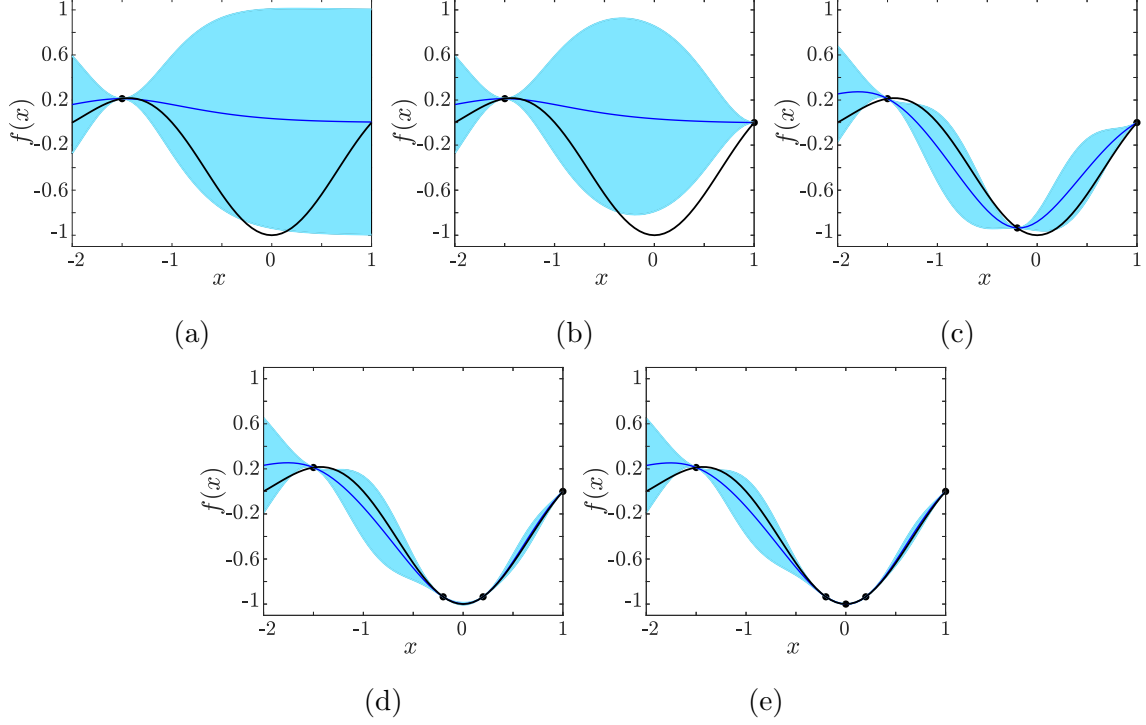


FIG. 3: An example of using Bayesian optimisation to locate a local minimum on a 1D toy problem. The black curve is the true objective function, black solid circles are the observed points, the dark blue curve is the posterior mean μ and the blue shaded area is the posterior uncertainty ($\mu \pm \sigma$).

IV. RESULTS

Small-amplitude wall-normal blowing was applied to the turbulent boundary-layer flow, for a range of blowing amplitudes and coverages, in search of a wall-normal blowing control strategy which would yield a skin-friction drag reduction with a net-power saving. The location and spatial extent of the control domain was the same as the validation case shown in figure 1. The Bayesian optimisation algorithm was used to search over three control parameters simultaneously. The parameters chosen were the wall-normal blowing coefficient, C_B , the number of streamwise blowing areas, N_B and the blowing coverage coefficient, $\alpha := \lambda_1 N_B / L_B$. These are indicated schematically in figure 4. Steady wall-normal blowing was applied uniformly in the spanwise direction, across the full extent of the control area, with λ_1 and λ_0 denoting the streamwise extent of blowing and non-blowing sections, respectively. Wall-normal blowing was imposed smoothly at the start and end of each blowing

section for a streamwise distance of $x \approx 3.5\delta_0$ using a hyperbolic tangent function. To account for these transition regions, where C_B gradually increases and decreases, the constraint $\lambda_1 > 3.5$ was added to the optimisation routine. To ensure that the Bayesian optimisation framework converged within a relatively small number of simulations, the wall-normal blowing amplitude was limited to $0 \leq C_B \leq 0.01$, and the number of blowing areas along the streamwise direction was limited to $1 \leq N_B \leq 10$. These two parameters were coupled with a range in blowing coverage coefficient of $0.045 \leq \alpha \leq 1$ meaning that both a uniform wall-normal blowing region and intermittent wall-normal blowing regions (i.e. wall-normal standing waves) are contained in the set of possible blowing strategies. These parameter ranges were chosen to coincide with recent successful applications of low-amplitude wall-normal blowing control (Kametani et al. [22], Stroh et al. [41]). Selecting the Bayesian optimisation hyper-parameters, i.e. l and θ_0 in equation 6, is challenging with no prior knowledge. Experimenting with a 3D toy problem, similar to the 1D example discussed earlier, showed that using a dynamic θ_0 leads to more reliable convergence when compared to a study with a constant θ_0 . Therefore in the Bayesian optimisations which follow, θ_0 was dynamically changed to be equal to the standard deviation of the newly evaluated objective functions. Similarly, the length scale l for C_B , N_B and α were set to 0.005, 5 and 0.5, respectively, which corresponds to half of the search space in each parameter. A poor choice for the hyper-parameters can result in slow convergence of the optimisation routine.




FIG. 4: Schematic of the boundary condition of the wall-normal velocity in the control region for the Bayesian optimisation studies.

The optimisation process was parallelised by running up to four simulations simultaneously, requiring up to 8,192 computational cores. Each simulation ran for up to 48 hours to provide sufficiently converged data to evaluate the objective function (i.e. the net-power saving) for each set of parameters trialled by the Bayesian optimisation framework. When running up to four simulations at the same time during the Bayesian optimisation study, the

true objective function for each simulation was not available to select the parameters for the next iteration. Therefore, the Bayesian optimisation algorithm used the mean value of the posterior as a proxy for the mean value of the true objective function. Initially, each simulation ran for a duration of $T = 1600\delta_0/U_\infty$. During the first half of each simulation, the flow adapts to the wall-normal blowing boundary condition. In the latter half of each simulation, data were collected to provide an estimate of the time-averaged skin-friction coefficient and thereby net-power saving to within an uncertainty of $\pm 1\%$. If the preliminary estimation of the net-power saving was less than 2.5%, data were collected over an additional $1600\delta_0/U_\infty$, which reduced the uncertainty in the time-averaged skin-friction coefficient and net-power saving to within $\pm 0.5\%$. For cases of interest, when the Bayesian optimisation framework converged to an optimum solution, for example, data were collected over $T = 9000\delta_0/U_\infty$, which reduced the uncertainty in the time-averaged skin-friction coefficient and net-power saving to within $\pm 0.1\%$.

The net-power saving (S) generated by each wall-normal blowing control strategy was assessed by taking into account the input power required to generate the wall-normal blowing plus any power saving due to a reduction in skin-friction drag. Due to the long-lasting downstream effects of the low-amplitude wall-normal blowing control, a global skin-friction drag coefficient ($\overline{C_f}$) was evaluated over a streamwise distance of $L = 615\delta_0$, such that,

$$\overline{C_f} = \frac{1}{L} \int_{35\delta_0}^{650\delta_0} \overline{c_f}(x) dx, \quad (10)$$

$$S = \frac{\overline{C_w} - \overline{C_{f0}}}{\overline{C_{f0}}}. \quad (11)$$

Here, $\overline{C_{f0}}$ is the global skin-friction coefficient of the canonical turbulent boundary-layer flow, and $\overline{C_w} = \overline{C_{fb}} + C_{wb}$. $\overline{C_{fb}}$ is the global skin-friction coefficient with wall-normal blowing control and C_{wb} is a coefficient to take into account the power required to generate each wall-normal blowing control strategy.

