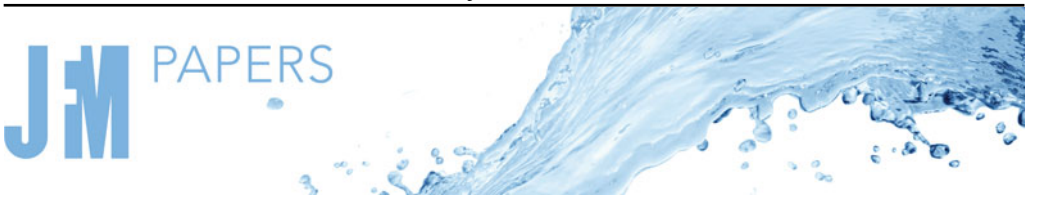
*J. Fluid Mech.* (2026), *vol.* 1036, A65, doi:10.1017/jfm.2026.11660




# Policy-based optimisation of spanwise wall oscillation waveforms for dissimilar-heat-transfer enhancement

**Lou Guérin**[1], **Cédric Flageul**[1], **Laurent Cordier**[1] **and Lionel Agostini**[1]

[1]Pprime Institute, Curiosity Team, CNRS, Université de Poitiers, Poitiers, France

**Corresponding author:** Lionel Agostini, lionel.agostini@univ-poitiers.fr

(Received 16 December 2025 UTC; revised 3 March 2026 UTC; accepted 4 May 2026 UTC)

The application of spanwise wall oscillations (SWO) to turbulent flow control has been largely confined to sinusoidal forcing and drag-reduction objectives; the optimisation of waveform topology across a continuous parameter space, and the application of SWO to dissimilar heat transfer (DHT), wherein the analogy factor $A_n$ quantifies the ratio of thermal to drag enhancement, remain unexplored. A policy-based optimisation framework coupled with implicit large-eddy simulations and Lagrangian-polynomial waveform parameterisation is employed to investigate four optimisation problems at $Re_\tau = 200$ and $Pr = 1$, where $Re$ is the friction Reynolds number and $Pr$ is the Prandtl number. Drag-reduction benchmarks validate the framework and establish that non-sinusoidal waveforms lead to no advantage for net energetic gain, whilst recovering a quasi-plateau topology that outperforms sinusoidal forcing when drag savings alone are maximised. Applied to DHT, waveform optimisation at 25 % reduced amplitude identifies a quasi-plateau configuration that achieves $\overline{A_n} = 1.087$, matching the sinusoidal optimum ($\overline{A_n} = 1.091$) whilst employing 7.4 % less control energy and producing 37.9 % greater absolute heat-transfer enhancement. These gains establish that sustained Stokes-layer strain during plateau phases governs preferential thermal transport through mechanisms fundamentally distinct from drag reduction. Phase-resolved analysis of the variance modulation reveals that the temperature variance regenerates more effectively and over a broader temporal interval than the velocity variance during the lingering phases, whilst the velocity field is suppressed earlier at the strain reversal. This asymmetric regeneration, hypothesised to originate in the pressure-strain redistribution term present exclusively in the momentum equations, constitutes the phase-resolved fingerprint of Reynolds analogy breaking even at $Pr = 1$. Waveform optimisation is thereby established as a superior, energy-efficient alternative to amplitude escalation for thermal management applications such as concentrated solar power receivers.

**Key words:** boundary layer control, mixing enhancement, turbulent mixing

© The Author(s), 2026. Published by Cambridge University Press.



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

## 1. Introduction

The mitigation of climate change represents one of the foremost challenges of the 21st century, necessitating the development and optimisation of renewable energy technologies. Concentrated Solar Power (CSP) plants form an essential component of this clean energy transition, wherein the thermal management of receivers demands enhanced convective heat transfer to maximise energy-conversion efficiency (David, Toutant & Bataille 2023). In solar pressurised-air receivers, air ($Pr \approx 0.71$) is utilised as the heat-transfer fluid, operating at friction Reynolds numbers estimated at around $Re_\tau \approx 1000$. At such moderate Reynolds numbers, near-wall turbulent structures remain the dominant contributors to both drag and heat transfer, as the large-scale outer-layer structures, whose strength increases as $Re_\tau$ increases, are not sufficiently energetic to substantially alter the wall-bounded flow dynamics. This regime is therefore particularly favourable for wall-based control strategies targeting the near-wall structures. The control of wall-bounded turbulent flows in such systems presents a fundamental fluid mechanics challenge: increasing heat transfer without incurring a prohibitive penalty in skin-friction drag.

Spanwise wall oscillations (SWO) represent a well-established open-loop control strategy for turbulent flow manipulation (Ricco, Skote & Leschziner 2021). In wall-bounded turbulent flows, the dominant contribution to skin-friction drag arises from near-wall coherent structures, specifically the low- and high-speed streaks that populate the buffer layer ($y^+ \approx 10$–$15$) , $y^+ = y \cdot u_\tau / \nu$ where $u_\tau = \sqrt{\tau_w}/\rho$. These streaks, which are sustained through a self-regenerating cycle involving streamwise vortices (Hamilton, Kim & Waleffe 1995; Waleffe 1997; Jiménez 2013), are responsible for the intense Reynolds shear stress that generates turbulent friction. Disrupting these structures constitutes the primary mechanism through which active flow-control strategies can achieve drag reduction. Unlike discrete actuation strategies that target specific coherent structures, SWO manipulate the flow over a large surface area, thereby interfering with the entire population of near-wall streaks simultaneously. The oscillating wall motion generates a Stokes layer, the penetration depth of which governs the effectiveness of the flow manipulation. Whilst several physical mechanisms have been proposed to explain the drag-reduction effect, including direct disruption of streak coherence, modification of the turbulence regeneration cycle and alteration of the near-wall vorticity dynamics, no universally accepted paradigm has emerged (Quadrio & Ricco 2004; Touber & Leschziner 2012; Agostini, Touber & Leschziner 2015). Nevertheless, the empirical effectiveness of SWO is well established, with control performance governed by the oscillation period through two competing effects. The Stokes-layer penetration depth, scaling as $\ell_s^+ \propto \sqrt{T^+}$, $T^+ = T/t_\tau$ where $t_\tau = \nu/u_\tau^2$, through viscous momentum diffusion, must be sufficient to reach the buffer region ($y^+ \approx 10$–$15$) where streaks are generated. However, excessively long periods enable complete regeneration of streaks between oscillation cycles, thereby negating the control effect. The optimal period thus represents a compromise: sufficiently long to ensure adequate penetration, yet sufficiently short to maintain disruption of the streak-generation cycle. Research on SWO has focused predominantly on drag reduction using sinusoidal forcing of the form $w_w(t) = W \sin(2\pi t/T)$, where $W$ is siamply the amplitude of the sinusoidal signal, with optimal performance achieved at relatively short periods ($T^+ \approx 100$–$125$) and maximum drag reductions approaching 45 % at low Reynolds numbers ($Re_\tau \approx 200$) (Quadrio & Ricco 2004; Quadrio, Ricco & Viotti 2009; Touber & Leschziner 2012; Gatti & Quadrio 2016).

Notwithstanding these advances, two significant research gaps remain unaddressed. The first pertains to signal waveform optimisation beyond the sinusoidal paradigm.



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

Whilst sinusoidal forcing is the standard, Cimarelli *et al.* (2013) demonstrated that non-sinusoidal waveforms, specifically those exhibiting plateau-like characteristics at maximum amplitude, can outperform sinusoidal forcing for pure drag reduction at fixed amplitudes. However, their study relied on a discrete set of nine pre-defined waveforms evaluated via direct numerical simulation (DNS). Consequently, the continuous parameter space of waveform topologies has never been systematically explored. It remains unknown whether an optimal waveform exists between these discrete points, or if specific non-sinusoidal shapes can lead to superior net energetic gains ($P_{net}$) in parameter regions where sinusoids are suboptimal.

The second, and more substantial, gap concerns the application of SWO for heat-transfer enhancement. In thermal control applications, the objective is to achieve dissimilar heat transfer (DHT), wherein heat-transfer enhancement exceeds the drag penalty. This condition is quantified by an analogy factor $\overline{A_n} = Nu_r/C_{f,r} > 1$, where $Nu_r$ and $C_{f,r}$ denote the ratios of Nusselt number and skin-friction coefficient relative to the uncontrolled baseline. The theoretical barrier to this objective is the Reynolds analogy, which predicts that modifications to turbulent-transport mechanisms affect momentum and thermal fields proportionally ($\overline{A_n} \approx 1$) for fluids with a Prandtl number near unity (Kader & Yaglom 1981). The present study adopts $Pr = 1$ to establish the most rigorous baseline in which the Reynolds analogy holds exactly; for air ($Pr \approx 0.71$), the thermal sublayer is marginally thicker, which would slightly reduce the absolute effectiveness of near-wall forcing on the thermal field, although the qualitative conclusions regarding waveform optimisation and DHT enhancement are expected to remain unchanged. Achieving DHT necessitates a fundamental breakdown of this analogy, selectively amplifying thermal transport whilst attenuating momentum transport – a challenging objective as the same turbulent mixing processes that generate skin friction also govern convective heat transfer. Whilst strategies such as surface roughness (Ligrani, Oliveira & Blaskovich 2003) typically augment drag disproportionately to heat transfer ($\overline{A_n} < 1$), the potential for SWO to achieve DHT was recently demonstrated by Guérin *et al.* (2024). At a specific, non-optimised operating point ($T^+ = 500$, $W^+ = 30$), they observed $\overline{A_n} = 1.064$. However, no systematic optimisation has been undertaken to maximise this effect, and the fundamental question of whether non-sinusoidal waveforms can facilitate this thermal-momentum decoupling remains entirely unexplored.

The present investigation addresses these gaps through a comprehensive optimisation study employing a policy-based optimisation (PBO) framework coupled with implicit large eddy simulations (LES) (Guérin *et al.* 2025), extended with a novel waveform parameterisation capability. A signal parameterisation based on Lagrangian polynomials is introduced to enable the exploration of signal topology as a continuous variable alongside period and amplitude. The investigation addresses four distinct optimisation problems.

(i) Waveform optimisation for drag reduction ($P_{sav}$): this case validates the framework by determining if continuous optimisation can independently recover or surpass the high-performance non-sinusoidal waveforms identified via discrete sampling by Cimarelli *et al.* (2013).
(ii) Waveform optimisation for net energetic gain ($P_{net}$): this case tests the hypothesis that non-sinusoidal waveforms might yield superior energetic efficiency compared with the sinusoidal optima established in previous parametric studies.
(iii) Sinusoidal optimisation for DHT: a rigorous baseline for DHT is established by identifying the optimal sinusoidal parameters ($T^+$, $W^+$) that maximise $\overline{A_n}$.

https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

(iv) Waveform optimisation for DHT: the primary investigation applies the waveform framework to DHT to determine if non-sinusoidal signals can outperform the optimal sinusoidal baseline. This addresses whether temporal waveform shaping constitutes an essential design parameter for energy-efficient thermal flow control.

This systematic approach establishes rigorous performance benchmarks for both sinusoidal and waveform-optimised configurations across multiple control objectives. The outcomes elucidate the role of waveform shape in SWO performance, provide the first optimised solutions for DHT and advance the understanding of thermal–momentum-transport coupling in oscillatory wall-bounded flows.

The remainder of this paper is organised as follows. Section 2 presents the methodology, encompassing the Lagrangian-polynomial waveform parameterisation, the PBO algorithm, performance metrics, numerical framework and the complete optimisation strategy. Section 3 details the results for each of the four optimisation cases: cases 1 and 2 address drag-reduction benchmarks that validate the framework and delineate the role of waveform topology across distinct control objectives; cases 3 and 4 constitute the primary investigation, establishing the sinusoidal and waveform-optimised baselines for DHT. The results section concludes with a comparative phase-resolved analysis that elucidates the temporal organisation of variance modulation under sinusoidal and plateau actuation, and advances a hypothesis for the mechanism responsible for the observed Reynolds analogy breakage. Section 4 summarises the principal contributions and outlines directions for future research. Appendix C investigates the energetic penalties associated with DHT enhancement under control-cost constraints, and Appendix A provides the mathematical derivations underpinning the waveform parameterisation.

## 2. Methodology

### 2.1. *Waveform parameterisation*

The systematic exploration of non-sinusoidal waveforms necessitates a parameterisation methodology capable of representing arbitrary temporal structures whilst maintaining mathematical tractability for gradient-based optimisation. To this end, a novel parameterisation based on Lagrangian-polynomial interpolation is introduced. The spanwise wall velocity signal $w_w(t)$ is constructed over one complete period $T$ subject to two fundamental constraints: periodicity and anti-symmetry. The anti-symmetry constraint, expressed as $w_w(t + T/2) = -w_w(t)$, guarantees symmetric positive and negative phases. This ensures zero net momentum injection over each complete cycle, a requisite for achieving equilibrium states, which closed-loop strategies often find challenging to attain (Cavallazzi *et al.* 2024). Furthermore, the restriction to time-continuous signals facilitates practical implementation compared with discontinuous waveforms.

This mathematical structure leads to the unique definition of the complete waveform via control points specified exclusively over the first half-period, $t \in [0, T/2]$. The five control points (the three free parameters plus the two boundaries) are positioned at equally spaced temporal intervals $t_i^1 = (i-1)T/8$ for $i \in \{1, \ldots, 5\}$, as illustrated in figure 1. To ensure continuity at period boundaries, the first and last control points are constrained to zero velocity

$$y_1^1 = 0 \quad (\text{at} t = 0) \quad \text{and} \quad y_5^1 = 0 \quad (\text{at } t = T/2). \tag{2.1}$$

Therefore, the waveform topology is determined entirely by the three interior control point amplitudes $(y_2^1, y_3^1, y_4^1)$. The selection of Lagrangian-polynomial interpolation, rather than



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

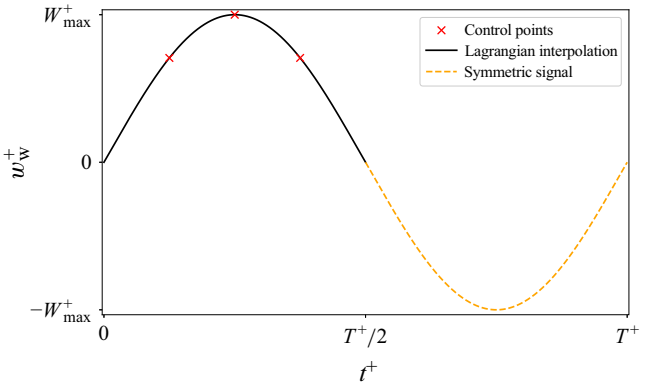



Figure 1. Illustration of the Lagrangian-polynomial parameterisation methodology. The boundary control points are fixed at zero amplitude to ensure periodicity, leaving three free interior control points ($y_2^1$, $y_3^1$, $y_4^1$, red crosses) that define the waveform shape over $[0, T/2]$. The second half-period is obtained through anti-symmetric reflection.

alternative techniques such as splines, is motivated by its computational efficiency for a small number of evenly spaced control points. The polynomial passing exactly through the design points is constructed directly, thereby circumventing the need to solve complex systems of equations. Furthermore, with only five control points, the primary limitations of high-degree interpolation (such as Runge's phenomenon) are effectively avoided.

Having established the control point structure and interpolation methodology, attention is now directed towards the parameter mapping strategy itself. To establish a robust parameter mapping whilst ensuring injectivity (whereby each parameter pair $(a_1, a_2)$ corresponds to a unique waveform topology, as detailed in Appendix A), the optimisation process avoids direct manipulation of the interior control point amplitudes. Instead, a reduced parameterisation is adopted, relying on only two parameters, $a_1 \in [0, 1]$ and $a_2 \in [0, 1]$. These parameters fully define the normalised waveform shape (prior to amplitude scaling) through determination of the three interior control point amplitudes according to the following transformation:

$$\begin{cases} y_1^1 = 0, \quad y_5^1 = 0, \\ y_2^1 = a_1, \\ y_3^1 = a_2(1 - a_1), \\ y_4^1 = (1 - a_1)(1 - a_2). \end{cases} \tag{2.2}$$

Worth highlighting is that this formulation explicitly enforces the normalisation constraint $\sum_{i=1}^{5} y_i^1 = 1$. As demonstrated rigorously in Appendix A, this constraint plays a fundamental role in guaranteeing the injectivity of the mapping from parameter pairs $(a_1, a_2)$ to waveform topologies. Injectivity, in turn, represents a cornerstone requirement for the effective deployment of gradient-based optimisation algorithms, fulfilling three objectives.

(i) Bijective mapping: it ensures that the parameter space defined by $T^+$, $W^+$ and $(a_1, a_2)$ is bijectively mapped to the space of admissible waveforms, establishing a well-defined correspondence between parameters and flow-control signals.
(ii) Elimination of redundancy: it prevents distinct parameter combinations from leading identical performance metrics. Such redundancy would introduce ambiguity in



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

gradient computation and potentially cause the algorithm to oscillate between equivalent parameter sets.

(iii) Convergence stability: it guarantees that the PBO algorithm (see § 2.2) can navigate the parameter space without encountering degenerate regions, thereby enabling systematic exploration and stable convergence towards optimal configurations.