In the sections which follow, results are presented for two different Bayesian optimisation studies. In each study, wall-normal blowing control is applied in exactly the same way, as described above. However, in the first study, the input power required to generate the wall-normal blowing, and therefore assess control performance within the Bayesian optimisation framework, is predicted using the experimental data from Kornilov and Boiko [26]. In the

second study, the input power required to generate the wall-normal blowing is measured experimentally by generating low-amplitude blowing with a miniature electromagnetic speaker in quiescent air. Different values of input power cause a different response of the posterior function within the Bayesian optimisation framework, leading to different optimum control solutions for each blowing device.

A. First Bayesian optimisation Study

Kornilov and Boiko [26] showed that low-amplitude wall-normal blowing with $C_B = 0.00287$ can reduce the skin-friction drag of a turbulent boundary-layer flow by up to 70% at a Reynolds number of $Re_\theta \sim 10^3$. In these experiments, data were collected with hot-wire anemometry directly above the region of blowing, and the low-amplitude wall-normal velocity was generated by blowing compressed air through a perforated plate. The perforated plate had an area of $420 \times 250 \text{ mm}^2$ and was 1.1 mm thick. The plate was populated with 0.17 mm holes providing a porosity of 17.1 % and a hole-to-plate-thickness ratio of 6.47. Alongside the hot-wire anemometry measurements, the pressure coefficient $C_p := \Delta p / 0.5 \rho U_\infty^2$ across the perforated plate was measured at a free-stream velocity of $U_\infty = 21 \text{ m/s}$ and was shown to be directly proportional to the wall-normal blowing amplitude, $C_p = 124 C_B$. These experimental data were used to estimate the power consumption required to generate the wall-normal blowing in this first Bayesian optimisation study. It was assumed that a constant and uniform wall-normal velocity boundary condition is a suitable representation for the wall-normal blowing through a perforated plate. In addition, it was assumed that the blowing power coefficient was $C_{wb} = (C_p C_B + C_B^3) \alpha_t$, where α_t is the ratio of the blowing areas to the total control area used to calculate the global skin-friction coefficient, $\overline{C_f}$. Most numerical studies assume that C_p is equal to zero, for example Kametani et al. [21]. Neglecting C_p makes the blowing power proportional to the cube of the blowing velocity, hence virtually zero. Under this assumption, the net-power saving corresponds to the global drag reduction. However, C_p could be significant and has potential to generate net-power losses. It should be noted that the cost of compressing the air is not included in the energy budget here, as this was unknown in the experiments of Kornilov and Boiko [26]. Therefore even with the contributions of C_p included, any estimation on net-power saving may be overestimated.

In this first Bayesian optimisation study, 18 DNS with wall-normal blowing control were conducted for comparison with the canonical turbulent boundary-layer flow. The blowing parameters for the first three simulations (with each simulation also referred to as a Case herein) were selected arbitrarily to pre-test the behaviour of the wall-normal blowing and to initialize the Bayesian optimisation. The choice of parameters for the three initialization cases can influence the search path and rate of convergence of the Bayesian Optimisation study; however, the final optimised result is independent of the initialization cases, assuming a single local minimum over the parameter space search. The parameters of the remaining 15 simulations were determined by the Bayesian optimisation algorithm. Table I shows the parameters chosen by the Bayesian optimisation framework for each simulation, alongside the corresponding maximum and global skin-friction drag reduction, and the net-power saving. Case 0, highlighted black in table I, corresponds to the canonical turbulent boundary-layer flow. The Bayesian optimisation framework predicts that the optimum control strategy achieves a maximum drag reduction of 37% and a net-power saving of 5% with low-amplitude uniform blowing. These set of control parameters are highlighted blue in table I, and are obtained after 13 simulations within the Bayesian optimisation framework. As anticipated, increased wall-normal blowing intensity yields increased maximum and global drag reduction but at the expense of a net-power loss: see Case 5, which is highlighted red in table I. Here, the local maximum drag reduction peaks at 76% but with a net-power loss of 9.5%.

Figure 5 shows the development of the turbulent boundary-layer flow during and downstream of control for Case 5 (red line) and Case 13 (blue line). The canonical turbulent boundary-layer flow (Case 0) is indicated by the black line. Figure 5 (a) shows the local skin-friction coefficient as function of Re_θ . Substantial drag reduction over the control region is observed for these two cases of interest, with the ringing on the skin-friction coefficient for Case 5 due to the intermittent blowing, or standing wave control strategy located by the Bayesian optimisation framework. One of the most striking observations is the long-lasting downstream effects of the wall-normal blowing on the skin-friction coefficient. It is the slow spatial recovery of the turbulent boundary-layer flow back to its canonical value which results in the net-power savings for Case 13. A skin-friction drag reduction persists beyond $Re_\theta = 1940$, which is equivalent to streamwise distance downstream of control of $\sim 650\delta_0$. Accompanying the long-lasting reduction in skin-friction drag is a thickening of the turbulent boundary-layer flow. This is shown in figure 5 (b) where the growth in boundary-layer

TABLE I: Wall-normal blowing parameters for the first Bayesian optimisation study with corresponding maximum local skin-friction drag reduction (%), global skin-friction drag reduction (%) and net-power saving (%). The blue row indicates the optimum control strategy to achieve a net-power saving. The red row indicates the wall-normal blowing parameters to achieve the highest local maximum skin-friction drag reduction. The black row indicates the canonical turbulent boundary-layer flow. For convenience, drag reduction and net-power saving are positive.

Case	C_B	N_B	α	Max Drag Red.	Global Drag Red.	Power Saving S
0	0	—	—	—	—	—
1	0.005	1	1	52	13.2	3.2
2	0.005	1	0.47	48.5	6.0	1.3
3	0.005	1	0.24	49	2.9	0.6
4	0.0037	5	0.9	39.3	9.3	4.3
5	0.00997	10	0.78	75.5	19.8	-9.5
6	0.00032	1	1	39.5	0.9	0.9
7	0.00424	3	0.95	45	11.0	3.8
8	0.00404	6	0.88	41.4	9.6	3.1
9	0.00233	8	0.82	39.5	5.4	3.3
10	0.00136	1	0.06	39.5	0.4	0.3
11	0.000003	10	0.78	39.5	-0.5	-0.5
12	0.0059	10	0.78	54	12.3	1.5
13	0.00289	1	1	36.5	8.3	5.0
14	0.00386	1	1	41.5	10.4	4.5
15	0.00296	1	1	39.5	8.4	4.8
16	0.0029	1	1	39.5	8.1	4.8
17	0.00289	1	1	39.5	7.9	4.6
18	0.00278	1	1	39.5	7.6	4.5

thickness is represented by $Re_\delta = \delta U_\infty / \nu$. With lower blowing amplitude, the distortion of the boundary-layer thickness at the onset of control is less prominent. This result was observed by Stroh et al. [41] and was thought to be due to a streamwise shift in the virtual origin of the turbulent boundary-layer flow.




FIG. 5: (a) Streamwise evolution of the skin-friction coefficient as a function of Re_θ for all cases. The thick solid lines corresponds to the three cases of interest in this study: Case 0 (black), Case 5 (red) and Case 13 (blue). (b) Streamwise evolution of Re_δ for Case 0, Case 5 and Case 13. The region of wall-normal blowing control is highlighted in pink.