The continuous waveform is constructed via fourth-order Lagrangian-polynomial interpolation. Specifically, for any time $t$ within the first half-period (denoted by superscript $k = 1$), the wall velocity is computed as

$$P_1(t) = \sum_{i=1}^{5} y_i^1 L_i^1(t), \quad \text{with} \quad L_i^1(t) = \prod_{\substack{j=1 \\ j\neq i}}^{5} \frac{t - t_j^1}{t_i^1 - t_j^1}. \tag{2.3}$$

For the second half-period (superscript $k = 2$), anti-symmetry is enforced by constructing control points with reflected temporal coordinates $t_i^2 = t_i^1 + T/2$ and amplitudes $y_i^2 = -y_{6-i}^1$, yielding the polynomial $P_2(t)$. The complete signal over one period is assembled as

$$w_w(t) = \mathbb{1}_{t\leqslant T/2}(t)\ P_1(t) + \mathbb{1}_{t>T/2}(t)\ P_2(t), \tag{2.4}$$

where $\mathbb{1}$ denotes the indicator function. Finally, this pre-scaled waveform is multiplied by a scaling factor $\mathrm{sc} = W_{max}^+ / \max_t(w_w(t))$ to ensure the maximum amplitude attains the target value $W_{max}^+$. The complete parameterisation mapping $\mathcal{L}$ thus incorporates $W_{max}^+$ as an additional degree of freedom controlling the absolute magnitudes of the interior control points $(y_2^1, y_3^1, y_4^1)$, whilst $(a_1, a_2)$ govern their relative proportions. This mapping is formally expressed as

$$\mathcal{L} : \left(a_1, a_2, T^+, W_{max}^+\right) \mapsto \mathrm{sc} \times P_1(t). \tag{2.5}$$

This methodology enables the representation of diverse waveform topologies, ranging from sinusoidal configurations to plateau-like structures with extended constant-velocity phases, through systematic variation of the parameters $(a_1, a_2)$ and $W_{max}^+$, providing sufficient expressiveness to discover non-trivial temporal structures whilst maintaining a compact parameter space. Before application to waveform optimisation, the Lagrangian parameterisation was validated against high-fidelity DNS benchmarks for sinusoidal oscillations at $T^+ = 125$ and $W^+ = 18$ (Quadrio & Ricco 2004), demonstrating that the method reproduces known drag-reduction performance ($P_{sav} = 39.27\,\%$ versus DNS reference of 39.3 %) within acceptable tolerances, as detailed in Appendix B.

### 2.2. *Optimisation algorithm: policy-Based Optimisation*

The optimisation is performed using the PBO algorithm (Viquerat *et al.* 2023), a gradient-based reinforcement learning method specifically designed for open-loop control in fluid dynamics applications. PBO is conceived as a hybrid framework that integrates the gradient-based learning efficiency of deep reinforcement learning with the population-based robustness of evolutionary strategies, drawing particular inspiration from the covariance matrix adaptation evolution strategy (CMA-ES) (Hansen 2023). This synthesis permits systematic exploration of high-dimensional parameter spaces whilst maintaining computational tractability through gradient-based updates. Benchmark comparisons have demonstrated that PBO performance is comparable to CMA-ES and substantially exceeds that of simplified single-step methods when minimising complex analytical functions (Viquerat *et al.* 2023). Furthermore, the algorithm has been successfully applied to



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

sinusoidal SWO parameter optimisation in prior work (Guérin *et al.* 2025), establishing its suitability for turbulent flow-control applications.

In the context of SWO, the control problem is formulated as a single-step decision process. In contrast to closed-loop control strategies, wherein an agent continuously adjusts actions based on time-varying flow states, the present open-loop scenario necessitates the determination of a single optimal parameter vector (comprising $a_1$, $a_2$, $T^+$, $W^+_{max}$) that is applied throughout the simulation. Consequently, instantaneous state observations are rendered unnecessary, and the optimisation relies exclusively on the terminal reward (e.g. drag reduction or analogy factor) obtained once the flow has attained a statistically steady state.

The PBO agent maintains a stochastic control policy $\pi_{\boldsymbol{\theta}}$, parameterised as a multivariate Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. This distribution is defined by three distinct neural network components that generate:

(i) the mean vector ($\boldsymbol{\mu}$), representing the current estimate of optimal parameters;
(ii) the standard deviation vector ($\boldsymbol{\sigma}$), governing the magnitude of exploratory sampling; and
(iii) rhe off-diagonal covariance coefficients ($\boldsymbol{\Sigma}$), enabling the algorithm to identify and exploit correlations between control parameters.

This explicit parameterisation of the covariance structure constitutes a principal advantage of PBO relative to simplified reinforcement learning algorithms, such as Proximal Policy Optimization-1 (Ghraieb *et al.* 2021), as it permits the sampling distribution to align with the principal axes of the objective function's topography. This alignment proves particularly beneficial in landscapes wherein parameters exhibit strong coupling, as is characteristic of turbulent flow-control problems where period and amplitude interactions govern performance.

At the algorithmic level, PBO proceeds as follows. At each generation, a population of $n_{ind}$ individuals is sampled from the current policy distribution $\pi_{\boldsymbol{\theta}}$ and evaluated through independent LES. The resulting rewards are utilised to update the policy via gradient ascent on the loss function

$$L(\theta) = \mathbb{E}_{a\sim\pi_\theta}[\log \pi_\theta(a)\ \hat{R}(a)], \quad \text{with} \quad \hat{R}(a) = \max\left(\frac{r(a)-\mu_r}{\sigma_r},\ 0\right), \tag{2.6}$$

where $r(a)$ denotes the reward (e.g. $P_{sav}$ or $\overline{A_n}$), and $\mu_r$ and $\sigma_r$ represent the reward mean and standard deviation over the current generation. This normalised, clipped formulation reduces variance and discards negative-advantage actions to stabilise the learning process. Three separate neural networks, each consisting of 3 hidden layers with 2 neurons, independently generate $\boldsymbol{\mu}$, $\boldsymbol{\sigma}$ and the off-diagonal covariance coefficients. The covariance matrix is constructed via the hypersphere decomposition method (Viquerat *et al.* 2023), which guarantees that valid positive semi-definite matrices are always produced from the network outputs. Each network is trained independently using the ADAM optimiser (Kingma & Ba 2014), with network-specific learning rates, epoch counts, batch sizes and generation history depths, as summarised in table 1. An exponential decay parameter $\alpha$ is employed to weight recent generations more heavily than earlier ones, thereby balancing off-policy data re-use with responsiveness to the current reward landscape. Complete algorithmic details, including the pseudo-code, sensitivity analyses with respect to population size and network architecture, and benchmark comparisons against CMA-ES, are available in Viquerat *et al.* (2023) and Guérin *et al.* (2025).



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

| $n_{ind}$ | $n_{para}$ | $n_{CPUs}$ | $\alpha$ | $\lambda_r(\mu, \sigma, \Sigma)$ | $n_e$ | $n_b$ | $n_g$ |
|---|---|---|---|---|---|---|---|
| 10 | 1 | 256 | 0.5 | $(5, 5, 1) \times 10^{-3}$ | (128,16,16) | (1,4,8) | (1,8,16) |

Table 1. The PBO training configuration for waveform optimisation. Here, $n_{ind}$: individuals per generation; $n_{para}$: parallel simulations; $n_{CPUs}$: CPU cores per simulation; $\alpha$: action sampling parameter; $\lambda_r$: learning rates for $(\mu, \sigma, \Sigma)$; $n_e$: epochs; $n_b$: mini-batch sizes; $n_g$: generation history depth.

The optimisation procedure operates in discrete generations. At each iteration, a population of parameter configurations (termed individuals) is sampled from the current policy $\pi_{\boldsymbol{\theta}}$ and evaluated through independent LES. The collected rewards are subsequently employed to update the policy parameters via gradient ascent, thereby concentrating probability mass within high-performing regions of the parameter space. This population-based evaluation strategy renders PBO particularly well suited to high-performance computing environments, as the computational cost is dominated by parallel fluid evaluations rather than sequential sampling (an inherent advantage relative to methods such as Bayesian optimisation). The complete PBO training configuration is summarised in table 1. Relative to the sinusoidal optimisation study presented in Guérin *et al.* (2025), wherein the parameter space comprised only two dimensions ($T^+$ and $W^+$) with 8 individuals per generation, the present waveform optimisation involves a four-dimensional parameter space. This increased action-space dimensionality necessitated elevating the population size to $n_{ind} = 10$ to ensure adequate sampling density, as convergence quality is substantially influenced by the ratio of population size to parameter space dimensionality (Viquerat *et al.* 2023). The computational configuration was adapted to accommodate the expanded parameter space and available computational resources. The strategy employs serial execution of simulations ($n_{para} = 1$), evaluating the 10 individuals sequentially rather than simultaneously, with each simulation allocated 256 CPU cores, i.e. representing a fourfold increase relative to the configuration considered in Guérin *et al.* (2025). These modifications produce contrasting effects on computational efficiency: wall time per generation is reduced by more than 50 % to approximately 2.8 h owing to the elevated per-simulation core allocation, whilst total CPU-time consumption per generation increases to 720 CPU hours owing to serial execution. This configuration reflects a deliberate trade-off, wherein optimisation throughput (reduced wall time) is prioritised over computational resource efficiency (increased CPU time). This strategy proved essential for exploring the higher-dimensional waveform parameter space within practical time scales whilst achieving convergence typically within 40–72 generations depending on problem complexity.

The hyperparameters listed in table 1, with the exception of $n_{ind}$, correspond to the default values provided by Viquerat *et al.* (2023), which were demonstrated to lead to robust convergence across benchmark problems of varying dimensionality (from two- to ten-dimensional Rosenbrock functions) without problem-specific tuning. The unique modification relative to the sinusoidal optimisation of Guérin *et al.* (2025) is the increase in population size from 8 to 10, motivated by the expanded four-dimensional parameter space. This limited tuning requirement is a practical advantage for the broader applicability of the PBO-LES framework to other flow configurations. Regarding initialisation, the PBO algorithm is designed to mitigate sensitivity to the first guess. The initial policy distribution is deliberately centred within the imposed action boundaries with a large variance and a diagonal covariance matrix, thereby ensuring broad exploration during the early generations. As the optimisation progresses, the covariance matrix adapts to the local



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

topology of the reward landscape through gradient-based updates of the correlation network. The population-based nature of PBO, i.e. wherein multiple individuals are evaluated per generation and the entire distribution is updated, provides an intrinsic mechanism for escaping suboptimal initial conditions, analogous to classical evolution strategies.

The choice of PBO over alternative optimisation techniques, such as gradient-based methods relying on finite-difference approximations, is motivated by several considerations. First, the inherent stochasticity of the reward signal (arising from finite-time averaging over limited oscillation cycles) renders finite-difference gradient estimates unreliable, as their accuracy depends on the difference between two noisy evaluations. The PBO, by contrast, extracts gradient information from the collective reward distribution across an entire population, effectively filtering statistical noise. Second, PBO naturally exploits parallel evaluation of independent simulations, avoiding the sequential bottleneck of perturbation-based gradient estimation. Third, the full covariance matrix adaptation provides directional search information analogous to second-order methods, without requiring explicit Hessian computation. Finally, the population-based exploration reduces the risk of convergence to local optima, which is particularly relevant for the DHT optimisation cases wherein the reward landscape is not known *a priori*.

### 2.3. *Performance metrics*

Having established the waveform parameterisation methodology and the optimisation algorithm, it is pertinent to define the performance metrics through which control efficacy is quantified. The assessment of SWO necessitates metrics that capture both the beneficial outcomes and the energetic costs inherent to active flow-control strategies. Throughout this investigation, five distinct metrics are employed, partitioned into two fundamental categories: momentum-transport metrics ($P_{sav}$, $P_{req}$, $P_{net}$) characterising drag reduction and energetic efficiency, and thermal-transport metrics ($Nu_r$, $\overline{A_n}$) characterising heat-transfer enhancement and thermal-momentum correlation, respectively. Throughout this section, the overbar notation $\overline{(\cdot)}$ denotes averaging over time and the homogeneous spatial directions ($x$ and $z$), whilst $\langle\cdot\rangle_{x,z}$ represents spatial averaging over the streamwise and spanwise directions only.

The drag-reduction metric ($P_{sav}$) quantifies the relative reduction in skin-friction coefficient achieved through SWO, expressed as a percentage

$$P_{sav} = 1 - \frac{\overline{C_{f,ac}}}{\overline{C_{f,unac}}}. \tag{2.7}$$

The skin-friction coefficient is defined as $C_f = \tau_w/((1/2)\rho U_b^2) = 2\nu U_b^{-2}\,\partial\langle u\rangle_{x,z}/\partial y|_w$, where $\tau_w$ denotes the wall shear stress. Positive values of $P_{sav}$ indicate successful attenuation of wall shear stress, with the magnitude reflecting the extent to which turbulent momentum transport has been suppressed relative to the uncontrolled baseline. This metric serves as the primary objective for validation cases wherein drag reduction is maximised independently of energetic expenditure considerations.

The control-cost metric ($P_{req}$) represents the power required to actuate the oscillating walls, normalised by the frictional power dissipation in the unactuated channel (Quadrio & Ricco 2004)

$$P_{req} = \frac{\overline{\tau_{w,z} w_w}_{ac}}{U_b\,\overline{\tau_{w,x}}_{unac}}, \tag{2.8}$$

where $U_b$ denotes the bulk velocity (maintained constant through the constant flow rate (CFR) strategy) and $\tau_{w,x} = \nu(\partial\langle u\rangle_{x,z}/\partial y)$ and $\tau_{w,z} = \nu(\partial\langle w\rangle_{x,z}/\partial y)$ denote the



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

streamwise and spanwise wall shear stresses, respectively, with summation accounting for contributions from both channel walls. This quantity is invariably positive and increases monotonically with both oscillation amplitude and frequency, thereby constituting the energetic penalty inherent to active flow-control strategies. Whilst drag reduction quantifies the beneficial reduction in pumping power, this metric captures the required investment of energy to maintain the control action itself.

The net energetic gain metric ($P_{net}$) represents the comprehensive energetic balance achieved through flow control, reconciling the competing contributions of drag reduction and control cost

$$P_{net} = P_{sav} - P_{req}. \tag{2.9}$$

Positive values indicate that the reduction in pumping power required to maintain CFR exceeds the power expended for wall actuation, thereby rendering the control strategy energetically advantageous for practical applications. This metric proves particularly pertinent for engineering implementations wherein energy efficiency constitutes a primary design constraint, serving as the objective function for validation cases addressing energetically optimal control configurations. The maximisation of $P_{net}$ necessitates a delicate balance between achieving substantial drag reduction whilst simultaneously minimising the energetic investment required to produce the oscillatory wall motion.

The heat-transfer enhancement metric ($Nu_r$) quantifies the relative augmentation of convective heat transfer induced by SWO. The Nusselt number $Nu$ is defined as $Nu = 2Hh/\lambda$, where $h$ denotes the convective heat-transfer coefficient and $\lambda$ the thermal conductivity; in the present dimensionless formulation, this reduces to $Nu = 2\,\partial\langle\Theta\rangle_{x,z}/\partial y|_w$, evaluated at the wall. The relative enhancement is thus expressed as

$$Nu_r = \frac{\overline{Nu_{ac}}}{\overline{Nu_{unac}}}. \tag{2.10}$$

Values exceeding unity indicate enhanced thermal transport relative to the uncontrolled baseline; however, this metric provides no information regarding the associated momentum-transport penalties. It is pertinent to recognise that simultaneous increases in both heat transfer and drag would yield elevated $Nu_r$ values without constituting desirable performance for practical applications, as thermal enhancement achieved through proportional momentum-transport augmentation offers no engineering advantage. Consequently, the evaluation of DHT performance necessitates a metric that explicitly accounts for the relative magnitudes of thermal- and momentum-transport modifications.

The analogy factor ($\overline{A_n}$) represents the ratio of relative heat-transfer enhancement to relative drag increase, thereby quantifying the extent to which the Reynolds analogy is broken by the control action

$$\overline{A_n} = \frac{Nu_r}{C_{f,r}} = \frac{\overline{Nu_{ac}}/\overline{Nu_{unac}}}{\overline{C_{f,ac}}/\overline{C_{f,unac}}}. \tag{2.11}$$

This metric constitutes the primary objective for DHT optimisation, with $\overline{A_n} > 1$ indicating preferential enhancement of thermal transport relative to momentum transport. The achievement of $\overline{A_n} > 1$ demonstrates that thermal mixing has been augmented to a greater extent than momentum mixing, thereby breaking the classical Reynolds analogy which traditionally couples these transport processes under dynamically similar conditions. Higher values of $\overline{A_n}$ signify more effective decoupling of thermal and momentum transport, representing superior performance for applications wherein heat-transfer enhancement is desired without commensurate drag penalties. The



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

maximisation of this metric thus constitutes a fundamentally different optimisation objective compared with drag reduction or net energetic gain, as it prioritises the selective manipulation of thermal structures whilst minimising momentum-transport modifications.

Mixed boundary conditions (MBC) are employed for the thermal field (Kasagi, Tomita & Kuroda 1992), imposing zero temperature at both walls. The unsteady source term associated with MBC ensures that a constant heat flux enters the domain. This configuration provides a robust baseline Reynolds analogy under uncontrolled conditions, thereby rendering the achievement of DHT ($\overline{A_n} > 1$) non-trivial, as both momentum and thermal transport are strongly coupled through the turbulent flow structures. The challenge of breaking this analogy under MBC conditions provides a stringent test for the efficacy of SWO in selectively manipulating thermal-transport mechanisms independently of momentum-transport processes.

The thermal-energy efficiency metric quantifies the thermal enhancement per unit of control energy expended, providing a comprehensive measure of the energetic efficiency with which thermal transport is augmented through active forcing. This metric is expressed as the ratio of relative heat-transfer enhancement to the absolute magnitude of required control power

$$\text{Thermal-energy efficiency} = \frac{\Delta Nu}{|P_{req}|}, \tag{2.12}$$

where $\Delta Nu = Nu_r - 1$ represents the fractional heat-transfer enhancement relative to the unactuated baseline. This formulation enables direct comparison of thermal performance across configurations operating at different actuation amplitudes and periods, as it explicitly accounts for the energetic investment required to achieve the observed thermal-transport augmentation. Higher values indicate superior thermal benefit per unit of control energy, thereby identifying actuation strategies that maximise heat-transfer enhancement whilst minimising energetic expenditure. This metric proves particularly pertinent for engineering applications wherein both thermal performance and energy efficiency constitute paramount design constraints, such as CSP systems requiring enhanced heat transfer without excessive parasitic power consumption.

### 2.4. *Numerical simulation framework and grid configurations*

The optimisation framework described in the preceding sections necessitates the execution of hundreds of flow simulations to systematically explore the waveform parameter space. To render this approach computationally tractable whilst maintaining physical fidelity, a multi-fidelity numerical strategy is employed, combining computationally efficient implicit LES for optimisation with high-fidelity DNSs for validation.

All simulations are performed using the `Xcompact3d` framework (Bartholomew *et al.* 2020), wherein the incompressible Navier–Stokes equations are solved in a turbulent channel flow at friction Reynolds number $Re_\tau = 200$. For DHT investigations, the momentum equations are coupled with a passive scalar transport equation for temperature at unity Prandtl number ($Pr = 1$). The numerical methodology employs sixth-order compact finite-difference schemes for spatial derivatives combined with a third-order Runge–Kutta scheme for temporal advancement, ensuring high accuracy whilst maintaining computational efficiency. For further details on the numerical configurations and governing continuous equations, refer to Guérin *et al.* (2024) for the DNS set-up and Guérin *et al.* (2025) for the LES configuration.



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

| Parameter | $S_1$ (LES) | $S_2$ (DNS) | Physical motivation |
|---|---|---|---|
| Usage | DR, $P_{net}$, $\overline{A_n}$ | Validation | — |
| Domain $(L_x, L_y, L_z)$ | (8, 2, 6) | (24, 2, 6) | Extended $L_x$ for largest structures |
| Grid points $(N_x, N_y, N_z)$ | $128 \times 65 \times 96$ | $400 \times 221 \times 200$ | High-fidelity resolution |
| $\Delta x^+$ | 11.0 | 10.7 | Resolved streamwise scales |
| $\Delta z^+$ | 11.0 | 5.3 | Resolved spanwise thermal gradients |
| $\Delta y^+_{min}$ (wall) | 2.0 | 0.43 | Viscous-sublayer resolution |
| $\Delta y^+_{max}$ (centreline) | 15.7 | 6.2 | — |
| Scalar transport | Active ($Pr = 1$) for DHT | Active ($Pr = 1$) | Temperature field for DHT |

Table 2. Grid specifications for the numerical configurations employed in the present investigation, at $Re_\tau \approx 200$ and $Pr = 1$. Here, $S_1$ is an implicit LES configuration used for all optimisation workflows (drag reduction and DHT), whilst $S_2$ provides DNS-level resolution for final validation.

The dimensionless governing equations, using the channel half-height $H$ and the bulk velocity $U_b$ as reference scales, are expressed as

$$\begin{cases} \boldsymbol{\nabla}^\star \cdot \boldsymbol{u}^\star = 0, \\ \dfrac{\partial \boldsymbol{u}^\star}{\partial t^\star} + \left(\boldsymbol{u}^\star \cdot \boldsymbol{\nabla}^\star\right) \boldsymbol{u}^\star = -\boldsymbol{\nabla}^\star p^\star + \dfrac{1}{Re}\, \Delta^\star \boldsymbol{u}^\star + \boldsymbol{f}^\star, \\ \dfrac{\partial \Theta^\star}{\partial t^\star} + \boldsymbol{u}^\star \cdot \boldsymbol{\nabla}^\star \Theta^\star + A^\star u^\star = \dfrac{1}{Pe}\, \Delta^\star \Theta^\star, \end{cases} \tag{2.13}$$

where $\boldsymbol{u}^\star = \boldsymbol{u}/U_b$, $p^\star = p/U_b^2$, $t^\star = tU_b/H$, $Re = U_b H/\nu$ is the bulk Reynolds number, $Pe = Re\, Pr$ the Péclet number and $\boldsymbol{f}^\star$ the body force maintaining CFR. The temperature $\Theta$ is non-dimensionalised following the MBC formulation of Kasagi *et al.* (1992), with $\Theta = (T - \langle T_w \rangle_t)/(T_b - \langle T_w \rangle_t)$, and $A^\star$ represents the streamwise temperature gradient arising from the constant wall-heat-flux condition. Periodic boundary conditions are imposed in the streamwise ($x$) and spanwise ($z$) directions, with no-slip conditions for velocity and $\Theta_w = 0$ at both channel walls. The flow is driven at CFR through an adaptive streamwise body force. The divergence-free condition is enforced through a fractional step method, wherein a Poisson equation for the pressure is solved using three-dimensional fast Fourier transforms (Laizet & Lamballais 2009). Hereafter, the star superscripts are dropped for brevity.

The investigation employs two distinct grid configurations, each designed to fulfil specific computational requirements within the optimisation-validation workflow. Table 2 summarises the complete specifications for all configurations, with resolution requirements motivated by the physical scales that must be adequately captured for each problem category.

The $S_1$ configuration is employed for all optimisation campaigns, encompassing drag reduction, net energy gain and DHT investigations. This configuration comprises a computational domain of $(L_x, L_y, L_z) = (8, 2, 6)$ discretised with $128 \times 65 \times 96$ grid points, leading to streamwise and spanwise resolutions of $\Delta x^+ = 11.0$ and $\Delta z^+ = 11.0$, respectively. The wall-normal resolution varies from $\Delta y^+_{min} = 2.0$ at the wall to $\Delta y^+_{max} = 15.7$ at the channel centreline, ensuring adequate representation of near-wall turbulent structures. The relatively compact streamwise domain proves sufficient for parametric studies, as the SWO primarily influence the near-wall region wherein turbulent streaks and quasi-streamwise vortices develop over shorter streamwise distances. This configuration balances computational efficiency with adequate physical fidelity for both momentum- and thermal-transport phenomena, requiring approximately 96 CPU hours per simulation. This



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

set-up employs a fourth-order accurate spectral vanishing viscosity (Lamballais, Fortuné & Laizet 2011) to provide controlled numerical dissipation at the smallest resolved scales, thereby characterising these simulations as implicit LES. This approach avoids explicit subgrid-scale modelling whilst maintaining numerical stability, a strategy that has been rigorously validated against DNS benchmarks in prior work (Guérin *et al.* 2025). For DHT investigations, an additional passive scalar temperature equation is resolved at $Pr = 1$, with numerical dissipation increased by a factor of 2.25 through the spectral vanishing viscosity scheme to address small instabilities observed in some drag/heat-transfer increase cases. The LES methodology demonstrates accurate prediction of primary performance metrics, with typical discrepancies in friction coefficient and Nusselt number remaining below 3 % and 5 %, respectively, relative to fully resolved DNS. It is noted that the $S_1$ resolution is already close to standard DNS resolution for this Reynolds number; the implicit LES classification arises from the use of spectral vanishing viscosity rather than from substantially coarsened grid spacing. Preliminary tests with coarser grids revealed excessive numerical dissipation during the oscillation-cycle phases in which near-wall structures are fully tilted in the spanwise direction by the Stokes layer, corrupting in particular the thermal-transport metrics. The present resolution thus represents a practical lower bound for reliable performance evaluation within the optimisation framework, as it ensures adequate capture of the spanwise-distorted streaks during maximum-tilt phases without excessive artificial damping.

The $S_2$ configuration provides DNS-level resolution for final validation of optimal configurations identified through the optimisation workflow. This high-fidelity grid comprises $400 \times 221 \times 200$ points distributed over the extended domain $(L_x, L_y, L_z) = (24, 2, 6)$, leading to substantially refined resolution ($\Delta x^+ = 10.7$, $\Delta z^+ = 5.3$, $\Delta y^+_{min} = 0.43$) capable of fully resolving all dynamically significant scales. The $S_2$ configuration serves as the authoritative standard against which final performance metrics are established, ensuring that the conclusions drawn from the optimisation-based exploration are substantiated through rigorous high-fidelity validation. Each DNS simulation requires approximately 18 000 CPU hours, rendering systematic parameter space exploration infeasible without the multi-fidelity strategy.

This hierarchical approach enables a computationally pragmatic workflow: extensive parameter exploration is performed through hundreds of LES at reduced computational cost, whilst final reported performance metrics are verified against DNS standards. The computational advantage is substantial, with each LES evaluation requiring approximately 0.5 % of the resources necessary for an equivalent DNS run, thereby permitting comprehensive exploration of the continuous waveform parameter space that would otherwise be prohibitively expensive. The subsequent validation of optimal configurations through DNS ensures scientific rigour whilst maintaining practical feasibility for discovering novel control strategies.

### 2.5. *Computational strategy and statistical convergence*

The PBO-LES framework employs a fundamentally different computational strategy compared with traditional parametric studies. Whilst Cimarelli *et al.* (2013) performed 90 high-fidelity DNS runs, each requiring integration times of approximately 10 000 viscous time units with 95 wash-out times to ensure statistical convergence, the present investigation executes hundreds of shorter LES per optimisation case (typically 400–720 simulations depending on problem complexity and convergence behaviour).

Each PBO-LES comprises two distinct phases. First, a warm-up period of 3200 wall units is enforced to ensure the flow adapts to the imposed actuation, thereby



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

reaching a statistically steady state prior to data collection. This extended warm-up duration, substantially exceeding typical transient time scales, guarantees that the modified flow regime is fully established before statistical sampling commences. Second, statistical collection proceeds over multiple complete oscillation cycles, with the duration determined adaptively according to the period through the relation $t^+_{collect} = T^+\lfloor 1630/T^+ \rfloor$, where $\lfloor\cdot\rfloor$ denotes the floor function. This formulation ensures that data are collected over an integer number of complete cycles, typically spanning 4–15 cycles depending on the period (corresponding to approximately 1630 wall units of statistical data per simulation). For the drag-reduction case at $T^+ = 111$, this leads to approximately 14–15 complete cycles; for the DHT waveform case at $T^+ = 350$, this corresponds to 4–5 complete cycles. This duration, whilst substantially shorter than traditional DNS validation runs, proves sufficient for the optimisation framework owing to a fundamental distinction between single-simulation convergence and population-based convergence. Individual simulations of limited duration inevitably retain residual statistical noise in their performance metrics ($P_{sav}$, $\overline{A_n}$), and 4–5 cycles would be insufficient for fully converged second-order statistics in a traditional validation study. However, the PBO algorithm does not rely on the absolute accuracy of individual evaluations; rather, it exploits the aggregate signal emerging from simultaneous evaluation of multiple individuals per generation ($n_{ind} = 10$) within the parameter neighbourhood. The algorithm updates its policy based on the collective reward distribution across these individuals, such that statistical noise from any single short-duration simulation is effectively filtered through population averaging. Convergence is evidenced by the consistency of rewards across final generations, wherein both the best individual and the generation average plateau at comparable performance levels. This population-based approach contrasts fundamentally with single-point evaluation strategies, wherein each simulation must achieve high statistical fidelity independently. In the PBO framework, individual simulations serve primarily for relative comparison within and across generations, with absolute performance subsequently validated through high-fidelity DNS. The optimal configurations identified through this efficient exploration methodology undergo rigorous validation using high-fidelity DNS. By integrating computationally efficient LES for exploration with DNS for final validation, this two-stage approach ensures both scientific rigour and practical feasibility. The computational advantage is significant: each LES requires approximately 96 CPU hours, compared with 18 000 CPU hours for an equivalent DNS run. This efficiency enables the execution of hundreds of simulations, facilitating a comprehensive exploration of the continuous parameter space that would otherwise be infeasible with DNS alone.

### 2.6. *Optimisation strategy and problem formulation*

Having established the computational infrastructure, i.e. waveform parameterisation, optimisation algorithm, numerical framework, performance metrics and statistical convergence strategy, it is pertinent to delineate the structured experimental programme through which these methodological components are deployed. The investigation is organised around four distinct optimisation problems, which must be addressed in a specific sequential order dictated by methodological necessity. The experimental design follows a systematic progression from validation on partially established benchmarks (wherein optimal-waveform topologies are known from discrete sampling but continuous parameter space exploration remains incomplete) to discovery on unexplored problems, with each case building upon the insights and confidence established by its predecessors. This hierarchical approach ensures that the framework's capabilities are rigorously



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

| Case | Objective | Signal Type | Parameters |
|---|---|---|---|
| 1 (validation) | Maximise $P_{sav}$ | Waveform | $T^+, a_1, a_2$ (fixed $W^+ = 15$) |
| 2 (validation) | Maximise $P_{net}$ | Waveform | $T^+, W^+, a_1, a_2$ |
| 3 (DHT) | Maximise $\overline{A_n}$ | Sinusoidal | $T^+, W^+$ |
| 4 (DHT) | Maximise $\overline{A_n}$ | Waveform | $T^+, a_1, a_2$ (fixed $W^+ = 30$) |

Table 3. Summary of optimisation problems. Cases 1–2 validate the waveform framework against established results from drag-reduction studies. Cases 3–4 address the primary research objective: discovering optimal actuation signals for DHT. An auxiliary sinusoidal optimisation at $W^+ = 30$ was additionally performed to provide an iso-amplitude comparison baseline for case 4 analysis (see § 3.4).

established before application to novel research questions wherein ground-truth solutions remain unknown. Table 3 summarises the four cases, differentiated by objective (or reward) function, signal parameterisation and control parameter space.

2.6.1. *Validation strategy: cases 1–2*

Prior to application to unexplored problems, the waveform optimisation framework is subjected to rigorous validation on drag-reduction benchmarks wherein optimal solutions have been identified from prior literature through discrete sampling in restricted parameter spaces. This validation proves essential, as it demonstrates both the framework's capacity to recover complex non-trivial waveforms and its resistance to 'over-parametrisation' when simpler solutions are optimal. The complementary nature of cases 1 and 2, wherein complex waveforms and simple sinusoids are respectively optimal, ensures the framework is tested across diverse optimisation landscapes, providing comprehensive confidence in its reliability before application to dissimilar-heat-transfer problems.

Case 1 optimises the waveform at fixed amplitude ($W^+ = 15$) to maximise drag reduction, as quantified by the metric $P_{sav}$ (see (2.7)). The amplitude is deliberately selected to operate in the pre-relaminarisation regime, permitting waveform effects to be observed independently of regime transitions that occur at higher amplitudes. The investigation by Cimarelli *et al.* (2013), wherein nine predefined waveforms were evaluated through 90 discrete DNS simulations, demonstrated that square-wave (Heaviside-like) configurations with abrupt transitions achieve superior drag reduction compared with sinusoidal signals at this amplitude. Successful recovery of analogous waveform topologies characterised by extended constant-velocity phases through the present continuous optimisation would establish bidirectional validation: the framework possesses sufficient sensitivity to identify performance-enhancing waveform topologies, whilst simultaneously confirming through independent methodology the discrete sampling observations reported in prior work. This case thus serves to verify that the Lagrangian parameterisation can discover beneficial temporal structures when they exist.

Case 2 extends validation to the complete four-parameter space ($T^+$, $W^+$, $a_1$, $a_2$), maximising net energetic performance as quantified by $P_{net}$ (see (2.9)). For this objective, prior investigations suggest that sinusoidal waveforms achieve optimal or near-optimal performance at the global optimum (Quadrio & Ricco 2004; Cimarelli *et al.* 2013), as non-sinusoidal control signals typically incur elevated energy consumption without offering commensurate drag-reduction benefits when operating at the energetically optimal amplitude and period. Convergence to a sinusoidal or near-sinusoidal configuration would demonstrate that the algorithm navigates the expanded parameter space without spurious attraction to waveform complexity, thereby confirming it correctly identifies the true global optimum despite the availability of additional degrees of freedom. This



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

complementary test ensures that the framework does not artificially favour complex waveforms when simpler solutions are energetically superior, a crucial verification for establishing confidence in subsequent discovery applications. Although Cimarelli *et al.* (2013) suggested the potential existence of alternative non-sinusoidal waveforms that might achieve comparable $P_{net}$ performance at specific parameter combinations, systematic exploration of the continuous parameter space through optimisation provides the most rigorous test of this hypothesis.

#### 2.6.2. *Discovery strategy: cases 3–4*

Following successful validation on drag-reduction benchmarks, the framework is applied to the primary research objective: discovering optimal actuation signals for dissimilar-heat-transfer enhancement. Unlike drag reduction, wherein optimal parameters are well established through extensive prior investigations, the thermal flow-control problem lacks such foundational parametric knowledge. The discovery phase therefore requires a systematic two-stage approach to establish rigorous performance baselines before exploring waveform effects.

The investigation proceeds through two sequential cases. Case 3 first identifies the optimal sinusoidal parameters through comprehensive exploration of the $(T^+, W^+)$ space, establishing the best achievable performance using conventional sinusoidal forcing. This 'sinusoidal champion' then serves as the performance benchmark for case 4, wherein waveform shape is optimised whilst holding amplitude below the sinusoidal optimum. This sequential structure is essential for isolating waveform topology contributions: without first identifying optimal sinusoidal parameters, any performance improvements in case 4 could not be unambiguously attributed to beneficial temporal structures, as they might instead arise merely from improved parameter selection relative to an arbitrary baseline. The two-stage progression thus ensures that any superiority of non-sinusoidal waveforms is demonstrably attributable to shape effects rather than confounded by comparison against a suboptimal sinusoidal reference.

Case 3 optimises period and amplitude $(T^+, W^+)$ to maximise the analogy factor $\overline{A_n}$ over an extensive parameter space ($10 \leqslant T^+ \leqslant 750$, $10 \leqslant W^+ \leqslant 40$). This broad exploration ensures comprehensive coverage and reliable identification of the global sinusoidal optimum, avoiding premature convergence to local maxima that might arise from restricted search domains. The resulting configuration, designated the 'sinusoidal champion', provides substantial improvement over the non-optimised reference of $\overline{A_n} = 1.064$ reported by Guérin *et al.* (2024), thereby establishing a rigorous performance target for waveform-based strategies. Importantly, case 3 identifies the optimal amplitude $W^+_{opt}$, subsequently employed as a reference for parameter selection in case 4. The comprehensive nature of this exploration ensures that the sinusoidal baseline represents the true optimum rather than a convenient but potentially suboptimal comparison point.