FIG. 6: Instantaneous visualisations of the streamwise velocity in the streamwise-spanwise plane at $y^+ \approx 10$ for Case 0, Case 5 and Case 13. Dark blue corresponds to a zero velocity and red correspond to U_∞ .

Instantaneous visualisations of the instantaneous streamwise velocity in the streamwise-

spanwise plane at $y^+ \approx 10$ (over the control area) are plotted in figure 6 for Case 0, Case 5 and Case 13. Over the control region, the wall-normal blowing increases the number of areas where the streamwise velocity is close to zero (dark blue), especially for Case 5 where the blowing intensity is around three times larger than Case 13. These near-wall visualisations also highlight the long-lasting downstream signature of the wall-normal blowing control. Persistent areas with low-speed streamwise velocity can be seen far downstream of control, with fewer areas of high-speed streamwise velocity observed when compared with the canonical turbulent boundary-layer flow (Case 0). However, the injection of energy into the turbulent boundary-layer generates a surge in the turbulence activity as observed by the increased number of vortical flow structures in the developing enstrophy fields shown in figure 7. These results are consistent with the simulations of Kametani and Fukagata [20] and Kametani et al. [21].




FIG. 7: 3D instantaneous visualisations of the enstrophy field for Case 0, Case 5 and Case 13, coloured by the wall-normal velocity component.

Fukagata et al. [12] derived the so-called Fukagata-Iwamoto-Kasagi (FIK) identity, which is an expression for the skin-friction coefficient for incompressible flows. For turbulent boundary layers which are homogeneous in the spanwise direction, the FIK identity can be expressed as

$$c_f^{FIK}(x) = \underbrace{\frac{4(1 - \delta^*)}{Re_\delta}}_{c_f^\delta} + 4 \underbrace{\int_0^1 (1 - y)(-\overline{u'v'})dy}_{c_f^T} + 4 \underbrace{\int_0^1 (1 - y)(-\bar{U}\bar{V})dy}_{c_f^C} - 2 \underbrace{\int_0^1 (1 - y)^2 \left(\frac{\partial \bar{U}\bar{U}}{\partial x} + \frac{\partial \overline{u'u'}}{\partial x} - \frac{1}{Re_{\delta_0}} \frac{\partial^2 \bar{U}}{\partial x^2} + \frac{\partial \bar{p}}{\partial x} \right) dy}_{c_f^D}, \quad (12)$$

where all length and velocity scales are non-dimensionalised by the local boundary-layer thickness or free-stream velocity, respectively. Here δ^* denotes the normalised displacement thickness. The FIK identity decomposes the skin-friction coefficient into four contributions: a boundary-layer thickness contribution c_f^δ , a Reynolds shear stress contribution c_f^T , a mean wall-normal convection contribution c_f^C and a spatial development contribution c_f^D (based on four terms, $c_f^D = c_f^{D1} + c_f^{D2} + c_f^{D3} + c_f^{D4}$). Note that the mean wall-normal convection term and the spatial development term are absent in fully developed channel and pipe flows. For a more detailed discussion on each term of the FIK identity, the reader is referred to Kametani and Fukagata [20].

The FIK decomposition shows that the large levels of skin-friction drag reduction generated over the control region are associated with large negative contributions of the convection term. This is somewhat expected given that $c_f^C \sim -\bar{U}\bar{V} < 0$, with steady wall-normal blowing within a boundary-layer flow. However, an intriguing observation is found when comparing c_f^C downstream of control for cases 5 and 13 for which $C_B \approx 1$ and 0.3%, respectively. For Case 5, figure 8 (b), c_f^C becomes positive downstream of control which may be a consequence of the relatively large wall-normal velocity rolling-up within the boundary layer as the flow advects downstream. A similar small peak of positive c_f^C is observed in figure 8 (c), with much lower wall-normal blowing amplitude, however $c_f^C < 0$ with increasing distance downstream of control, as seen in the canonical turbulent boundary-layer flow, figure 8 (a).

A very similar trend can also be observed for c_f^D : this term is always positive over the control region with a very sharp increase at the start of control followed by a steady decay, which then becomes negative for Case 5 and positive for Case 13 (after a small region where it is negative just at the end of the control region). Therefore, it is possible to conclude that over the blowing region the reduction of the convection contribution is largely responsible for the skin-friction drag reduction, while the spatial development contribution is largely responsible for the skin-friction drag reduction downstream of control.

Figure 9 shows the individual contributions for each term of the spatial development contribution of the FIK identity for Case 0, Case 5 and Case 13. In agreement with previous simulations in a similar set-up (Kametani and Fukagata [20]), adverse and favourable pressure gradients (black lines) can be observed at the start and at the end of each blowing area. This shows that it is important to take into account the contribution of the pressure