Case 4 investigates the central research hypothesis: whether non-sinusoidal temporal structures can achieve superior DHT performance relative to the optimal sinusoidal baseline, ideally at reduced energetic cost. The amplitude is fixed at $W^+ = 30$, selected to operate within the moderate-amplitude regime whilst remaining below the upper bound explored in case 3. This strategic amplitude selection serves dual purposes: it enables assessment of whether waveform optimisation can compensate for reduced forcing intensity, potentially providing a more energy-efficient pathway to DHT enhancement, whilst simultaneously reducing optimisation dimensionality to concentrate computational effort on waveform parameter space exploration. Success is defined as identification of a waveform topology leading to $\overline{A_n}$ (see (2.11)) approaching or exceeding the sinusoidal



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

| Case | Objective | Strategy | Signal | $(a_1, a_2)$ | $T^+$ | $W^+_{max}$ | Reward |
|---|---|---|---|---|---|---|---|
| 1 | Drag reduction | Validation | Waveform | $\in [0, 1]^2$ | $\in [10, 300]$ | 15 | $P_{sav}$ |
| 2 | Net energetic gain | Validation | Waveform | $\in [0, 1]^2$ | $\in [10, 300]$ | $\in [0, 30]$ | $P_{net}$ |
| 3 | DHT baseline | Discovery | Sinusoidal | $\frac{1}{\sqrt{2}(1+\sqrt{2})}, \frac{2}{1+\sqrt{2}}$ | $\in [10, 750]$ | $\in [10, 40]$ | $\overline{A_n}$ |
| 4 | DHT optimisation | Discovery | Waveform | $\in [0, 1]^2$ | $\in [10, 750]$ | 30 | $\overline{A_n}$ |

Table 4. Systematic optimisation programme: parameter ranges, signal types and reward functions for the four sequential cases. Cases 1–2 (blue, drag reduction) validate the framework's capability to identify both complex waveforms (case 1) and simple solutions (case 2) without bias. Cases 3–4 (red, DHT) constitute the discovery phase, establishing optimal sinusoidal baseline (case 3) before exploring waveform effects (case 4). Parameter ranges $(a_1, a_2) \in [0, 1]^2$ apply to waveform cases; sinusoidal optimisation fixes $a_1 = 1/(\sqrt{2}(1+\sqrt{2}))$, $a_2 = 2/(1+\sqrt{2})$ to generate pure sinusoidal signals.

champion performance despite operating at moderate amplitude. Such an outcome would demonstrate that optimisation of the temporal waveform offers a viable alternative to simple amplitude escalation for enhancing DHT, with potential implications for energy-efficient thermal management applications.

The complete parameter ranges and reward functions employed for each optimisation case, with ranges selected to enable comprehensive exploration whilst maintaining physical realisability are conveyed by table 4. Waveform parameters are constrained to $(a_1, a_2) \in [0, 1]^2$ for all waveform cases, ensuring the generated signals remain physically realisable and bounded within the feasible domain defined by the Lagrangian parameterisation. Period and amplitude ranges are adjusted according to the specific physical regime under investigation and the particular objective function being optimised, with bounds informed by prior parametric studies and physical considerations regarding Stokes-layer penetration and control authority.

## 3. Results and discussion

The investigation comprises four sequential optimisation cases, structured according to the validation-then-discovery strategy detailed in § 2.6. Cases 1–2 validate the waveform framework against established drag-reduction benchmarks, demonstrating the framework's capacity to identify both complex waveforms (when beneficial) and simple solutions (when optimal). Cases 3–4 constitute the primary contributions, establishing optimal parameters for DHT through systematic optimisation of both sinusoidal and waveform configurations. This structured progression ensures that the framework is thoroughly validated before application to unexplored problems.

### 3.1. *Waveform optimisation for drag reduction (case 1)*

This investigation addresses case 1 of the validation strategy: waveform optimisation at fixed amplitude to maximise drag reduction ($P_{sav}$), serving to verify the framework's capacity to identify complex waveforms when they provide distinct performance advantages. This case extends the well-established sinusoidal drag-reduction problem into the continuous waveform parameter space defined by $(T^+, a_1, a_2)$. Whilst Cimarelli *et al.* (2013) demonstrated that certain non-sinusoidal waveforms from their discrete set achieved superior drag reduction compared with sinusoidal forcing, their investigation employed a fundamentally different methodology: ten predefined waveform topologies (including sinusoidal and nine non-sinusoidal configurations) were each evaluated at



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

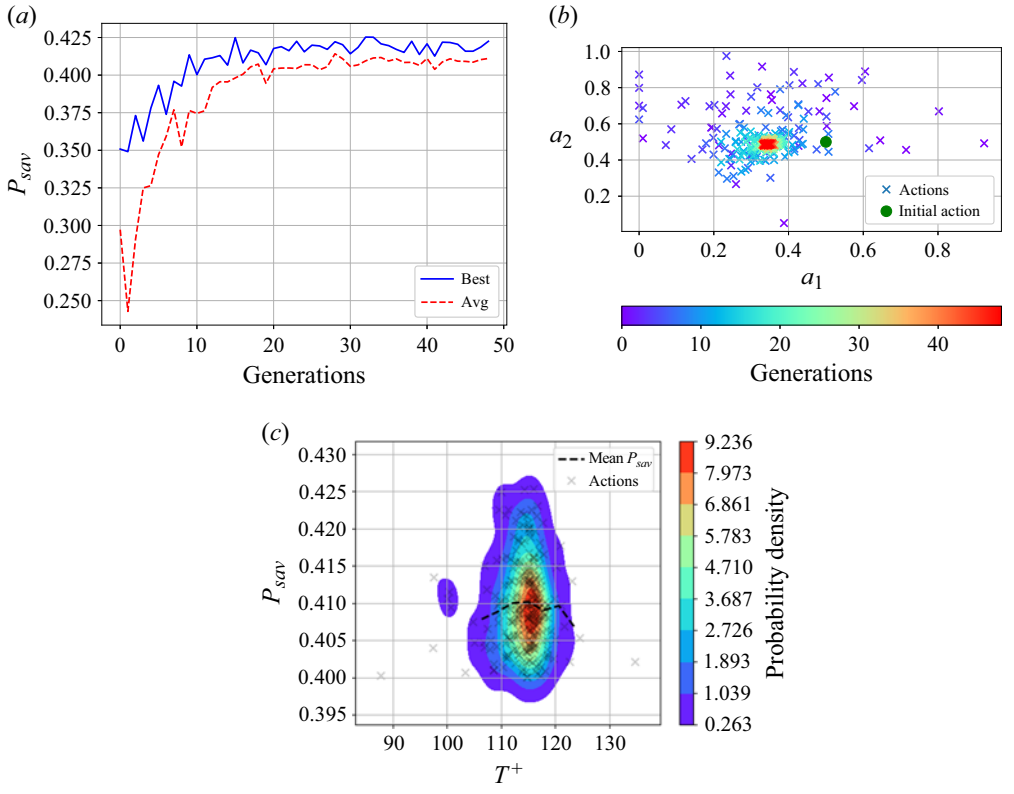



Figure 2. Convergence characteristics for drag-reduction optimisation (case 1). (*a*) Evolution of $P_{sav}$ through generations. (*b*) Parameter space exploration in the $a_1$–$a_2$ space. (*c*) Kernel density estimation in the $T^+$–$P_{sav}$ space, revealing optimal period $T^+ \in [110, 120]$.

nine discrete parameter combinations of period and amplitude ($T^+ \in \{62.5, 125, 250\}$ and $W^+ \in \{2.25, 4.5, 9\}$, selected as factors of $0.5\times$, $1\times$ and $2\times$ relative to the known sinusoidal optimal period), totalling 90 high-fidelity DNS simulations. In contrast, the optimisation-based approach employed here systematically explores the continuous waveform parameter space through hundreds of computationally efficient LES guided by the PBO algorithm, enabling discovery of optimal configurations without prior knowledge of promising regions and without restriction to predefined waveform topologies.

The amplitude was fixed at $W^+_{max} = 15$, a value slightly above the optimal amplitude of $W^+ \approx 12$ identified in prior sinusoidal studies (Quadrio & Ricco 2004), thereby providing the optimisation algorithm sufficient margin to explore waveform shapes without amplitude constraints unduly restricting the accessible performance envelope. The PBO agent was applied to optimise ($T^+$, $a_1$, $a_2$) for maximising $P_{sav}$. The framework converged within 48 generations (480 individual simulations), requiring approximately 140 wall-time hours and 35 000 CPU hours of computational resources. The convergence characteristics are shown in figure 2, demonstrating systematic exploration and convergence in the waveform parameter space. The reward evolution (figure 2*a*) exhibits steady improvement over the first 30 generations before stabilising, whilst the parameter space exploration (figure 2*b*) shows progressive clustering towards the optimal region at $a_1 \approx 0.34$, $a_2 \approx 0.49$. The best individual achieved $P_{sav} = 42.5\,\%$, with an average reward of $P_{sav} \approx 41\,\%$ across the final generation. Beyond the waveform parameters ($a_1$, $a_2$), the optimisation



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

| Configuration | $\%P_{sav}$ | $\%P_{req}$ | $\%P_{net}$ |
|---|---|---|---|
| Optimal sinusoidal | 38.12 | −118.39 | −80.27 |
| Optimal waveform | 40.64 | −184.92 | −144.28 |

Table 5. The DNS validation of optimal waveform vs. sinusoidal actuation for drag reduction (case 1). The waveform achieves superior drag reduction (highlighted), albeit at increased control cost, resulting in negative $P_{net}$ values wherein control energy expenditure exceeds drag-reduction benefits for both configurations.

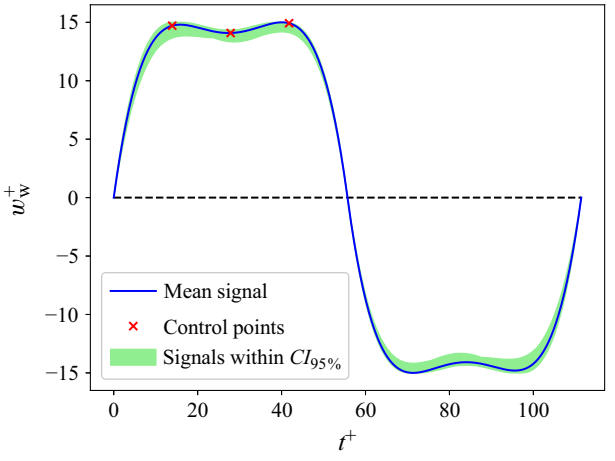



Figure 3. Optimal waveform for drag reduction (case 1) at $W^+ = 15$ with parameters $a_1 \approx 0.34$, $a_2 \approx 0.49$, $T^+ \approx 111$. The green shaded area highlights the 95 % confidence interval of the waveform parameters: $CI_{95\,\%} = \{a_1, a_2 \in [0, 1]^2 \mid \forall i \in \{1, 2\},\ a_i \in [\mu_{a_i} - 1.96\sigma_{a_i}, \mu_{a_i} + 1.96\sigma_{a_i}]\}$.

process simultaneously identified the optimal actuation period $T^+$. The kernel density estimation (figure 2*c*) reveals a well-defined optimal period range $T^+ \in [110, 120]$, with a converged period of $T^+ \approx 111$. This optimal period coincides precisely with that identified for sinusoidal signals, indicating period invariance to waveform shape. This represents a fundamental characteristic of drag-reduction optimisation: whilst waveform topology significantly influences performance magnitude, the optimal actuation frequency remains governed by the characteristic time scale of near-wall turbulence regeneration (Blesbois *et al.* 2013; Agostini *et al.* 2015), independent of signal shape.

The converged parameter configuration ($a_1 \approx 0.34$, $a_2 \approx 0.49$, $T^+ \approx 111$) produces the quasi-plateau waveform illustrated in figure 3, characterised by extended constant-velocity phases at extremum amplitudes. This waveform topology is consistent with that reported by Cimarelli *et al.* (2013), thereby validating the continuous optimisation framework's capacity to recover known optimal structures identified through discrete sampling.

A DNS validation of this optimal configuration was performed using the high-fidelity $S_2$ configuration at $T^+ = 111$, $W^+ = 15$. As summarised in table 5, the optimised quasi-plateau waveform achieves $P_{sav} = 40.64\,\%$, representing a 2.5 % improvement over the optimal sinusoidal signal at identical parameters (38.12 %). The Stokes-layer protrusion exhibits a 4.1 % increase relative to the sinusoidal baseline, indicating enhanced penetration of spanwise momentum into the flow field. However, this performance gain is accompanied by a 56.2 % increase in required control power ($P_{req}$), confirming that the quasi-plateau configuration is suboptimal for net energetic considerations. The elevated



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

control cost is attributable to the fundamental physics of plateau-phase actuation. The required control power is defined through $\int \tau_{w,z} w_w \, \mathrm{d}t$, representing the time-integrated product of spanwise shear stress and wall velocity. For the quasi-plateau waveform, the wall maintains constant extremum velocity ($w_w^+ = \pm 15$) over extended durations, in contrast to the sinusoidal signal wherein velocity varies continuously and only attains maximum amplitude instantaneously. During plateau phases, sustained spanwise shear stress must be maintained to preserve constant wall velocity against the viscous resistance of the already-displaced fluid within the Stokes layer. This contrasts with sinusoidal oscillations, wherein smooth acceleration and deceleration phases permit fluid momentum to assist wall motion during velocity reversals, thereby reducing instantaneous power requirements. The quasi-plateau topology therefore requires continuous energy injection to overcome viscous friction whilst maintaining constant velocity, analogous to maintaining constant velocity when pushing against persistent resistance rather than allowing natural momentum to assist during smooth transitions. Therefore, the extended plateau duration at maximum amplitude, coupled with sustained elevated shear stress, produces substantially increased time-integrated power expenditure relative to the sinusoidal baseline.

This result establishes bidirectional validation of the framework's capabilities. The methodology possesses sufficient sensitivity to distinguish performance differences arising from waveform topology variations, as evidenced by the recovery of the quasi-plateau configuration through independent systematic exploration of the continuous parameter space, thereby confirming the discrete sampling observations of Cimarelli *et al.* (2013). Furthermore, the consistency of the optimal period ($T^+ \approx 111$) across distinct waveform topologies demonstrates that drag reduction is fundamentally period dominated, with waveform shape providing secondary performance modulation. Case 1 thus validates the framework's capacity to identify complex waveforms when they provide performance advantages, establishing confidence for application to the unexplored thermal flow-control problems addressed in cases 3–4.

### 3.2. *Waveform optimisation for net energetic gain (case 2)*

In this section, the investigation addresses case 2 from the validation strategy: complete four-parameter optimisation to maximise net energetic gain, serving to verify that the framework does not exhibit an attraction to parametric complexity when simpler solutions are optimal. Whilst case 1 established that non-sinusoidal plateau waveforms achieve superior drag reduction at fixed amplitude, Cimarelli *et al.* (2013) demonstrated that sinusoidal signals become globally optimal when control costs are incorporated in the net energetic balance. The present investigation evaluates whether continuous waveform optimisation recovers this known sinusoidal optimum when exploring the complete four-dimensional parameter space ($T^+$, $W^+$, $a_1$, $a_2$) with net energetic gain ($P_{net} = P_{sav} - P_{req}$) as the objective function.

The convergence characteristics are illustrated in figure 4, which demonstrates systematic exploration across both $T^+$–$W^+$ and $a_1$–$a_2$ parameter spaces. The optimisation process required 72 generations to achieve parameter convergence, corresponding to 720 simulations, approximately 200 wall-time hours and nearly 50 000 CPU hours. This substantially greater computational investment relative to the drag-reduction case is attributable to the additional parameter dimension. Figure 4(*a*) shows the reward evolution, which exhibits stagnation beyond generation 20, with the best and average individuals achieving $P_{net}$ values of 9.3 % and 8 %, respectively. This stagnation is indicative of an extended near-optimal region within the parameter space, wherein the objective function exhibits relatively flat topography over a substantial domain. The $T^+$–$W^+$ exploration



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

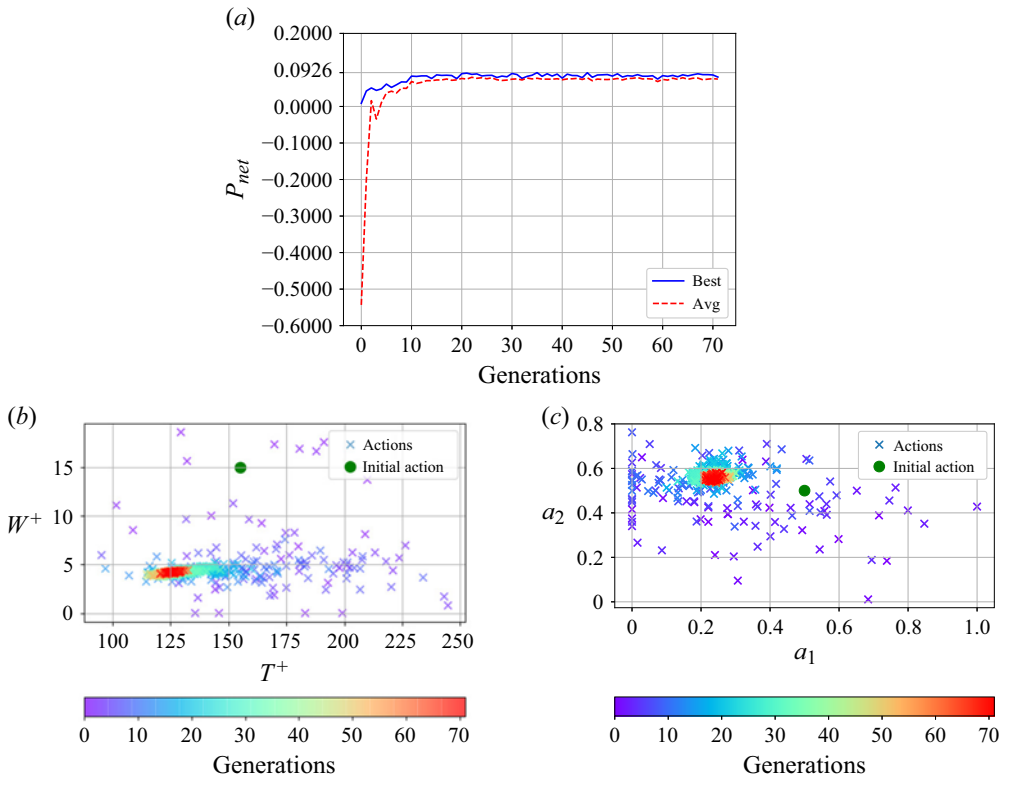



Figure 4. Convergence characteristics for net energetic gain optimisation (case 2). (*a*) Evolution of $P_{net}$ through generations. (*b*) Parameter space exploration in $T^+$–$W^+$ coordinates showing convergence at $T^+ \in [120, 140]$, $W^+ \in [3.5, 5]$. (*c*) Waveform parameter exploration in $a_1$–$a_2$ space exhibiting broader dispersion than case 1.

(figure 4*b*) shows convergence towards the parameter ranges $T^+ \in [120, 140]$ and $W^+ \in [3.5, 5]$, precisely coinciding with those identified for the optimal sinusoidal configuration (Guérin *et al.* 2025). The optimisation preferentially explored lower-amplitude regions, as elevated amplitudes substantially increase control costs. The $a_1$–$a_2$ exploration (figure 4*c*) demonstrates convergence to parameter values that produce near-sinusoidal waveforms (figure 5), with markedly broader dispersion relative to case 1 reflecting the extended near-optimal region wherein multiple waveform configurations lead to comparable $P_{net}$ performance. This wide parameter dispersion constitutes strong evidence of objective-function insensitivity to waveform topology: when control costs dominate net energetic balance, the precise waveform parameters become substantially less critical than period and amplitude, resulting in a diffuse convergence pattern wherein the algorithm identifies an extended optimal region rather than a singular optimal point. Crucially, the comparable $P_{net}$ values achieved across this dispersed parameter region (varying by less than 1.5 % points) confirm that the wide dispersion reflects genuine flatness of the objective-function landscape rather than incomplete convergence, thereby establishing that waveform exerts minimal influence on energetically constrained performance for drag-reduction objectives.

As illustrated in figure 5, the optimal waveform is indeed near sinusoidal. The 95 % confidence interval demonstrates greater waveform variability compared with the drag-reduction case, consistent with the broader spectrum of potentially optimal configurations within this extended parameter region. All waveforms within the confidence interval exhibit near-sinusoidal characteristics. The mean optimal signal displays slight



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

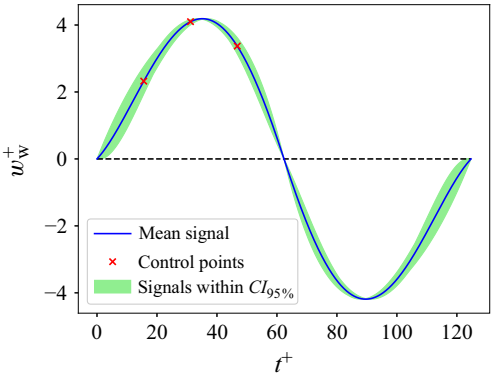



Figure 5. Optimal waveform for net energetic gain (case 2) converging to a near-sinusoidal configuration.

temporal asymmetry, featuring a marginally prolonged ascent phase relative to the descent phase. However, this asymmetry remains subtle, and the optimisation converged to a configuration nearly identical to the optimal sinusoidal signal without achieving measurable performance improvement.

Case 2 establishes that sinusoidal signals represent the global optimum for net energetic gain. The framework correctly identified this simple solution despite the freedom to explore complex waveform topologies across the four-dimensional parameter space, thereby validating that the methodology does not exhibit spurious attraction to parametric complexity when simpler solutions are optimal. Whilst Cimarelli *et al.* (2013) demonstrated that non-sinusoidal waveforms can locally outperform sinusoids at suboptimal parameter points, the present results confirm that no waveform exceeds sinusoidal performance even with varying $T^+$ and $W^+$. The wide dispersion observed in the $a_1$–$a_2$ parameter space (figure 4*c*) provides critical mechanistic insight: when control costs dominate the objective function, waveform topology exerts minimal influence on performance, rendering waveform parameters relatively inconsequential compared with period and amplitude. This insensitivity to shape establishes a fundamental characteristic of energetically constrained momentum-transport problems. The successful validation through cases 1–2 establishes confidence for application to unexplored problems where optimal solutions are unknown.

Having successfully validated the waveform optimisation framework through cases 1–2, wherein the framework correctly identified both complex optimal waveforms (drag reduction) and simple optimal solutions (net energetic gain) without systematic biases, the methodology is now applied with confidence to the primary research objective: identification of optimal actuation strategies for DHT enhancement. The discovery phase (cases 3–4) addresses a problem for which optimal solutions are not established in the literature.

### 3.3. *Sinusoidal optimisation for dissimilar heat transfer (case 3)*

This section focuses on case 3, which is part of the discovery strategy: establishing the optimal sinusoidal baseline that serves as the performance benchmark for subsequent waveform optimisation in case 4. Unlike the validation cases (cases 1–2), wherein optimal parameters were established through extensive prior investigations (Quadrio & Ricco 2004; Cimarelli *et al.* 2013; Guérin *et al.* 2025), the present thermal flow-control problem lacks such foundational parametric data. Although DHT has been achieved through various flow-control strategies (Yamamoto, Hasegawa & Kasagi 2013; Uchino, Mamori & Fukagata 2017), systematic investigation of sinusoidal SWO for this objective remains



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

largely unexplored. The feasibility of DHT via sinusoidal SWO was recently demonstrated by Guérin *et al.* (2024) at a single, non-optimised operating point ($T^+ = 500$, $W^+ = 30$), producing $\overline{A_n} = 1.064$. This preliminary result, however, provided no guidance regarding optimal parameters for maximising dissimilarity, necessitating systematic exploration of the parameter space to establish a rigorous performance baseline.

The investigation optimises sinusoidal parameters to maximise the time-averaged analogy factor $\overline{A_n}$, which quantifies the degree of preferential heat-transfer enhancement relative to the associated drag penalty, thereby serving as a direct measure of dissimilarity between thermal and momentum transport under actuation. The PBO-LES framework was applied to identify optimal parameters across the parameter space $T^+ \in [10, 750]$ and $W^+ \in [10, 40]$. The period range was selected to encompass both the drag-reduction optimal regime ($T^+ \approx 110$) and substantially longer periods potentially beneficial for thermal dissimilarity, whilst the upper bound was constrained by computational considerations to ensure adequate statistical sampling within each actuation cycle. The optimisation converged within 40 generations (400 simulations), requiring approximately 120 wall-time hours and 30 000 CPU hours of computational resources. This accelerated convergence, more than twice as rapid as the Nusselt number maximisation case reported in Guérin (2025), is attributed to improved statistical reliability at shorter optimal periods and the existence of a narrower optimal parameter region for DHT, compared with thermal enhancement objectives. The convergence characteristics are illustrated in figure 6, demonstrating systematic exploration and rapid convergence in the parameter space.

The reward evolution (figure 6*a*) exhibits rapid stabilisation following generation 15, with the best individual achieving $\overline{A_n} \approx 1.11$ and the generation average converging to comparable values. This close agreement between best and average performance indicates low statistical uncertainty and a well-defined optimal region, contrasting with the broader dispersion observed in net-energetic-gain optimisation (case 2). The reduced uncertainty is attributable to two factors: (i) the fourfold shorter optimal periods compared with thermal maximisation cases reported in Guérin (2025), permitting more complete actuation cycles within the statistical sampling window, and (ii) potential cancellation of statistical uncertainties between $Nu$ and $C_f$ fluctuations within the analogy factor formulation, as both quantities exhibit correlated variations that partially offset in the ratio.

The parameter space exploration (figure 6*b*) demonstrates that the framework initially explored the entire parameter domain before progressively concentrating evaluations within the near-optimal region. Convergence was observed at maximum amplitude ($W^+ = 40$) and moderate period ($T^+ \approx 347$), with systematic clustering towards elevated amplitudes. The convergence to maximum amplitude for sinusoidal signals indicates that, within the constraints of sinusoidal forcing, enhanced spanwise momentum penetration facilitates preferential thermal transport relative to momentum transport. However, as demonstrated subsequently in case 4, this amplitude–protrusion relationship observed for sinusoidal signals does not constitute a fundamental requirement for DHT enhancement when waveform topology is permitted to vary.

The optimal actuation period was identified through kernel density estimation in the $T^+$–$\overline{A_n}$ space, illustrated in figure 6(*c*). The analysis reveals a well-defined optimal period range $T^+ \in [320, 375]$, substantially narrower than the range identified for thermal maximisation ($T^+ \in [1300, 1600]$ as reported in Guérin 2025), whilst exhibiting comparable width to the drag-reduction range ($T^+ \in [110, 120]$). The observed uncertainty remains below 0.01 maximum difference in $\overline{A_n}$, markedly lower than in thermal enhancement cases. This narrow optimal range partially explains the accelerated convergence relative to other objectives and demonstrates that DHT exhibits sharply defined parameter sensitivity. The optimal period for sinusoidal DHT ($T^+ \in [320, 375]$)



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

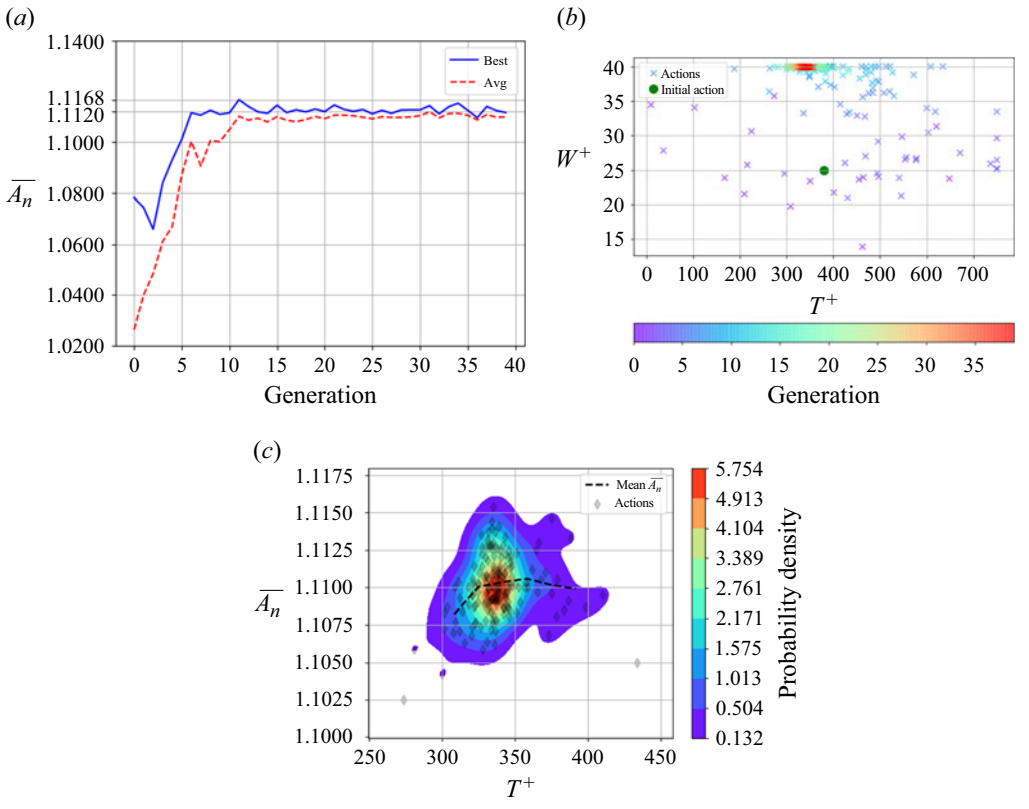



Figure 6. Convergence characteristics for sinusoidal DHT optimisation (case 3). (*a*) Evolution of $\overline{A_n}$ through generations. (*b*) Parameter space exploration in $T^+$–$W^+$ coordinates. (*c*) Kernel density estimation revealing optimal period $T^+ \in [320, 375]$.

occupies an intermediate position between drag reduction ($T^+ \in [110, 120]$) and thermal maximisation ($T^+ \in [1300, 1600]$), reflecting the distinct physical mechanism governing dissimilarity generation. Whereas drag reduction relies upon high-frequency disruption of near-wall turbulent structures through rapid Stokes-layer oscillations, and thermal maximisation benefits from extended periods permitting maximal Stokes-layer protrusion (Guérin 2025), DHT requires a balanced intermediate time scale that maintains sustained turbulence regeneration whilst achieving preferential thermal-transport enhancement relative to momentum transport.

A DNS validation of the optimal configuration ($T^+ = 325$, $W^+ = 40$) was performed using the high-fidelity $S_2$ mesh configuration at extended statistical collection time, with results presented in table 6. The optimised sinusoidal configuration achieves $\overline{A_n} = 1.091$, representing a 2.5 % improvement over the non-optimised baseline reported by Guérin *et al.* (2024) ($\overline{A_n} = 1.064$ at $T^+ = 500$, $W^+ = 30$). The heat-transfer enhancement reaches 16.40 %, whilst the drag penalty remains comparatively modest at 6.51 %, confirming that significant dissimilarity can be achieved through optimised sinusoidal actuation at moderate actuation amplitudes.

Case 3 establishes the first systematically optimised sinusoidal baseline for DHT through SWO. The converged parameters ($T^+ = 325$, $W^+ = 40$) and validated performance ($\overline{A_n} = 1.091$) provide a rigorous benchmark against which waveform optimisation can be evaluated in case 4. The substantially different optimal periods



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

| Configuration | $\%\Delta Nu$ | $\%\Delta C_f$ | $\overline{A_n}$ |
|---|---|---|---|
| Optimal sinusoidal ($T^+ = 325$, $W^+ = 40$) | 16.40 | 6.51 | 1.091 |

Table 6. The DNS validation of optimal sinusoidal actuation for DHT (case 3).

for DHT compared with drag reduction confirm that thermal flow control represents a distinct optimisation problem requiring dedicated investigation rather than straightforward extrapolation from momentum-transport results.

### 3.4. *Waveform optimisation for dissimilar heat transfer (case 4)*

This investigation addresses case 4, which constitutes the main contribution of this study: the systematic exploration of waveform optimisation for DHT through comprehensive parameter space exploration. The primary objective is to assess whether non-sinusoidal waveforms can surpass the DHT performance of the optimal sinusoidal baseline established in case 3, particularly when operating at reduced actuation amplitudes. This investigation tests the fundamental hypothesis that waveform topology optimisation provides an alternative, potentially more energy-efficient, pathway to DHT enhancement compared with the conventional strategy of amplitude escalation. This study represents the first reported systematic optimisation of non-sinusoidal waveforms for thermal flow control in the literature. Whilst cases 1–2 validated the framework through benchmarking against established optima for drag reduction, case 4 extends the analysis into unexplored parameter spaces wherein no prior optimal solutions exist. The outcomes therefore constitute novel contributions to the understanding of thermal–momentum-transport coupling under oscillatory wall forcing, with direct implications for energy-efficient thermal management applications.

The amplitude was fixed at $W^+_{max} = 30$ (representing a 25 % reduction below the optimal sinusoidal amplitude determined in case 3), and the PBO was applied to optimise $(T^+, a_1, a_2)$ for maximum $\overline{A_n}$. This strategic amplitude reduction serves dual purposes: it enables direct assessment of whether waveform optimisation can compensate for reduced actuation intensity, whilst simultaneously reducing optimisation dimensionality to concentrate computational effort on waveform parameter space exploration. The PBO converged rapidly to a novel quasi-plateau waveform within 40 generations (400 simulations), demonstrating remarkably low parameter uncertainty and achieving a best individual performance of $\overline{A_n} = 1.10$. The convergence characteristics are shown in figure 7, demonstrating systematic exploration and rapid convergence across all parameter dimensions. The exceptionally tight clustering in the $a_1$–$a_2$ parameter space (figure 7*b*) establishes that DHT exhibits sharply defined sensitivity to waveform topology, a fundamental characteristic distinguishing thermal- from momentum-transport optimisation. This behaviour contrasts markedly with the wide parameter dispersion observed for net energetic gain (case 2, figure 4*c*), wherein control costs dominated the objective function and rendered waveform shape relatively inconsequential. The present tight clustering demonstrates that thermal flow-control objectives exhibit fundamentally distinct sensitivity to waveform topology compared with energetically constrained momentum-transport problems: waveform optimisation provides substantially greater performance leverage for thermal objectives, wherein the precise temporal structure of spanwise actuation strongly influences the dissimilarity between thermal- and momentum-transport mechanisms. This observation supports the hypothesis that the temporal



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

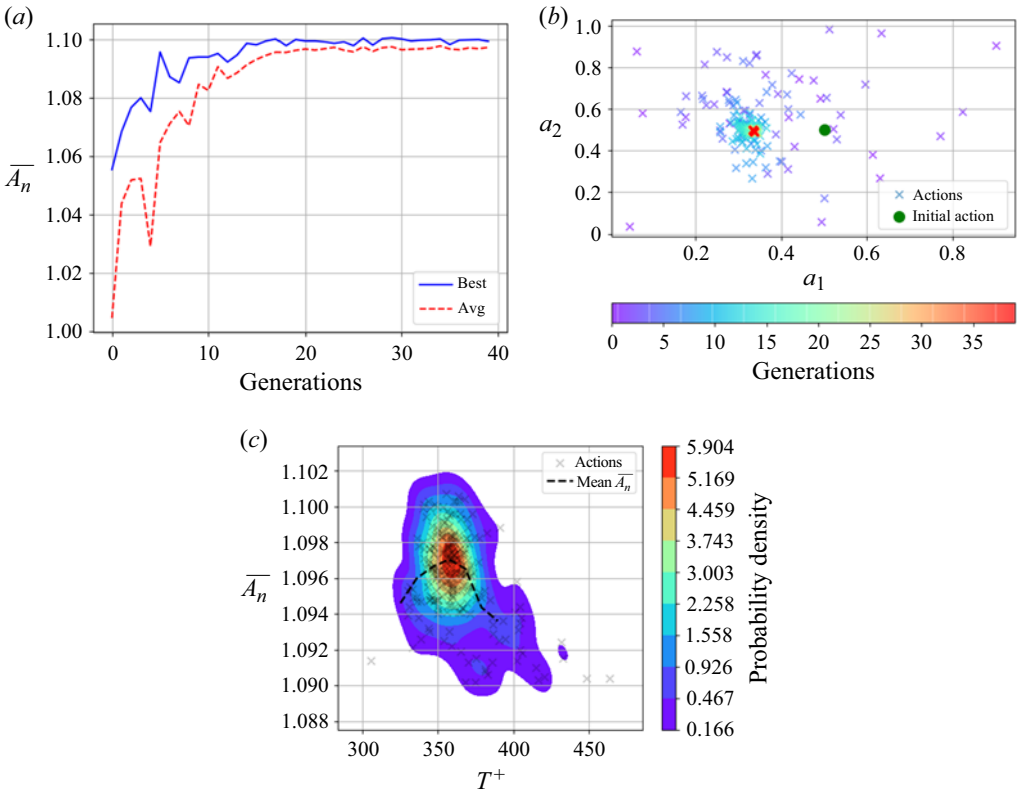



Figure 7. Convergence characteristics for DHT waveform optimisation (case 4). (*a*) Evolution of $\overline{A_n}$ through generations. (*b*) Parameter space exploration in $a_1$–$a_2$ coordinates showing exceptionally tight clustering. (*c*) Kernel density estimation revealing optimal period $T^+ \in [340, 365]$.

modulation of imposed spanwise shear, rather than mere amplitude escalation, constitutes the primary pathway for efficient DHT enhancement.