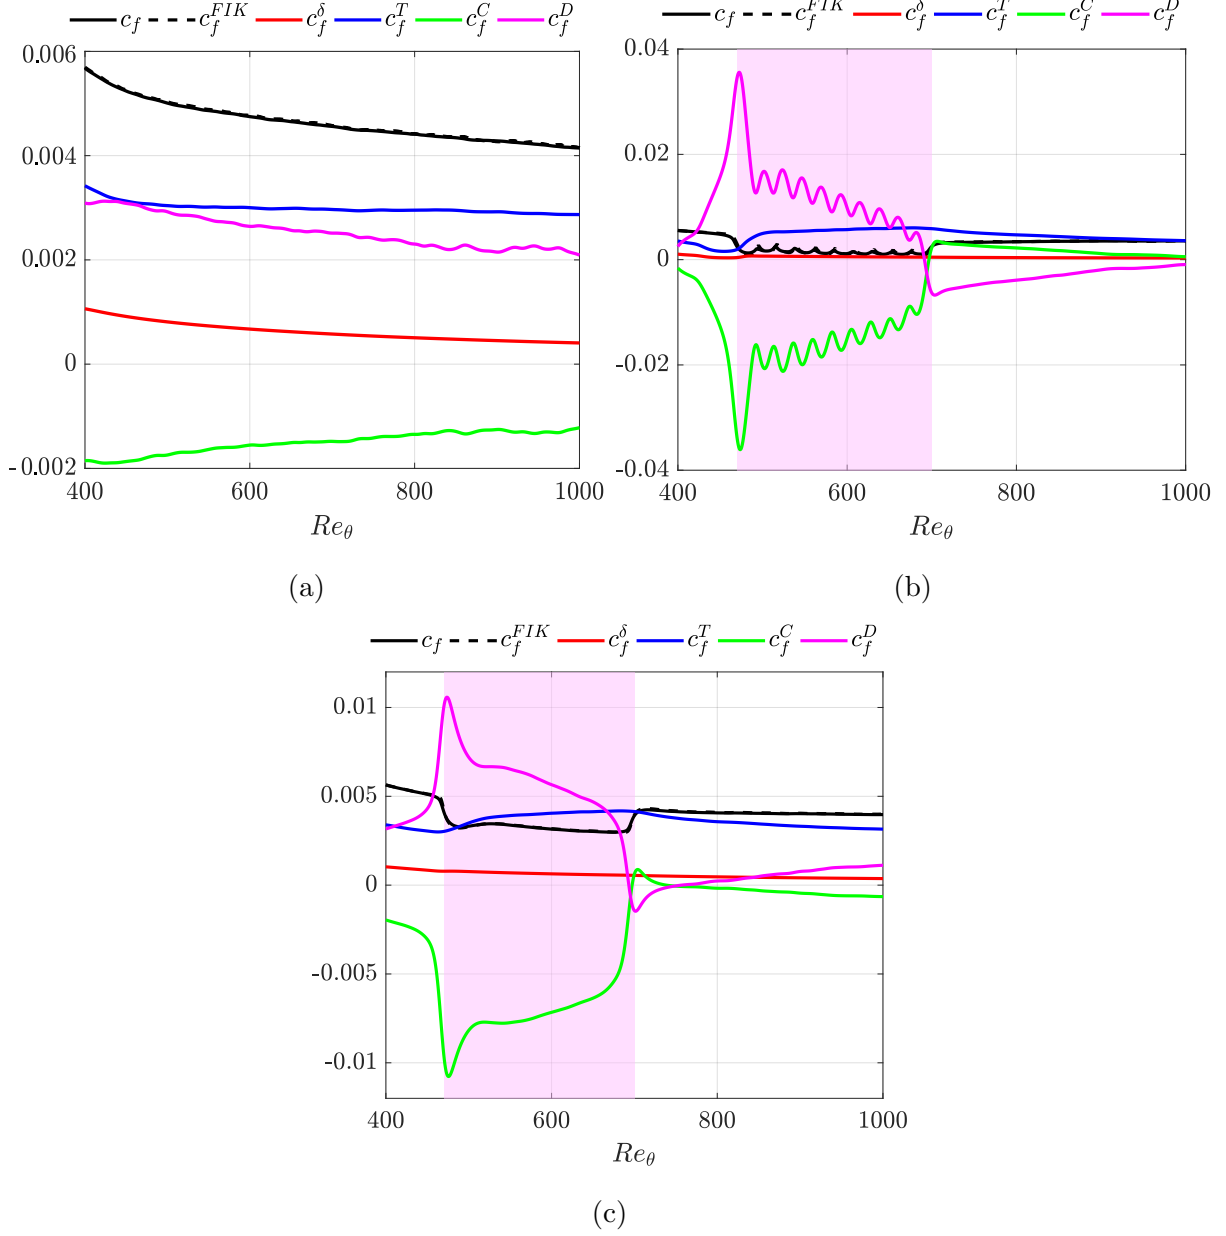


FIG. 8: Streamwise evolution of the FIK identity contributions for Case 0 (a), Case 5 (b) and Case 13 (c).

gradients in the FIK identity. For all cases, the spatial development contribution c_f^D is clearly dominated by the spatial development of the streamwise velocity gradient. c_f^{D2} and c_f^{D3} are virtually zero for the canonical case and very small for the other cases. Interestingly, c_f^D increases downstream of the wall-blowing with positive values for case 13 and negative values for case 5, as already observed in figure 8.




FIG. 9: Streamwise evolution of the four terms of the spatial development contribution of the FIK identity for Case 0 (a), Case 5 (b) and Case 13 (c).

B. Second Bayesian optimisation Study

In this second Bayesian optimisation study, a more realistic estimate of the net-power savings associated with a wall-normal blowing skin-friction control strategy is investigated. The total input power required to generate wall-normal blowing with a miniature electromagnetic

speaker is used to assess control performance in the turbulence resolving simulations within the Bayesian optimisation framework. It is assumed that a steady uniform velocity boundary condition is sufficient to model the wall-blowing generated by an array of miniature electromagnetic speakers.

The wall-normal blowing was generated by a Visaton K 50 miniature electromagnetic speaker. The speaker had a 50 mm diameter diaphragm which was covered with a 3 mm thick perforated acrylic plate. The plate contained 2020 holes which were $460\ \mu\text{m}$ in diameter, providing a porosity of 17.1% and plate-thickness-to-hole-diameter ratio of 6.52: see figure 10 (a). The plate dimensions were chosen to closely match previous successful wall-normal blowing studies which had optimised the blowing geometry through a series of parametric tests (Hwang [17], Kornilov and Boiko [26]). The miniature speaker was driven with a square wave ranging from 0.5 - 5 volts at excitation frequencies ranging from 400 - 500 Hz with a 50% duty cycle. To enable instantaneous measurements of the input power to the miniature speaker, an AD820 operational amplifier, connected in a non-inverting configuration, was used to amplify the voltage drop over a Caddock MP930 0.1Ω current sense resistor: see figure 10 (b). The blowing velocity generated by the miniature speaker was measured with a Dantec 55P15 hot-wire probe powered by a TSI IFA 100 constant temperature anemometry system. The hot-wire voltage was amplified with a gain of 10 prior to sampling. Before and after each experimental run, the hot-wire probe was calibrated against a Dantec ConfortSense 54N95 probe, accurate to $\pm 0.02\ \text{m/s}$, in a low-speed wind tunnel. Each calibration consisted of 15 points over a 0 - 1.6 m/s speed range, which were sampled for 30 seconds at 1 kHz and fit with a third-order polynomial. During each calibration the air temperature was measured to an accuracy of $\pm 0.04^\circ\text{C}$ with an Omega PT100 powered by an Omega PT-104A DAQ module. All blowing measurements were conducted in quiescent air in the centre of a wind tunnel test section. The test section had a cross-sectional area of $0.43 \times 0.35\ \text{m}^2$ and was sealed over a 1.3 m length to create a chamber which was not influenced by any external drafts. During blowing measurements, the instantaneous hot-wire voltage, speaker input voltage waveform, speaker input current and air temperature were sampled simultaneously with a National Instruments PXIe-6356 16-bit DAQ system. The hot-wire data were linearly interpolated across calibration curves to account for any small drifts in temperature during the course of the experiments. Data were sampled at a frequency 10 times greater than the speaker excitation frequency for 120

seconds at each spatial location. The hot-wire probe was positioned and traversed across the face of the miniature speaker by a custom-designed 3-axis traverse system which consisted of three hybrid stepper motors controlled by an National Instruments PXI-7334 motion controller via a LabView interface. Renishaw optical linear encoders measured the position of the hot-wire probe to an accuracy of $0.5 \mu m$.




FIG. 10: (a) Picture of the Visaton K 50 miniature speaker used to generate wall-normal blowing. (b) schematic circuit diagram for the miniature speaker setup.

An example of instantaneous wall-normal blowing velocity taken in the centre of the speaker at around 5 hole diameters downstream from the jets exits (e.g. $y/d \sim 5$, where y is the wall-normal distance from the perforated plate and d is the hole diameter) is shown in figure 11 (a). Plotted alongside the velocity is instantaneous input power to the speaker. Here, the speaker is driven with a square wave voltage input of 4 volts at 400 Hz with a 50% duty cycle. The speaker diaphragm moves up and down on each rising and falling edge of the voltage waveform, respectively, causing a sinusoidal velocity response downstream from the porous plate, similar in working principle to a synthetic jet (Glezer and Amitay [16]). During all measurements, the hot-wire probe is sufficiently far downstream of the jets exits to avoid the effects of any reverse flow entrained into the cavity during the down stroke of the diaphragm (Di Cicca and Iuso [9], Glezer and Amitay [16]). The time-averaged wall-normal blowing velocity at $y/d \sim 5$ is shown in figure 11 (b). The input power-velocity curves follow a non-linear response with greater efficiency in operation found for a driving frequency of around 440 Hz. Either side of this driving frequency, more power is expended driving the diaphragm to generate the same averaged wall-normal blowing velocity. This

optimised frequency is likely due to a coupled resonance that depends on the cavity flow, cavity geometry and structural characteristics of the diaphragm (Glezer and Amitay [16]).

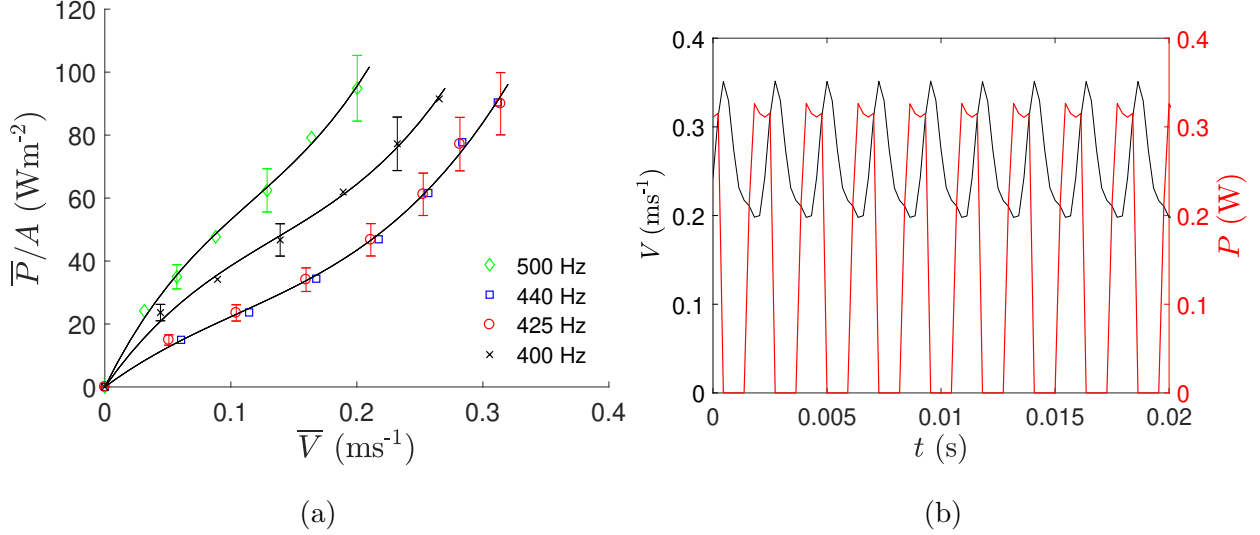


FIG. 11: (a) Instantaneous wall-normal velocity and input power to the speaker driven at 4 volts, 440 Hz at 50% duty cycle. (b) Time-averaged wall-normal velocity versus power per unit area from 1.5-5 volt over 400-500 Hz at 50% duty cycle. All data taken at the centre of the miniature speaker at $y/d \sim 5$.

In this second Bayesian optimisation study, the free-stream velocity is assumed to be $U_\infty = 21 \text{ m/s}$ to match the experimental data from Kornilov and Boiko [26] which were used in the first Bayesian optimisation study. In the turbulence resolving simulations which follow, the net-power saving is calculated as described previously using the 440 Hz power curve shown in figure 11 (b) with the blowing power coefficient defined as

$$C_{wB} = \frac{\bar{P}_B}{\frac{1}{2}\rho U_\infty^3 A_t}. \quad (13)$$

Here \bar{P}_B is the time-averaged total input power to each miniature speaker as shown in figure 11 (b), ρ is the density of the fluid and A_t is the total area over which the global drag reduction is estimated.

In this second Bayesian optimisation study, 11 DNS with wall-normal blowing control were conducted. The first three simulations used the same blowing parameters as the first Bayesian optimisation study to pre-test the behaviour of the wall-normal blowing control and to initialise the Bayesian optimisation framework. Table II shows the parameters chosen

by the Bayesian optimisation framework as it searches through parameter space to achieve a global skin-friction drag reduction with a net-power saving. After 6 iterations the Bayesian optimisation framework predicts that a short intense uniform blowing strategy will achieve a local maximum skin-friction drag reduction of 60.5% and a global skin-friction drag reduction of 4.1% with a net-power saving (S) of 0.1%. These parameters are highlighted blue in table II. To test the sensitivity of the net-power savings to input power from the blowing device, the lower uncertainty limit values from the power per unit area curve shown in figure 11 (b) were used to re-calculate the net-power savings for the 11 iterations found by the Bayesian optimisation framework. These net-power savings are listed under S1 in table II and show that a 1.2% net-power saving is achieved for Case 1. This is perhaps unexpected, illustrating that any change in input power to the blowing device requires the Bayesian optimisation framework search again through parameter space to find a new optimised control strategy. This result is encouraging in that a small improvement in efficiency of the blowing device could yield significant net-power savings. Moreover, a new Bayesian optimisation search may yield a larger net-power saving for a different control strategy. The second Bayesian optimisation study also predicts in Case 8, highlighted red in table II, that strong, short intermittent blowing (or wall-normal standing wave) yields a significant local maximum drag reduction of 51.5% with a potential net-power saving of up to 0.7%.

Figure 12 (a) shows the local skin-friction coefficient versus Re_θ for Case 6 (blue line), Case 8 (red line) and the canonical turbulent boundary-layer flow (Case 0). Similar to the first Bayesian optimisation study, substantial drag reduction is observed over the control region, which persists for several hundred boundary-layer thicknesses downstream of control with a slow spatial recover back to the canonical wall-turbulent flow. Figure 12 (b) shows the growth in boundary-layer thickness with wall-normal blowing control strategies for Case 6 and Case 8. Small undulations are observed over the control region due to the injection of wall-normal momentum by the blowing control followed by an almost constant increase in boundary-layer thickness far downstream, qualitative similar to the observations for the first Bayesian optimisation study, figure 5.

Note finally that the decomposition of the FIK identity for Cases 6 and 8, along with the near-wall behaviour of the streamwise velocity and evolution of the downstream vortical flow structures look qualitatively similar to the data obtained for Case 13 from the first Bayesian optimisation study, therefore are omitted for brevity for the second Bayesian optimisation

TABLE II: Wall-normal blowing parameters for BO2, with corresponding maximum local drag reduction (%), global drag reduction (%) and net-power saving (%). The blue row indicates one optimum control strategy with uniform blowing. The red row highlights one optimum control strategy with an intermittent blowing. The black row indicates the canonical turbulent boundary-layer flow. The net-power savings S and S1 use the mean and lower uncertainty limit values for power per unit area, respectively, as shown in figure 11 (b). For convenience, drag reduction and net-power saving are positive.

Case	C_B	N_B	α	Max Drag Red.	Global Drag Red.	Power Saving S	Power Saving S1
0	0	0	—	—	—	—	—
1	0.005	1	1	52	13.2	-0.9	1.2
2	0.005	1	0.47	48.5	6.0	-0.6	0.3
3	0.005	1	0.24	49	2.9	-0.3	0.0
4	0.0041	2	0.24	37	2.5	-0.2	0.2
5	0.0051	1	0.25	47.5	3.1	-0.3	0.2
6	0.007	1	0.23	60.5	4.1	0.1	0.5
7	0.0042	1	0.22	39	2.2	-0.3	0.0
8	0.0093	5	0.23	51.5	5.6	0.0	0.7
9	0.0088	4	0.24	55	5.3	0.0	0.7
10	0.0079	3	0.23	57	4.6	-0.1	0.5
11	0.0080	2	0.23	63.5	4.7	-0.1	0.6

study.