The optimal period was determined within $T^+ \in [340, 365]$ (figure 7*c*), representing a substantially narrower range than observed in case 3. The period–amplitude–waveform coupling is demonstrated by comparing the present optimal waveform at $W^+ = 30$ (converged at $T^+ \approx 350$) against an auxiliary sinusoidal optimisation performed in Guérin (2025) at the same amplitude, which identified optimal parameters ($T^+ = 450$, $W^+ = 30$) with $\overline{A_n} = 1.068$. This 100-unit period difference at fixed amplitude establishes that optimal period varies with waveform topology, whilst the intermediate position of both configurations relative to drag-reduction optima ($T^+ \approx 111$) and high-amplitude sinusoidal DHT ($T^+ \in [320, 375]$ for $W^+ = 40$) confirms period–amplitude coupling. This represents a fundamental distinction from drag reduction, wherein the optimal period remained largely invariant to waveform modifications ($T^+ \approx 111$ for both sinusoidal and plateau configurations). This period–amplitude–waveform coupling indicates that the temporal dynamics of Stokes-layer interaction with near-wall turbulence exhibits complex dependencies for thermal transport that differ fundamentally from momentum transport, thereby establishing thermal flow control as a distinct optimisation problem requiring dedicated parametric investigation. The optimal waveform (figure 8) exhibits a quasi-plateau topology analogous to that identified for drag reduction, characterised by extended constant-velocity phases at extremum amplitudes. Quantitative analysis reveals



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

| Configuration | $\%\Delta Nu$ | $\%\Delta C_f$ | $\overline{A_n}$ |
|---|---|---|---|
| Waveform ($T^+ = 350$, $W^+ = 30$) | 22.31 | 12.37 | **1.087** |

Table 7. The DNS validation of optimal waveform for DHT (case 4).

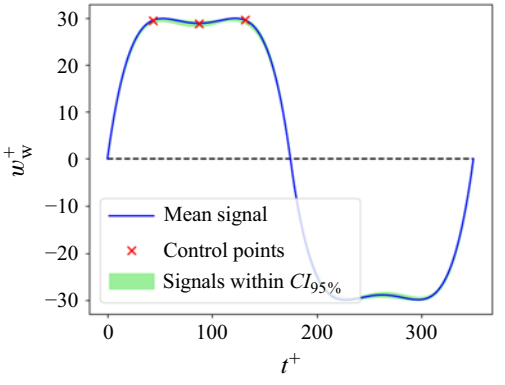



Figure 8. Optimal waveform for DHT (case 4) at $W^+ = 30$ exhibiting quasi-plateau topology.

that the waveform maintains $|w_w^+| > 0.95 W_{max}^+$ (i.e. within 5 % of maximum amplitude) for approximately 38 % of each half-period, thereby establishing sustained extremum-velocity phases that distinguish this topology from continuously varying sinusoidal signals. This extended plateau duration constitutes the defining characteristic enabling the observed thermal-transport enhancement through prolonged phases of slowly varying Stokes strain. The narrow 95 % confidence interval range confirms the high sensitivity of DHT performance to waveform parameters.

A DNS validation was performed at the optimal parameters ($T^+ = 350$, $W^+ = 30$), with results presented in table 7. The optimised waveform achieves $\overline{A_n} = 1.087$ with a substantial heat-transfer enhancement of $\%\Delta Nu = +22.31$ %.

A principal result of the present study emerges from a triangulated comparison isolating both waveform and amplitude effects. As established previously, an auxiliary sinusoidal optimisation at $W^+ = 30$ identified optimal parameters ($T^+ = 450$, $W^+ = 30$) with $\overline{A_n} = 1.068$ within 40 generations, demonstrating the period–amplitude–waveform coupling characteristic of DHT. This supplementary optimisation, employing the same PBO-LES framework with sinusoidal parameterisation (varying only $T^+$ and fixing $W^+ = 30$), provides the essential iso-amplitude reference point required to isolate waveform topology contributions from amplitude effects. Whilst this configuration is suboptimal relative to the best sinusoidal case at elevated amplitude (case 3), it enables rigorous decomposition of shape versus amplitude effects. Table 8 presents three configurations: the best sinusoidal configuration at $W^+ = 40$ (case 3), the sinusoidal baseline at $W^+ = 30$ (auxiliary optimisation) and the optimal waveform at $W^+ = 30$ (case 4). This structure enables two important comparisons: (i) at fixed amplitude ($W^+ = 30$), waveform versus sinusoid isolates the pure shape effect; (ii) across amplitudes, the waveform at $W^+ = 30$ versus sinusoid at $W^+ = 40$ demonstrates the efficiency advantage of waveform optimisation relative to amplitude escalation.

The triangulated comparison in table 8 establishes two complementary results that illuminate the mechanisms through which waveform optimisation enhances DHT efficiency.



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

| Configuration | $T^+$ | $W^+$ | $\overline{A_n}$ | $\%\Delta Nu$ | $\%P_{req}$ | $\%P_{net}$ | $\ell^+_{0.01}$ |
|---|---|---|---|---|---|---|---|
| Best sinusoidal | 325 | 40 | 1.091 | 16.16 | −499.1 | −505.43 | 54.18 |
| Sinusoidal baseline | 450 | 30 | 1.068 | 12.42 | −375.9 | −389.97 | 71.80 |
| Optimal waveform | 350 | 30 | **1.087** | **22.31** | −462.0 | **−474.65** | 62.45 |
| Comparison 1 – waveform effect at fixed amplitude ($W^+ = 30$): | | | | | | | |
| Waveform vs. sinusoid baseline | | | +1.8 % | +79.6 % | −22.9 % | −21.7 % | −13.0 % |
| Comparison 2 – efficiency advantage (waveform $W^+ = 30$ vs. best sinusoidal $W^+ = 40$): | | | | | | | |
| Waveform vs. best sinusoidal | | | −0.4 % | +37.9 % | +7.4 % | +6.1 % | +15.3 % |

Table 8. Performance comparison isolating waveform effects at fixed amplitude and efficiency gains relative to amplitude escalation. Blue: best sinusoidal configuration (case 3). Red: optimal waveform (case 4). Bold highlights superior performance in principal metrics. Negative $P_{req}$ and $P_{net}$ indicate energy expenditure (control cost). Lower values of $P_{net}$ are worse (higher net deficit). Protrusion height: $\ell^+_{0.01}$ is the wall-normal distance where spanwise velocity reaches 1 % of wall velocity amplitude

3.4.1. *Waveform-shape effect at fixed amplitude (comparison 1)*

Comparing the optimal waveform against the sinusoidal baseline (both at $W^+ = 30$) demonstrates that waveform topology independently contributes a 1.8 % improvement in $\overline{A_n}$ (from 1.068 to 1.087), achieved through 79.6 % greater heat-transfer enhancement ($\%\Delta Nu$ increases from 12.42 % to 22.31 %). This performance gain is accompanied by a 22.9 % increase in control power requirement ($P_{req}$ becomes more negative), attributable to the sustained extremum-velocity plateau phases characteristic of the optimised waveform topology. This comparison isolates the waveform-shape contribution from any amplitude effects, thereby confirming that the quasi-plateau topology possesses intrinsic advantages for DHT enhancement at fixed actuation intensity. The performance characteristics reveal distinct mechanistic pathways through which waveform optimisation enhances DHT efficiency: the temporal structure of spanwise actuation constitutes a more effective control lever for thermal transport than mere amplitude increase. The sustained plateau phases maintain constant extremum velocity over extended durations, thereby imposing prolonged periods of low strain rate that permit turbulence regeneration patterns fundamentally distinct from those induced by continuously varying sinusoidal forcing. The 22.9 % increase in control power relative to the sinusoidal baseline at identical amplitude represents the energetic cost of sustained plateau-phase actuation; however, this penalty is substantially offset by the 79.6 % enhancement in thermal transport, thereby establishing a more favourable energy-performance trade-off than that achievable through sinusoidal actuation at any amplitude.

3.4.2. *Efficiency advantage relative to amplitude increase (comparison 2)*

The optimal waveform at $W^+ = 30$ achieves virtually identical dissimilarity ($\overline{A_n} = 1.087$) to the best sinusoidal configuration at $W^+ = 40$ ($\overline{A_n} = 1.091$, deviation $< 0.4$ %), whilst employing 7.4 % less control energy ($P_{req} = -462.0$ versus $-499.1$) and producing 37.9 % greater absolute heat-transfer enhancement ($\%\Delta Nu = +22.31$ % versus $+16.16$ %). This demonstrates that waveform optimisation provides a more energy-efficient pathway to high-dissimilarity performance than sinusoidal actuation at elevated amplitude. The energy-performance efficiency can be quantified through the ratio of thermal enhancement to energetic expenditure: the waveform achieves $\Delta Nu/|P_{req}| = 0.0483$ compared with 0.0324 for the best sinusoidal configuration, representing a 49.1 % improvement in thermal-energy efficiency. This metric establishes that waveform optimisation provides substantially greater thermal benefit per unit of control energy invested, thereby offering a



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

more economically viable pathway for thermal management applications operating under energetic constraints. The optimised waveform achieves 6.1 % improved net energetic performance ($P_{net} = -474.65$ versus $-505.43$), confirming superior efficiency across all performance criteria when both drag penalty and actuation cost are incorporated.

3.4.3. *Decoupling of DHT performance from Stokes-layer protrusion height*

A phenomenon warranting particular scrutiny emerges from comparison of Stokes-layer protrusion height characteristics across the optimal configurations. The optimised waveform at $W^+ = 30$ exhibits Stokes-layer protrusion height $\ell^+_{0.01} = 62.45$, defined as the wall-normal distance where spanwise velocity reaches 1 % of the wall velocity amplitude, which must be evaluated relative to two distinct sinusoidal baselines to fully elucidate the observed behaviour. First, when compared against the sinusoidal baseline at identical amplitude ($W^+ = 30$, with $\ell^+_{0.01} = 71.8$), the waveform displays reduced protrusion height (13 % decrease) whilst simultaneously achieving superior DHT performance ($\overline{A_n} = 1.087$ versus 1.068). Second, when compared against the best sinusoidal configuration at elevated amplitude ($W^+ = 40$, with $\ell^+_{0.01} = 54.18$), the waveform exhibits enhanced protrusion height (15 % increase) alongside comparable dissimilarity ($\overline{A_n} = 1.087$ versus 1.091). This comparison reveals a fundamental decoupling of DHT performance from protrusion height: the sinusoidal champion at $W^+ = 40$ achieves superior performance ($\overline{A_n} = 1.091$) despite possessing the shallowest protrusion ($\ell^+_{0.01} = 54.18$) amongst all three configurations, whilst the sinusoidal baseline at $W^+ = 30$ exhibits the deepest protrusion ($\ell^+_{0.01} = 71.8$) whilst delivering the poorest DHT performance ($\overline{A_n} = 1.068$). Since protrusion height scales primarily with period ($\ell^+_{0.01} \propto \sqrt{T^+}$) rather than amplitude, the observed ranking – wherein shorter periods lead to both shallower protrusion and superior performance – establishes that spatial penetration depth does not constitute the primary mechanism for DHT enhancement. The optimal configurations converge to moderate periods ($T^+ \in [325, 350]$ for cases 3–4) rather than the extended periods ($T^+ = 450$ for sinusoidal $W^+ = 30$) that would maximise protrusion height, demonstrating that DHT optimisation favours temporal modulation characteristics over spatial penetration depth. One interpretation of this behaviour is that spatial protrusion height and temporal modulation of imposed shear constitute independent mechanisms for transport enhancement. Whilst sinusoidal signals at extended periods achieve deep spatial penetration, this advantage proves insufficient to compensate for suboptimal temporal actuation characteristics. Conversely, waveform topology optimisation operates through temporal modulation of the imposed spanwise shear gradient ($\partial w/\partial y$), wherein the quasi-plateau configuration maintains sustained extremum velocity over extended durations, thereby imposing prolonged phases of slowly varying Stokes strain. This temporal structure enables preferential thermal-transport enhancement through mechanisms fundamentally distinct from spatial protrusion, thereby decoupling DHT performance from protrusion height. The observation that optimal configurations systematically favour moderate periods over extended periods, despite the latter producing deeper spatial penetration, establishes temporal actuation structure as the dominant control parameter for DHT enhancement.

Waveform optimisation therefore represents a fundamentally more effective and energy-efficient strategy for DHT enhancement than amplitude increase. The optimal waveform exhibits superior performance across all evaluated metrics when control costs are accounted for, thereby establishing waveform topology as the foremost design parameter for thermal flow-control applications where both heat-transfer augmentation and energy efficiency are of utmost importance.



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

| Case | Objective | Strategy | Optimal Config. | Outcome |
|---|---|---|---|---|
| 1 | Max $P_{sav}$ | Validation | Plateau ($T^+ = 111$) | Framework validated |
| 2 | Max $P_{net}$ | Validation | Sinusoidal ($T^+ = 140$) | Framework validated |
| 3 | Max $\overline{A_n}$ | Discovery | Sinusoidal ($T^+ = 325$, $W^+ = 40$) | Baseline established |
| 4 | Max $\overline{A_n}$ | Discovery | Plateau ($T^+ = 350$, $W^+ = 30$) | Novel optimum found |

Table 9. Summary of outcomes across all four cases.

Table 9 synthesises the principal outcomes across all four optimisation cases investigated in this study, demonstrating the progression from validation through established benchmarks to discovery of novel optimal configurations for thermal flow control.

### 3.5. *Phase-resolved analysis of Stokes-layer temporal organisation*

The preceding optimisation results established that the quasi-plateau waveform at $T^+ = 350$, $W^+ = 30$ achieves substantial DHT enhancement whose performance decouples from the Stokes-layer penetration depth (§ 3.4), indicating that transport mechanisms fundamentally distinct from those governing spatial protrusion are operative. Analysis of sublayer dynamics leads to a thinning ratio $R_{\alpha,\nu} = 1.043$, defined as the ratio of relative viscous- to thermal-sublayer thinning under actuation, which only marginally exceeds the value of unity corresponding to proportional thinning; sublayer thinning is accordingly excluded as the primary driver of DHT performance. The Fukagata, Iwamoto & Kasagi (2002) identity decomposition performed by Guérin *et al.* (2024) further demonstrated that forcing-term contributions to the dissimilarity are negligible, thereby identifying modifications to the turbulent-transport terms as the origin of the observed Reynolds analogy breakage. The present simulations are conducted at $Pr = 1$, the condition under which molecular diffusivities are identical and the Reynolds analogy is most robust; any observed dissimilarity must therefore arise from structural differences in the turbulent-transport equations rather than from disparate diffusivities, constituting the most demanding possible test of the DHT mechanism.

To demonstrate that the temporal organisation of the Stokes strain, rather than its spatial penetration depth, governs the degree of dissimilarity, a comparative phase-resolved analysis is presented between the optimal sinusoidal configuration ($T^+ = 325$, $W^+ = 40$) and the optimal waveform ($T^+ = 350$, $W^+ = 30$). These two cases are selected as they achieve comparable dissimilarity ($\overline{A_n} = 1.091$ and 1.087, respectively) through fundamentally different actuation strategies: the sinusoidal configuration relies upon elevated amplitude and deep Stokes-layer penetration ($\ell^+_{0.01} = 54.18$), whereas the plateau waveform operates at reduced amplitude with shallower penetration ($\ell^+_{0.01} = 62.45$) and must compensate through temporal modulation of the imposed shear. The physical context is provided by the interaction between the Stokes strain and the self-sustaining process (SSP) of near-wall turbulence (Hamilton *et al.* 1995; Waleffe 1997), whose regeneration cycle operates on a characteristic time scale $\tau^+_{SSP} \approx 100$–$200$. A triple decomposition framework is employed throughout, whereby the instantaneous field $f(\boldsymbol{x}, t)$ is separated into three distinct components

$$f(\boldsymbol{x}, t) = \underbrace{\overline{f}(\boldsymbol{x})}_{\textit{Time average}} + \underbrace{\hat{f}(\boldsymbol{x}, t)}_{\textit{Periodic fluctuation}} + \underbrace{f''(\boldsymbol{x}, t)}_{\textit{Stochastic fluctuation}}, \tag{3.1}$$



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

where $\overline{f}$ denotes the time average over the complete observation period. To isolate the coherent motion locked to the actuation cycle, the phase-averaging operator is introduced. For an actuation period $T$, the phase average of a quantity $f$ at a specific normalised phase $t^* = t/T$ is defined as the ensemble average over $N$ oscillation cycles

$$\widetilde{f}(\boldsymbol{x}, t^*) \equiv \langle f(\boldsymbol{x}, t)\rangle_\phi = \frac{1}{N}\sum_{n=0}^{N-1} f(\boldsymbol{x}, t + nT). \tag{3.2}$$

The phase-averaged field thus comprises both the time-mean and the coherent periodic component, $\widetilde{f} = \overline{f} + \hat{f}$, whilst the stochastic turbulent fluctuation is defined as the deviation of the instantaneous field from its phase average, $f'' = f - \widetilde{f}$. By construction, $\widetilde{f''} = 0$; however, the phase average of products of stochastic fluctuations, such as the phase-averaged Reynolds stresses $\widetilde{u''u''}$, is non-zero and varies throughout the actuation cycle, thereby reflecting the modulation of turbulence intensity by the Stokes strain. The ensemble averaging is performed over $N = 20$ complete oscillation cycles for both the sinusoidal and waveform cases.

### 3.5.1. *Stokes-layer structure and near-wall response*

The phase-resolved evolution of the near-wall fields and the Stokes-layer spatio-temporal distribution is presented in figure 9, comparing the sinusoidal case ($T^+ = 325$, $W^+ = 40$; left column) and the optimal waveform ($T^+ = 350$, $W^+ = 30$; right column). The first row displays the temporal distribution of the phase-averaged quantities $\widetilde{u}^+$, $\widetilde{\theta}^+$ and $\widetilde{w}^+$ at $y^+ \approx 1$; the second row presents the phase-resolved maps of the spanwise velocity $\widetilde{w}^+$; the third row presents the Stokes strain rate (SSR), defined as $\partial^2\widetilde{w}^+/\partial y^+\partial t^+$, which quantifies the temporal rate of change of the imposed Stokes shear.

All quantities in figures 9(*a*) and 9(*b*) are normalised by the friction velocity $u_{\tau,0}$ and friction temperature $\theta_{\tau,0}$ of the unactuated reference case. At $y^+ \approx 1$ and $Pr = 1$, the unactuated fields satisfy $u^+ = \theta^+ = 1$ identically; deviations from unity therefore quantify the effect of the wall oscillation on the respective transport mechanisms, with values exceeding unity indicating enhancement of skin friction or Nusselt number relative to the unactuated baseline. In the sinusoidal case (figure 9*a*), $\widetilde{w}^+$ varies smoothly and continuously, and both $\widetilde{u}^+$ and $\widetilde{\theta}^+$ exhibit periodic oscillations synchronised with the spanwise wall displacement. The phase-averaged temperature $\widetilde{\theta}^+$ exceeds $\widetilde{u}^+$ at all phases, with substantially larger peak-to-trough excursions (approximately 0.35 for $\widetilde{\theta}^+$ versus 0.22 for $\widetilde{u}^+$), confirming that preferential thermal enhancement persists throughout the entire actuation cycle rather than being restricted to isolated phase intervals. A notable temporal asymmetry is observed: the descent from maximum to minimum transport occupies a longer fraction of the cycle than the subsequent recovery, indicating that turbulence suppression proceeds more gradually than its recommencement following the Stokes strain reversal. Furthermore, the $\widetilde{\theta}^+$ peaks arrive slightly before the $\widetilde{u}^+$ peaks within each half-cycle, suggesting that the thermal field responds more rapidly to the release of suppression than the velocity field. During a narrow phase interval, $\widetilde{\theta}^+$ dips marginally below unity, indicating a transient suppression of heat transfer that partially offsets the enhancement achieved during the remainder of the cycle.

In the optimal-waveform case (figure 9*b*), the plateau phases of nearly constant $\widetilde{w}^+$ are clearly discernible, separated by rapid reversal intervals. Both $\widetilde{u}^+$ and $\widetilde{\theta}^+$ attain substantially larger peak values than in the sinusoidal case, and their temporal profiles are qualitatively different: both fields display a more symmetric rise-and-fall pattern, and the phase ordering is reversed relative to the sinusoidal case, with $\widetilde{u}^+$ peaking slightly



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

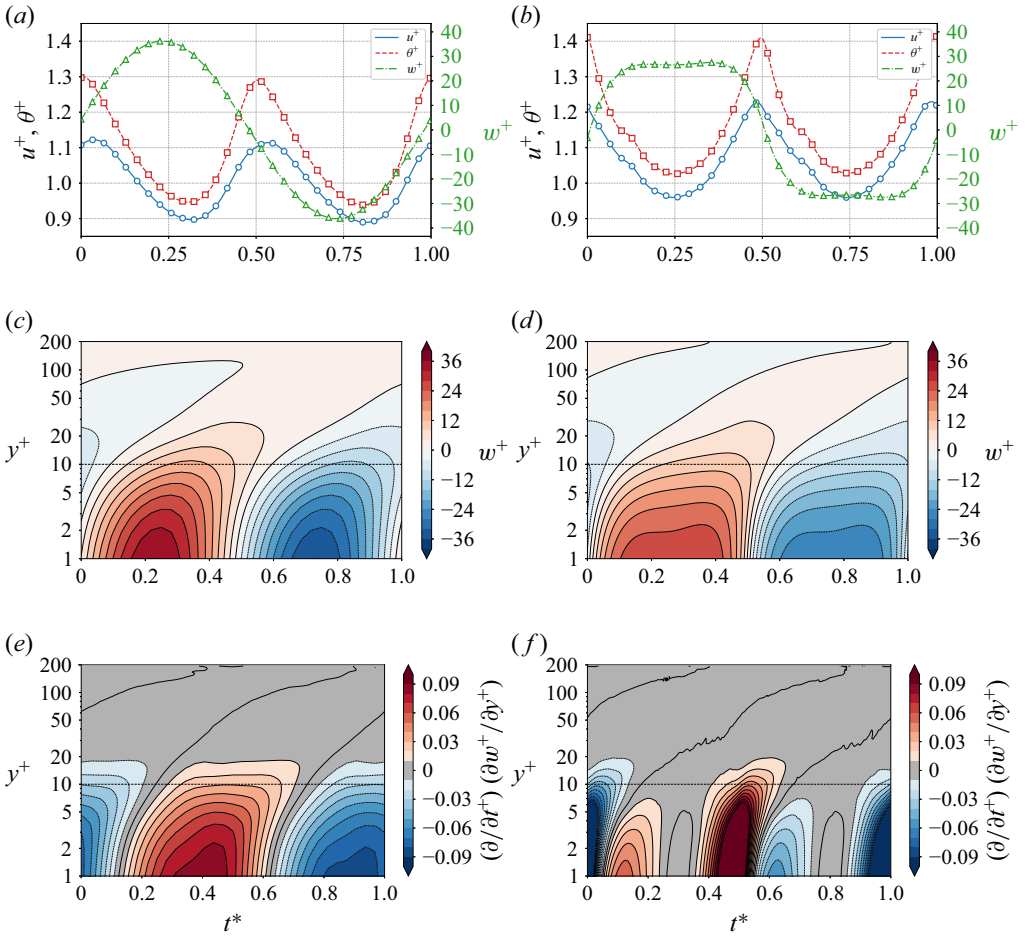



Figure 9. Comparative phase-resolved Stokes-layer structure. Left column: sinusoidal case ($T^+ = 325$, $W^+ = 40$). Right column: optimal waveform ($T^+ = 350$, $W^+ = 30$). (*a*, *b*) Temporal distribution of the phase-averaged wall quantities $\widetilde{u}^+$, $\widetilde{\theta}^+$ and $\widetilde{w}^+$ at $y^+ \approx 1$, normalised by the unactuated friction scales; deviations from unity quantify the forcing effect on skin friction and heat transfer. (*c*, *d*) Phase-resolved maps ($y^+$ vs $t^*$) of the spanwise velocity $\widetilde{w}^+$, revealing the spatial extent and temporal structure of the Stokes layer. (*e*, *f*) Stokes strain rate $\mathrm{SSR} = \partial^2\widetilde{w}^+/\partial y^+\partial t^+$: grey-shaded regions indicate $|\mathrm{SSR}| < 0.01$, corresponding to lingering phases where the Stokes strain remains approximately constant; concentrated bands indicate rapid reversal intervals. The dashed horizontal line denotes $y^+ = 10$.

before $\widetilde{\theta}^+$. This reversal of the peak phasing between the two configurations may indicate that the mechanisms governing the temporal coupling between velocity and temperature responses differ under sinusoidal and plateau actuation, an aspect warranting further investigation. During the plateau phases, both quantities decrease initially as the buffer-layer Stokes strain weakens, before undergoing a pronounced recovery associated with the recommencement of the SSP; $\widetilde{\theta}^+$ exhibits a markedly stronger amplification than $\widetilde{u}^+$ during this recovery, and the gap between the two quantities widens progressively through the late-plateau phase, reaching its maximum at the transport peaks. In contrast to the sinusoidal case, $\widetilde{\theta}^+$ remains above unity throughout the entire cycle, indicating that the Nusselt number exceeds its unactuated value at every phase without exception. This persistent excess of $\widetilde{\theta}^+$ over $\widetilde{u}^+$, observed in both configurations but more pronounced in the waveform case, constitutes the phase-resolved manifestation of the cycle-averaged



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

dissimilarity $\overline{A_n} > 1$ documented in table 8, and establishes that preferential thermal enhancement is concentrated during phases of high turbulent activity associated with the late-plateau recovery.

The SSR is recalled to be defined as $\mathrm{SSR} = \partial^2 \widetilde{w}^+ / \partial y^+ \partial t^+$, i.e. the mixed spatio-temporal derivative of the spanwise velocity, which quantifies the rate at which the imposed Stokes shear evolves in time at a given wall-normal position. When $|\mathrm{SSR}| \approx 0$, the local Stokes strain $\partial \widetilde{w}/\partial y$ remains approximately constant: the flow experiences a sustained, nearly time-invariant spanwise shear rather than a continuously evolving one. This regime is designated as lingering, in analogy with the sustained forcing conditions characteristic of spatial steady actuation (Viotti, Quadrio & Luchini 2009). The lingering fraction of the cycle may accordingly be quantified as the temporal interval during which the SSR magnitude remains below a prescribed threshold in the buffer-layer region where turbulent structures are most active ($y^+ \gtrsim 10$). In figures 9(*e*) and 9(*f*), the grey-shaded regions delineate precisely this condition, corresponding to phases and wall-normal positions where $|\mathrm{SSR}| < 0.01$; the extent of these grey regions therefore provides a direct visual measure of lingering. In the sinusoidal case (figure 9*e*), the SSR oscillates smoothly and symmetrically throughout the cycle, with approximately equal fractions of the period devoted to positive and negative strain-rate values; the strain is continuously varying, and the grey regions of near-zero SSR are confined to narrow transitional bands between the positive and negative lobes. In the optimal-waveform case (figure 9*f*), the SSR is concentrated in narrow, intense bands during the reversal intervals, separated by extensive grey regions during the plateau phases. In the buffer-layer domain above $y^+ \approx 10$, i.e. the region populated by the turbulent structures of the SSP, these grey regions extend over a substantial fraction of the actuation cycle, demonstrating that the Stokes strain remains effectively constant throughout the plateau at wall-normal positions where the vorticity-tilting suppression mechanism operates. During these lingering phases, the buffer-layer strain $\partial \widetilde{w}/\partial y$, initially strong following the preceding reversal, diffuses and diminishes despite the sustained wall velocity; this progressive weakening of the vorticity-tilting mechanism (Agostini *et al.* 2015) permits the SSP to recommence during the latter portion of the plateau, affording sufficient time for differential turbulence regeneration to develop. The contrast between the sinusoidal and plateau SSR patterns establishes that it is the temporal organisation of the strain, i.e. the duration and intensity of lingering phases relative to reversal intervals, rather than spatial penetration depth, that distinguishes the two actuation strategies. This temporal organisation governs the fraction of the actuation cycle during which the SSP operates under conditions permitting differential resumption of momentum and thermal production.

### 3.5.2. *Phase modulation of velocity and temperature variance*

The phase-resolved distributions of the streamwise velocity variance $\widetilde{u''u''}^+$ and the temperature variance $\widetilde{\theta''\theta''}^+$ are conveyed by figure 10, comparing the sinusoidal and waveform configurations. Both quantities attain their maximum values in the vicinity of $y^+ \approx 10$. All four panels share the same colour scale ($0 \leqslant \cdot \leqslant 8$), permitting direct amplitude comparison across cases and fields. The variance magnitudes are substantially larger in the waveform case (figures 10*b* and 10*d*) than in the sinusoidal case (figures 10*a* and 10*c*), and in both configurations the temperature variance $\widetilde{\theta''\theta''}^+$ attains higher peak values than the velocity variance $\widetilde{u''u''}^+$, consistent with the preferential thermal enhancement documented in the preceding subsection. Beyond the amplitude differences,



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

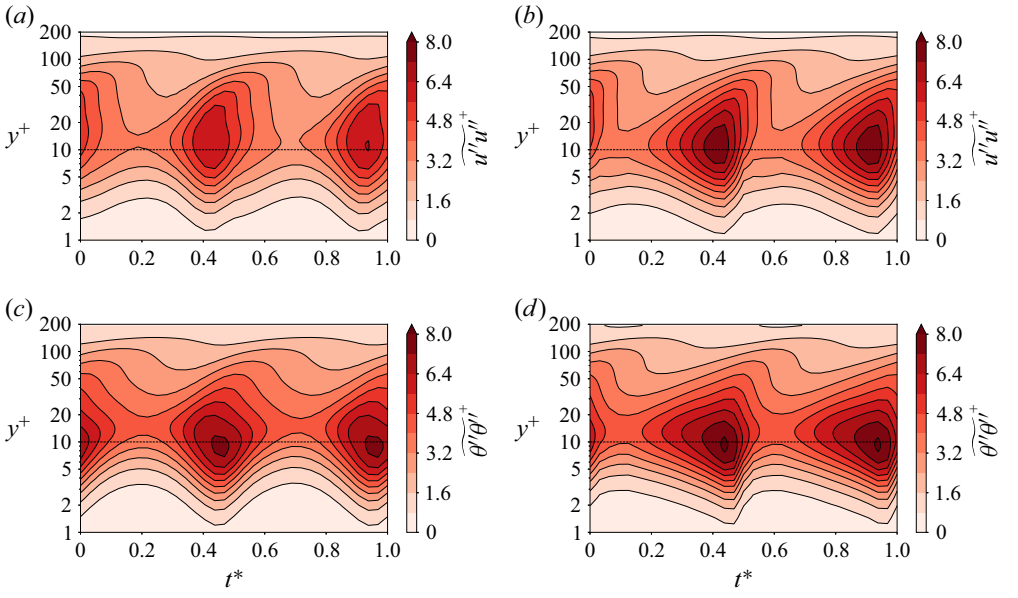



Figure 10. Phase-resolved maps ($y^+$ vs $t^*$) of the stochastic variance fields. Left column: sinusoidal case ($T^+ = 325$, $W^+ = 40$). Right column: optimal waveform ($T^+ = 350$, $W^+ = 30$). (*a*, *b*) Streamwise velocity variance $\widetilde{u''u''}^+$. (*c*, *d*) Temperature variance $\widetilde{\theta''\theta''}^+$. The dashed horizontal line denotes $y^+ = 10$.

the temporal symmetry of the variance modulation differs between the two configurations and requires detailed examination. In the sinusoidal case (figures 10*a* and 10*c*), the variance modulation appears approximately symmetric at first glance. However, closer scrutiny at $y^+ \approx 10$ reveals a subtle asymmetry, most clearly discernible in $\widetilde{u''u''}^+$ (figure 10*a*): the variance increase (regeneration phase) commences earlier and proceeds more rapidly than the subsequent decrease (suppression phase), which extends over a longer fraction of the cycle. This asymmetry is also reflected in the shape of the isocontours below $y^+ = 10$ and is consistent with the temporal asymmetry of $\widetilde{u}^+$ observed in figure 9(*a*), where the descent from maximum to minimum transport was noted to occupy a longer fraction of the cycle than the recovery. The same asymmetry is present in $\widetilde{\theta''\theta''}^+$ (figure 10*c*), although less pronounced.

In the optimal-waveform case (figures 10*b* and 10*d*), this temporal asymmetry is drastically amplified. The variance regeneration occupies an extended portion of the cycle during the lingering phase, with the increase commencing promptly as the Stokes strain weakens and building progressively throughout the plateau. The suppression phase, by contrast, is compressed into a narrow temporal interval coinciding with the rapid strain reversal. This pronounced asymmetry between a protracted regeneration and a brief suppression, far more extreme than the subtle imbalance observed in the sinusoidal case, is a direct consequence of the plateau waveform topology, which concentrates strain-rate variations into brief reversal intervals whilst maintaining near-constant strain over extended lingering phases. In the waveform case, the velocity-temperature asymmetry is further accentuated: the temperature variance maximum is more pronounced relative to its velocity counterpart, and the region of elevated $\widetilde{\theta''\theta''}^+$ extends further into the second half of the plateau.