For the first Bayesian optimisation study, the algorithm converged to a control strategy corresponding to low-intensity uniform blowing over the full streamwise extent of the control region. By only considering the pressure difference across the blowing wall, a very optimistic idealised net-power saving of 5% was observed for the best set of parameters. Figure 13 shows the shape of the mean and covariance of the objective function (net-power saving) as a function of C_B , α and N_B for the first Bayesian optimisation study (top) and combined with the extra data from the second Bayesian optimisation study (bottom), in order to refine the shape of the mean and the covariance of the net-power saving. It should be noted that with more observations the shape of the mean net-power saving does not vary significantly,




FIG. 12: (a) Streamwise evolution of the friction coefficient as a function of Re_θ for all the cases. The thick solid lines corresponds to the three cases of interest in this study: Case 0 (black), Case 6 (blue) and Case 8 (red). (b) Streamwise evolution of Re_δ (based on the boundary-layer thickness) for the uncontrolled case, Case 6 and Case 8. The blowing section is highlighted in pink.

and the predicted location of the minimum does not change which suggests that the first Bayesian optimisation is converged. The only difference can be seen close to $C_B = 1$ where more data are available from the second Bayesian optimisation. The algorithm anticipates that more than 4% of net-power saving (see the dark red iso-contour) is possible with uniform blowing (close to $N_B = 1$) with relatively small blowing velocity ($v_w \approx 0.003$). It can be seen that the shape of the objective function is relatively simple, mainly 2D, which suggest that intermittent blowing (N_B parameter) does not have a significant role compared to the other parameters, at least in the present study. The 2D shape for the objective function is confirmed by repeating Case 12 from the first Bayesian optimisation study, but with a single blowing strip ($N_B = 1$, $\alpha = 0.77$, and $C_B = 0.6\%$, respectively), see the green dot in figure 13. The shape can be attributed to the small streamwise extent of the control area, to the small number of parameters for the optimisation and to the procedure to compute the net-power saving from the experimental data. Additionally, the simple shape of the objective function possibly explain why the algorithm is converging very quickly, with order 10 simulations.

Similarly, figure 14 shows the shape of the mean and covariance of objective function (net-power saving) as a function of C_B , α and N_B for the second Bayesian optimisation




FIG. 13: Shape of the mean (left) and covariance (right) of the objective function (net-power saving) as a function of the three parameters C_B , α and N_B for the first Bayesian optimisation (top) and combined with the extra data from the second Bayesian optimisation (bottom). net-power saving is shown as positive.

study (top) and is combined with the extra data from the first Bayesian optimisation study (bottom), in order to refine the shape of the mean and the covariance of the net-power saving. It can be seen that the shape of the objective function is very different between the first and second Bayesian optimisation studies, even if the parameters for the simulations are the same. The difference are related to the estimation of the power needed by the blowing control solution. Depending on the blowing device, it might be more efficient for

low intensity blowing over a long streamwise area or it might be more efficient for more intense blowing over a short streamwise area. For the second Bayesian optimisation study, the shape of the objective function is more complex, with the possibility of different optimal solutions based on different sets of parameters. This seems to suggest that for a given number of parameters, and depending on the input power required for the blowing, there might be different drag reduction strategies to achieve net-power saving. As such, the present results predict a net-power saving which is reported are encouraging and future studies will aim to the level of net-power saving by increasing the number of input parameters used within the Bayesian optimisation algorithm.

V. DISCUSSION

The two Bayesian optimisation studies have shown that a net-power saving on the order of a few percent could be possible using a low-amplitude wall-normal blowing control strategy at Reynolds numbers of $Re_\theta \sim 10^3$. Previous experimental studies have shown the potential of using a uniform low-amplitude wall-normal blowing approach, achieving turbulent skin-friction drag reductions of more than 50% in subsonic flows and 80% in supersonic flows when using blowing coefficients of $v/U_\infty \sim 0.001$ (Hwang [17, 18]). However, the focus on all previous investigations was on achieving a skin-friction drag reduction and the associated convective drag reduction mechanism, rather than optimising a subset of the blowing parameters to achieve a net-power saving as demonstrated here. Given that a 3% reduction in the turbulent skin-friction drag acting on a long-range commercial aircraft would save £1.2M in jet fuel per aircraft per year and prevent the annual release of 3,000 tonnes of carbon dioxide (Bushnell and Hefner [4]), the predicted net-power savings on the order of a few percent offer significant potential for future applications, especially given the potential of the control approach for higher Reynolds number flows (Hwang [17, 18]). The low-amplitude wall-normal blowing control strategies found within the Bayesian optimisation framework are effective in achieving a net-power saving as the skin-friction drag reduction persists for very long distances downstream of control as already observed in Stroh et al. [41]. These types of control strategies, which have long-lasting downstream effects, are attractive for many practical flows of interest including the flows over high-speed trains, along the hulls of ships or down the fuselage of an aircraft. However, accompanying the long-lasting skin-




FIG. 14: Shape of the mean (left) and covariance (right) of objective function (net-power saving) as a function of the three parameters C_B , α and N_B for the second Bayesian optimisation (top) and combined with the extra data from the first Bayesian optimisation (bottom). net-power saving is shown as positive.

friction drag reduction is a thickening of the turbulent boundary layer. Therefore, care would be needed if applying these control techniques in situations where the pressure drag plays an important role in overall performance, for example, in the flow over an aircraft's wings (Atzori et al. [2]). Noting the growing interest in battery technology and the production of green energy, alternatively powering an active low-powered control solution, rather than relying on the usual fossil fuelled systems, could yield significant savings in transport emissions ahead of any full-scale alternative-energy revolution. More broadly, energy re-

duction in any system can only be seen as positive. It is also worth noting, as pointed out by Kornilov and Boiko [26], that by re-using exhausted air feeds on flight, although the same approach is applicable across the transportation sector, could provide the blowing power required for a low-amplitude wall-normal blowing control strategy with little energy penalty, mimicking the use of compressed air in a laboratory setting. This would open the possibility of achieving up to a 5% net-power saving with simplistic actuator technology, as found experimentally (Kornilov and Boiko [26]), and here numerically.

With the primary objective of the Bayesian optimisation framework being to achieve a net-power saving, establishing the potential power consumption of the two different blowing devices was an integral part of this investigation. The different power consumption needed for each blowing device created different search paths through parameter space within the Bayesian optimisation framework, which ultimately resulted in two different optimised wall-turbulence control strategies. This is an interesting result illustrating that idealised control power would likely lead to non-optimal control strategies, if achieving a net-power saving was the primary objective. In this investigation, each blowing device was simulated as having a uniform steady wall-normal velocity over some part of the spatial domain. In reality, it is likely that either blowing device will be spatially discrete in the streamwise and spanwise directions to some extent, and each miniature electromagnetic speaker will apply a small amount of suction at the wall to refill the speaker cavity during excitation. However, the aim of the present investigation was not to implement a high-fidelity model of any one particular blowing device. Instead, the uniform blowing boundary condition has provided a reasonable approximation of each blowing device, and has allowed this investigation to highlight the strengths of a Bayesian optimisation approach for wall-turbulence control.

The Bayesian optimisation framework optimised just three blowing parameters to locate significant skin-friction drag reduction with net-power saving. Importantly, the Bayesian optimisation technique was able to converge to an optimised solution with very few simulations, order 10 in this investigation, highlighting that a Bayesian optimisation approach may also provide a powerful tool for experimentalists, where a time-averaged skin-friction drag value for a given set of control parameters could be obtained within minutes. Having an ability to optimise parameters with a reduced number of observations, whether these be with experiments or simulations, is important. Other forms of optimisation, for example, machine learning control with evolutionary algorithms (Gautier et al. [14]), have shown to

take up to one week of wind tunnel testing, requiring on the order of 500 observations to locate an optimised control law for a single actuator system; however, impressively, this evolutionary approach did find a new way to control a backward facing step flow at modest Reynolds number via a new drag reduction mechanism, which also emphasises the potential of a statistically-driven approach to flow control. Whilst few simulations were needed in the present investigation to locate an optimised solution, a mere three control parameters were inputted into the Bayesian optimisation framework. It is likely that increasing the number of input parameters will increase the number of observations required to determine an optimised solution. However, optimising an increased number of input parameters may enable new innovative pathways to control wall-turbulence with the potential for higher net-power savings. From a numerical view point, instead of using costly DNS, it would be more feasible to expedite numerical observations by using Implicit Large-Eddy Simulations (ILES). Recently, a new method was implemented in **Incompact3d** to perform ILES, which targeted numerical dissipation introduced at the small scales through the discretisation of the second derivatives of the viscous terms (Dairay et al. [7], Lamballais et al. [30]). It was shown in these studies that it is possible to design a high-order finite-difference scheme in order to mimic a subgrid-scale model for LES based on the concept of Spectral Vanishing Viscosity (SVV, see for instance Karamanos and Karniadakis [24], Tadmor [43]), at no extra computational cost. ILES of turbulent channel flows have already been performed with **Incompact3d** (Lamballais et al. [30]), providing the framework for future simulations of turbulent boundary-layer flows.

VI. CONCLUSIONS

This paper presents DNS of low-amplitude wall-normal blowing control of zero-pressure gradient turbulent boundary-layer flows at Reynolds numbers of $Re_\theta \sim 1000$. The aim of this investigation was to optimise the blowing amplitude and blowing coverage to enable a skin-friction drag reduction with a net-power saving. optimisation of the control parameters was performed within a Bayesian optimisation framework which required as little as 6 observations to determine the optimum set of control parameters. Control performance was assessed by using the power consumption from two different sets of experimental data from two different types of blowing device. The first type of blowing device was used by

Kornilov and Boiko [26] and conveyed compressed air into a plenum chamber then through a perforated plate into the oncoming turbulent boundary-layer flow. The second type of blowing device was original and generated low-amplitude wall-normal blowing with a miniature electromagnetic speaker. The wall-normal blowing control was simulated by a steady and uniform velocity boundary condition within the high-order flow solver **Incompact3d**, with the effects of control monitored over $650\delta_0$, where δ_0 was the height of the laminar boundary-layer flow at the inlet of the simulation at $Re_\theta = 170$.

The simulations have demonstrated that the low-amplitude wall-normal blowing control can generate a skin-friction drag reduction which persists for up to $650\delta_0$ downstream of control, and that it is the slow spatial recovery of the skin-friction coefficient back to its canonical counterpart which generates the net-power savings in this study. In particular, by decomposing the skin-friction coefficient using the Fukagata-Iwamoto-Kasagi (FIK) identity, it has been shown that it is the changes in contribution to the convection and streamwise development terms of the turbulent boundary-layer flows which generates the net-power savings.

Using the power consumption for the compressed air type blowing device caused the Bayesian optimisation framework to identify a uniform blowing control strategy as optimum. Here, a blowing amplitude 0.289% of the free-stream velocity generated a local skin-friction drag reduction of 36.5% and a global skin-friction drag reduction of 8.3%, with a net-power saving of 5%. Although, it should be noted that the cost of compressing the air was not accounted for in this energy budget. Similarly, by using the lowest estimate of the power consumption for the miniature electromagnetic speaker obtained in the present experiments (S1 in table II), a uniform blowing strategy with a blowing amplitude of 0.5% of the free-stream velocity generated a local skin-friction drag reduction of 52% and a global skin-friction drag reduction of 13.2%, with a net-power saving of 1.2%. In addition, net-power savings of 0.7% were found using a wall-normal standing wave control strategy, which was unexpected, and highlights the potential of using a Bayesian optimisation approach to find new wall-turbulence control strategies. Interestingly, the search paths through parameter space differed depending on the power consumption used for each Bayesian optimisation study, even when starting from the same initial parameter conditions, illustrating that any change related to the objective function requires a new Bayesian optimisation study.

Local maximum skin-friction drag reductions of 75.5% and 63.5% were observed using

the power consumption from the compressed air type or miniature electromagnetic speaker blowing device, respectively. Although it should be noted that obtaining a local maximum skin-friction drag reduction was not the focus of either study and is purely a consequence of the search path through parameter space made by the Bayesian optimisation framework. Therefore it is possible that higher levels of local skin-friction drag reduction may be found if this was the objective of the Bayesian optimisation.

More broadly, the results presented illustrate that a Bayesian optimisation approach provides a powerful tool for optimising wall-turbulence control strategies, requiring few observations to converge on the optimum set of control parameters. It is likely that the number of required observations will increase as the number of control parameters increase.

ACKNOWLEDGMENTS

The authors would like to thank the EPSRC for the computational time made available on the UK supercomputing facility ARCHER via the UK Turbulence Consortium (EP/L000261/1). The authors also acknowledge PRACE for awarding them access to Marconi at CINECA, Italy (project 2016163847) and Hazel Hen at HLRS, Germany (project 2018184381). OM would like to thank Imperial College London for funding his PhD with an Imperial College President Scholarship. RDW would like to acknowledge the support from the Air Force Office of Scientific Research through Grant No. FA9550-17-1-0231.

-
- [1] P Henrik Alfredsson and Ramis Örlü. Large-eddy breakup devices—a 40 years perspective from a stockholm horizon. *Flow, Turbulence and Combustion*, 100(4):877–888, 2018.
 - [2] M Atzori, R Vinuesa, A Stroh, B Frohnappfel, and P Schlatter. Assessment of skin-friction-reduction techniques on a turbulent wing section. *arXiv preprint arXiv:1812.03762*, 2018.
 - [3] Eric Brochu, Vlad M Cora, and Nando De Freitas. A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599*, 2010.
 - [4] Dennis M Bushnell and Jerry N Hefner. *Viscous drag reduction in boundary layers*. American Institute of Aeronautics and Astronautics, 1990.

- [5] Cheng Chin, Ramis Örlü, Philipp Schlatter, Jason Monty, and Nicholas Hutchins. Influence of a large-eddy-breakup-device on the turbulent interface of boundary layers. *Flow, Turbulence and Combustion*, 99(3-4):823–835, 2017.
- [6] Kwing-So Choi, Timothy Jukes, and Richard Whalley. Turbulent boundary-layer control with plasma actuators. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 369(1940):1443–1458, 2011.
- [7] T. Dairay, E. Lamballais, S. Laizet, and J.C. Vassilicos. Numerical dissipation vs. subgrid-scale modelling for Large Eddy Simulation. *Journal of Computational Physics*, 337:252–274, 2017.
- [8] P Dawes and H Wheeler. Preliminary experiments on sweat cooling. *Jet Propulsion Lab./California Institute of Technology Progress Report*, pages 3–13, 1946.
- [9] Gaetano Maria Di Cicca and Gaetano Iuso. On the near field of an axisymmetric synthetic jet. *Fluid dynamics research*, 39(9-10):673, 2007.
- [10] Carlos Diaz-Daniel, Sylvain Laizet, and J Christos Vassilicos. Direct Numerical Simulations of a wall-attached cube immersed in laminar and turbulent boundary layers. *International Journal of Heat and Fluid Flow*, 68:269–280, 2017.
- [11] Carlos Diaz-Daniel, Sylvain Laizet, and J Christos Vassilicos. Wall shear stress fluctuations: Mixed scaling and their effects on velocity fluctuations in a turbulent boundary layer. *Physics of Fluids*, 29(5):055102, 2017.
- [12] Koji Fukagata, Kaoru Iwamoto, and Nobuhide Kasagi. Contribution of reynolds stress distribution to the skin friction in wall-bounded flows. *Physics of Fluids*, 14(11):L73–L76, 2002.
- [13] Ricardo García-Mayoral and Javier Jiménez. Drag reduction by riblets. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 369(1940):1412–1427, 2011.
- [14] Nicolas Gautier, J-L Aider, THOMAS Duriez, BR Noack, Marc Segond, and Markus Abel. Closed-loop separation control using machine learning. *Journal of Fluid Mechanics*, 770:442–457, 2015.
- [15] Michael A Gelbart, Jasper Snoek, and Ryan P Adams. Bayesian optimization with unknown constraints. *arXiv preprint arXiv:1403.5607*, 2014.
- [16] Ari Glezer and Michael Amitay. Synthetic jets. *Annual review of fluid mechanics*, 34(1):503–529, 2002.

- [17] Danny Hwang. An experimental study of turbulent skin friction reduction in supersonic flow using a microblowing technique. In *38th Aerospace Sciences Meeting and Exhibit*, page 545, 1999.
- [18] Danny Hwang. Review of research into the concept of the microblowing technique for turbulent skin friction reduction. *Progress in Aerospace Sciences*, 40(8):559–575, 2004.
- [19] Javier Jiménez, Sergio Hoyas, Mark P Simens, and Yoshinori Mizuno. Turbulent boundary layers and channels at moderate reynolds numbers. *Journal of Fluid Mechanics*, 657:335, 2010.
- [20] Yukinori Kametani and Koji Fukagata. Direct Numerical Simulation of spatially developing turbulent boundary layers with uniform blowing or suction. *Journal of Fluid Mechanics*, 681:154–172, 2011.
- [21] Yukinori Kametani, Koji Fukagata, Ramis Örlü, and Philipp Schlatter. Effect of uniform blowing/suction in a turbulent boundary layer at moderate reynolds number. *International Journal of Heat and Fluid Flow*, 55:132–142, 2015.
- [22] Yukinori Kametani, Koji Fukagata, Ramis Örlü, and Philipp Schlatter. Drag reduction in spatially developing turbulent boundary layers by spatially intermittent blowing at constant mass-flux. *Journal of Turbulence*, 17(10):913–929, 2016.
- [23] Yong-Duck Kang, Kwing-So Choi, and Ho Hwan Chun. Direct intervention of hairpin structures for turbulent boundary-layer control. *Physics of Fluids*, 20(10):101517, 2008.
- [24] G.S. Karamanos and G.E. Karniadakis. A spectral vanishing viscosity method for Large-Eddy Simulations. *Journal of Computational Physics*, 163(1):22–50, 2000.
- [25] Joongnyon Kim, Kyoungyoun Kim, and Hyung Jin Sung. Wall pressure fluctuations in a turbulent boundary layer after blowing or suction. *AIAA journal*, 41(9):1697–1704, 2003.
- [26] Vladimir I Kornilov and Andrey V Boiko. Efficiency of air microblowing through microperforated wall for flat plate drag reduction. *AIAA journal*, 50(3):724–732, 2012.
- [27] Vladimir I Kornilov and Andrey V Boiko. Flat-plate drag reduction with streamwise noncontinuous microblowing. *AIAA journal*, 52(1):93–103, 2013.
- [28] Sylvain Laizet and Eric Lamballais. High-order compact schemes for incompressible flows: A simple and efficient method with quasi-spectral accuracy. *Journal of Computational Physics*, 228(16):5989–6015, 2009.
- [29] Sylvain Laizet and Ning Li. Incompact3d: a powerful tool to tackle turbulence problems with up to $0(10^5)$ computational cores. *International Journal for Numerical Methods in Fluids*, 67

- (11):1735–1757, 2011.
- [30] Eric Lamballais, Véronique Fortuné, and Sylvain Laizet. Straightforward high-order numerical dissipation via the viscous term for Direct and Large Eddy Simulation. *Journal of Computational Physics*, 230(9):3270–3275, 2011.
 - [31] Sanjiva K Lele. Compact finite difference schemes with spectral-like resolution. *Journal of computational physics*, 103(1):16–42, 1992.
 - [32] HS Mickley, RoCo Ross, AoLo Squyers, and WE Stewart. Heat, mass, and momentum transfer for flow over a flat plate with blowing or suction. Technical report, Massachusetts Inst. of Tech., 1953.
 - [33] Maurizio Quadrio. Drag reduction in turbulent boundary layers by in-plane wall motion. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 369(1940):1428–1442, 2011.
 - [34] Maurizio Quadrio and Pierre Ricco. Critical assessment of turbulent drag reduction through spanwise wall oscillations. *Journal of Fluid Mechanics*, 521:251–271, 2004.
 - [35] Carl Edward Rasmussen. Gaussian processes in machine learning. In *Advanced lectures on machine learning*, pages 63–71. Springer, 2004.
 - [36] Morris W Rubesin. An analytical estimation of the effect of transpiration cooling on the heat-transfer and skin-friction characteristics of a compressible, turbulent boundary layer. *NACA TN 3341*, 1954.
 - [37] Philipp Schlatter and Ramis Örlü. Assessment of Direct Numerical Simulation data of turbulent boundary layers. *Journal of Fluid Mechanics*, 659:116–126, 2010.
 - [38] Roger L Simpson, RJ Moffat, and WM Kays. The turbulent boundary layer on a porous plate: experimental skin friction with variable injection and suction. *International Journal of Heat and Mass Transfer*, 12(7):771–789, 1969.
 - [39] Jasper Snoek, Hugo Larochelle, and Ryan P Adams. Practical bayesian optimization of machine learning algorithms. In *Advances in neural information processing systems*, pages 2951–2959, 2012.
 - [40] Philippe R Spalart and J Douglas McLean. Drag reduction: enticing turbulence, and then an industry. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 369(1940):1556–1569, 2011.
 - [41] A Stroh, Y Hasegawa, Philipp Schlatter, and B Frohnäpfel. Global effect of local skin friction

- drag reduction in spatially developing turbulent boundary layer. *Journal of Fluid Mechanics*, 805:303–321, 2016.
- [42] Yasushi Sumitani and Nobuhide Kasagi. Direct Numerical Simulation of turbulent transport with uniform wall injection and suction. *AIAA journal*, 33(7):1220–1228, 1995.
- [43] E. Tadmor. Convergence of spectral methods for nonlinear conservation laws. *SIAM Journal on Numerical Analysis*, 26(1):30–44, 1989.
- [44] Chaitanya Talnikar, Patrick Blonigan, Julien Bodart, and Qiqi Wang. Parallel optimization for Large Eddy Simulations. *arXiv preprint arXiv:1410.8859*, 2014.
- [45] Kaoru Torii, Niichi Nishiwaki, and Masaru Hirata. Heat transfer and skin friction in turbulent boundary layer with mass injection. In *Proceedings of the Third International Heat Transfer Conference*, volume 3, pages 34–48. Am. Inst. Chem. Engrs New York, 1966.
- [46] Richard D Whalley and Kwing-So Choi. Turbulent boundary-layer control with plasma span-wise travelling waves. *Experiments in Fluids*, 55(8):1796, 2014.