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

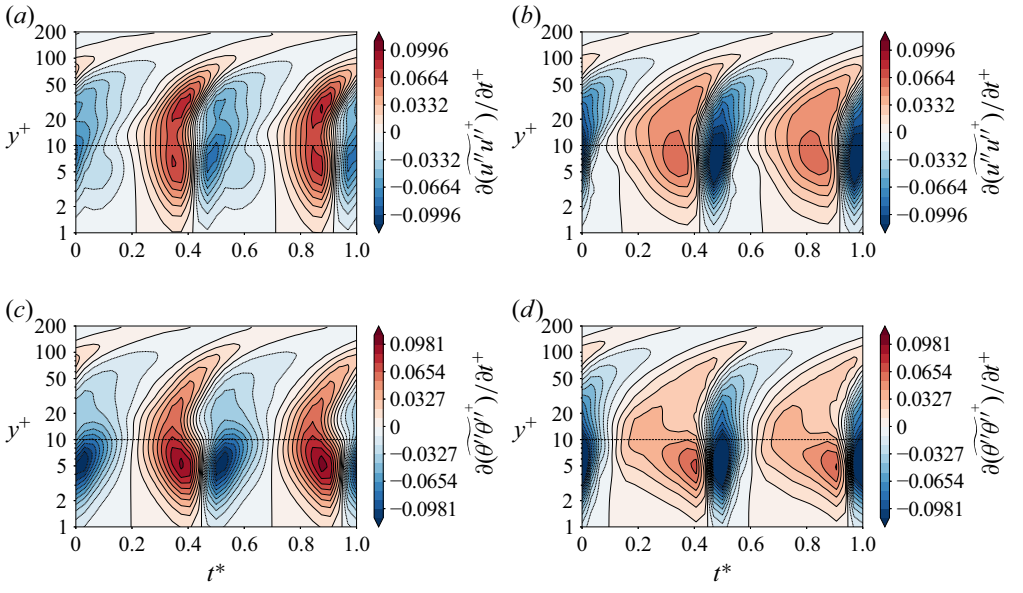



Figure 11. Phase-resolved maps ($y^+$ vs $t^*$) of the variance temporal derivatives. Left column: sinusoidal case ($T^+ = 325$, $W^+ = 40$). Right column: optimal waveform ($T^+ = 350$, $W^+ = 30$). (*a*, *b*) Rate of change of streamwise velocity variance $\partial\widetilde{u''u''}^+/\partial t^+$. (*c*, *d*) Rate of change of temperature variance $\partial\widetilde{\theta''\theta''}^+/\partial t^+$. Positive values (red) indicate variance strengthening; negative values (blue) indicate weakening. The dashed horizontal line denotes $y^+ = 10$.

To sharpen these observations and establish a direct connection between the lingering Stokes strain and the differential variance modulation, the phase-resolved temporal derivatives $\partial\widetilde{u''u''}^+/\partial t^+$ and $\partial\widetilde{\theta''\theta''}^+/\partial t^+$ are examined in figure 11; positive values indicate phases of variance strengthening (regeneration), whilst negative values indicate variance weakening (suppression). In the sinusoidal case (figures 11*a* and 11*c*), the strengthening and weakening phases alternate smoothly, with regions of broadly comparable temporal extent and magnitude for both the velocity and temperature fields; a subtle asymmetry is nevertheless discernible, consistent with the variance maps (figures 10*a* and 10*c*), whereby the strengthening lobes are slightly more compact and intense than the weakening lobes. In the optimal-waveform case (figures 11*b* and 11*d*), the rate-of-change patterns differ substantially: variance weakening is concentrated in narrow temporal bands coinciding with the reversal intervals, during which the rapidly changing Stokes strain actively disrupts coherent structures, whilst variance strengthening occupies extended temporal intervals during the lingering phases identified in figure 9(*f*), during which the near-constant Stokes strain permits the SSP to recommence and turbulent structures to regenerate. The asymmetry between the velocity and temperature rates of change becomes apparent upon comparison of the right-column panels: the strengthening of $\partial\widetilde{\theta''\theta''}^+/\partial t^+$ (figure 11*d*) is sustained over a broader temporal interval and extends further from the wall than its velocity counterpart $\partial\widetilde{u''u''}^+/\partial t^+$ (figure 11*b*) during the lingering phases. Furthermore, the transition from strengthening to weakening at the onset of the reversal occurs earlier for the velocity variance than for the temperature variance, producing a temporal offset that provides direct evidence of asymmetric regeneration and constitutes the phase-resolved fingerprint of the dissimilarity mechanism. A plausible

https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

explanation for this differential behaviour resides in the structural difference between the two variance transport equations: the streamwise velocity variance budget contains a pressure-strain redistribution term that transfers energy between velocity components through the divergence-free constraint, whereas no analogous mechanism exists in the scalar variance equation. If this redistribution acts as an additional sink for $\widetilde{u''u''}^{+}$ at the reversal, the velocity variance would be suppressed sooner whilst the temperature variance, governed by an equation devoid of such a term, would continue to strengthen unimpeded over a slightly extended window, consistent with the observed temporal offset. Verification of this hypothesis requires a full phase-resolved decomposition of the variance budgets, which is deferred to a dedicated investigation.

### 3.5.3. *Synthesis and physical interpretation*

The comparative phase-resolved analysis across figures 9, 10 and 11 establishes the following physical sequence connecting waveform topology to DHT. The quasi-plateau waveform concentrates the Stokes strain-rate variations into narrow reversal intervals, creating extended lingering phases (figure 9*f*) during which the near-wall Stokes strain remains approximately constant. During these lingering intervals, the progressive weakening of the vorticity-tilting suppression permits the SSP to resume, and both velocity and temperature variance regenerate (figures 11*b* and 11*d*). The temperature variance regenerates more effectively than the velocity variance, as demonstrated by the broader temporal extent and greater wall-normal reach of the positive thermal rate-of-change regions relative to their velocity counterparts. At the reversal, the velocity variance weakens earlier than the temperature variance, producing a temporal offset that generates an excess of thermal variance over velocity variance.

The origin of this asymmetric regeneration is hypothesised to reside in the fundamental difference between the transport equations governing the two variance fields: the streamwise velocity variance budget contains a pressure-strain redistribution term $\Pi_{uu}$ that transfers energy between velocity components through the divergence-free constraint ($\boldsymbol{\nabla} \cdot \boldsymbol{u} = 0$), whereas no analogous redistribution mechanism exists in the temperature variance equation, which evolves under a pure advection–diffusion dynamics. Confirmation of this hypothesis and quantification of the respective roles of the diagonal and off-diagonal pressure-strain terms throughout the actuation cycle require a comprehensive phase-resolved investigation of the turbulent stress and scalar variance budgets, which is beyond the scope of the present study and constitutes a natural direction for future research.

The lingering mechanism thereby provides a resolution to the protrusion paradox identified in § 3.4: the optimal waveform achieves comparable dissimilarity to the high-amplitude sinusoidal configuration despite reduced spatial penetration, as the prolonged lingering of the Stokes strain affords the SSP more time to develop the asymmetric regeneration conditions that produce preferential thermal enhancement. The sinusoidal configuration, operating at elevated amplitude, achieves deep spatial penetration and produces a strong Stokes strain throughout the cycle; the smooth, continuous variation of this strain, however, does not provide extended phases of near-constant strain during which the asymmetric regeneration can develop fully. The quasi-plateau waveform compensates for its reduced amplitude by concentrating the strain variation into brief reversal intervals, thereby maximising the fraction of the actuation cycle during which lingering conditions prevail and the conjectured asymmetry between momentum and scalar variance regeneration can develop.



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

## 4. Conclusion

This study introduced and validated a novel framework for optimising the temporal waveform topology of SWO for turbulent flow control. Through the association of PBO with implicit LES and polynomial-based signal parameterisation, the optimisation methodology was extended beyond conventional period and amplitude parameters to encompass signal shape as a continuous design variable. The framework was applied to establish optimal actuation strategies for DHT, a problem for which no prior systematic optimisation had been performed. A principal result is that the optimal-waveform topology depends fundamentally on the control objective: when energy expenditure is accounted for, the drag-reduction optimal waveform is sinusoidal ($P_{net}$ maximisation), whereas the DHT-optimal waveform is a quasi-plateau signal characterised by sustained extremum-velocity phases separated by rapid reversals. This contrast demonstrates that signal shape constitutes a genuine design variable whose optimal configuration is dictated by the underlying physics of the targeted transport mechanism, rather than being universal across objectives. The investigation establishes four principal contributions, progressing from identification of the optimal forcing to the elucidation of the temporal organisation mechanism governing preferential thermal enhancement.

(i) First, rigorous validation through benchmark drag-reduction problems confirmed the framework's capacity to identify both complex optimal waveforms when they provide performance advantages (quasi-plateau topology for $P_{sav}$ maximisation, recovering results of Cimarelli *et al.* 2013) and simple optimal solutions when parametric complexity offers no benefit (near-sinusoidal configurations for $P_{net}$ maximisation, confirming Guérin *et al.* 2025). This bidirectional validation established confidence for application to unexplored thermal flow-control problems.

(ii) Second, the first systematically optimised sinusoidal baseline for DHT was established through comprehensive exploration of the period–amplitude parameter space, achieving $\overline{A_n} = 1.091$ at $T^+ = 325$ and $W^+ = 40$. This represents a 2.5 % improvement over the non-optimised baseline reported by Guérin *et al.* (2024) and provides a rigorous performance benchmark against which waveform optimisation can be evaluated. The optimal period ($T^+ \in [320, 375]$) occupies an intermediate position between drag-reduction optima ($T^+ \approx 110$) and thermal maximisation optima ($T^+ \approx 1600$ as reported in Guérin 2025), thereby establishing DHT as a distinct optimisation problem requiring dedicated parametric investigation, rather than straightforward extrapolation from momentum-transport results.

(iii) Third, the optimised quasi-plateau waveform at reduced amplitude ($W^+ = 30$) achieves virtually identical dissimilarity ($\overline{A_n} = 1.087$) to the optimal sinusoidal configuration at elevated amplitude ($W^+ = 40$, $\overline{A_n} = 1.091$), whilst employing 7.4 % less control energy and producing 37.9 % greater absolute heat-transfer enhancement. Triangulated comparison at fixed amplitude confirms that waveform shape alone contributes 79.6 % greater heat-transfer enhancement relative to the sinusoidal baseline, establishing that observed performance gains are intrinsic to signal topology. The energy-performance efficiency, quantified as $\Delta Nu/|P_{req}| \approx 0.048$ compared with 0.032 for optimal sinusoidal forcing, represents approximately 50 % improvement in thermal benefit per unit of control energy invested. This constitutes the first optimised non-sinusoidal waveform for thermal flow control reported in the literature. Furthermore, DHT performance was observed to decouple from the Stokes-layer protrusion height: the optimal sinusoidal configuration at $W^+ = 40$ achieves superior dissimilarity despite possessing the shallowest protrusion ($\ell^+_{0.01} = 54.18$) amongst all configurations, whilst the sinusoidal baseline at $W^+ = 30$ exhibits



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

the deepest protrusion ($\ell^+_{0.01} = 71.8$) with the weakest performance, a paradox whose resolution is provided by the phase-resolved analysis of the fourth contribution.

(iv) Fourth, a comparative phase-resolved analysis between the optimal sinusoidal ($T^+ = 325$, $W^+ = 40$) and optimal-waveform ($T^+ = 350$, $W^+ = 30$) configurations identified the temporal organisation of the Stokes strain as the governing parameter for dissimilarity. The quasi-plateau waveform concentrates the strain-rate variations into narrow reversal intervals, creating extended lingering phases during which the near-constant Stokes strain permits the SSP of near-wall turbulence to recommence. During these lingering phases, the temperature variance regenerates more effectively and over a broader temporal interval than the velocity variance, and the transition from variance strengthening to weakening at the onset of the reversal occurs earlier for the velocity field than for the temperature field. This asymmetric regeneration is hypothesised to originate in the structural difference between the momentum and scalar transport equations: the streamwise velocity variance budget contains a pressure-strain redistribution term $\Pi_{uu}$ that may drain energy through the divergence-free constraint at the Stokes strain reversal, whereas no analogous redistribution mechanism exists in the temperature variance equation. The lingering mechanism resolves the protrusion paradox identified in the third contribution: DHT enhancement is governed not by spatial penetration depth but by the fraction of the actuation cycle during which conditions favourable to asymmetric variance regeneration prevail.

These results place the DHT mechanism in a broader context when contrasted with drag-reduction physics. At the drag-reduction optimal period ($T^+ \approx 100$–$125$), the plateau duration is insufficient for the buffer-layer strain to decay, and the SSPs remains suppressed throughout; superior drag reduction is then achieved by compressing the reversal intervals during which the SSP may resume (Cimarelli *et al.* 2013). At the substantially longer DHT-optimal period ($T^+ = 350$), the progressive weakening of the buffer-layer Stokes strain during the extended plateau affords sufficient time for the SSP to recommence, and it is this differential resumption of momentum and thermal variance regeneration that generates the observed preferential thermal enhancement. Waveform topology optimisation for DHT therefore operates through a fundamentally different physical pathway than drag-reduction optimisation: cyclic temporal modulation of the conditions governing asymmetric variance regeneration, rather than duty-cycle compression of SSP suppression. Quantification of the role of $\Pi_{uu}$ in mediating this asymmetry, including the phase at which it transitions to negative values and the temporal extent over which it acts as a velocity variance sink, requires a dedicated phase-resolved budget analysis and constitutes a principal objective of future work.

The demonstrated performance characteristics possess direct implications for energy-efficient thermal management applications. Waveform optimisation enables performance comparable to high-amplitude sinusoidal actuation whilst simultaneously reducing energy consumption and enhancing absolute thermal transport by 37.9 %. This dual benefit, maintained dissimilarity with improved efficiency, addresses the design challenge for applications wherein both heat-transfer enhancement and energy efficiency are of considerable importance, such as CSP receivers and thermal management systems operating under energetic constraints.

**Acknowledgements.** This project was provided with computing HPC and storage resources by GENCI at TGCC thanks to the grants AD012A14284, A0172A07624 and A0152A07624 on the supercomputer Joliot Curie's SKL partition.



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

**Funding.** This work was supported by the French National Research Agency (ANR) under grant ANR-21-CE50-0031. Additionally, the last author received funding from the ANR Young Researchers (JCJC) programme under grant ANR-23-CE46-0004.

**Data availability statement.** The simulations were performed using the open-source Xcompact3d framework, publicly available at https://github.com/xcompact3d/Incompact3d. The PBO optimisation algorithm implementation and the data supporting the results of this study are available from the corresponding author upon reasonable request.

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

## Appendix A. Injectivity of the waveform parameterisation function

The purpose of this appendix is to provide a formal mathematical proof for the uniqueness of the signal generation method described in § 2.1. Specifically, it is demonstrated that the function $\mathcal{L}$, which maps the parameters $(a_1, a_2, T^+, W^+_{max})$ to the resulting signal shape (represented by the polynomial $P_1(t)$), is injective. Injectivity, in this context, guarantees that distinct combinations of input parameters consistently produce distinct output signal shapes, ensuring that no two different parameter sets lead to the exact same waveform. Establishing this property is essential for the optimisation process. It ensures a well-defined search space by confirming that each point explored by the optimisation algorithm corresponds to a unique wall signal. This uniqueness eliminates ambiguity and supports the reliable convergence of the algorithm towards an optimal parameter set.

*Proposition.* The application $\mathcal{L}$ defined as

$$\begin{aligned} \mathcal{L} : [0,\,1] \times [0,\,1] \times [10,\,750] \times\, ]0,\,30] &\to \mathbb{R} \\ (a_1,\, a_2,\, T^+,\, W^+_{max}) &\mapsto w_w(t) \end{aligned} \tag{A1}$$

is injective.

*Proof.* Since the second half-period of the signal is an antisymmetric reflection of the first, proving injectivity for the first half-period suffices to establish injectivity for the entire signal. Let us assume that $Q(t)$ (the first half-period polynomial of $w_w(t)$) is the output of two different sets of parameters, i.e.

$$Q(t) = \mathcal{L}(a_1, b_1, T_1, W_1) = \mathcal{L}(a_2, b_2, T_2, W_2). \tag{A2}$$

Given that $W_1 > 0$ and $W_2 > 0$, the equalities $T_1 = T_2$ and $W_1 = W_2$ follow directly, as the period and maximal amplitude of $Q(t)$ are uniquely determined by these parameters.

Let $\{y_{1i}\}_{i=1,\ldots,5}$ represent the pre-scaling control point weights associated with parameters $a_1$ and $b_1$, and $\{y_{2i}\}_{i=1,\ldots,5}$ those arising from $a_2$ and $b_2$ (see (2.2) for the mapping). Since the same scaled polynomial $Q(t)$ is produced by these control weights, they must be equal to each other after scaling is applied, i.e. they must be proportional before this operation. In other words, there exists $k \in \mathbb{R}$ such that

$$y_{1i} = k\, y_{2i} \quad \forall\, i \in \{1, \ldots, 5\}. \tag{A3}$$

Now, as defined in (2.2), the sum of the control weights in both cases is equal to 1. That is,

$$\sum_{i=1}^{5} y_{1i} = k \sum_{i=1}^{5} y_{2i} = 1. \tag{A4}$$



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

Since $\sum_{i=1}^{5} y_{2i} = 1$, it follows that $k = 1$. Thus, it is concluded that

$$y_{1i} = y_{2i} \quad \forall\, i \in \{1, \ldots, 5\}. \tag{A5}$$

Given that the mapping of (2.2) is evidently injective (each unique pair $(a_1, a_2)$ produces a unique set of weights $\{y_i\}$ satisfying the constraint $\sum_i y_i = 1$), it can now be deduced that $a_1 = a_2$ and $b_1 = b_2$. Thus, it has been proven that

$$\mathcal{L}(a_1, b_1, T_1, W_1) = \mathcal{L}(a_2, b_2, T_2, W_2) \quad \Longrightarrow \quad (a_1, b_1, T_1, W_1) = (a_2, b_2, T_2, W_2), \tag{A6}$$

which demonstrates that the application $\mathcal{L}$ is injective.

Therefore, it has been formally demonstrated that the signal parameterisation method employed, based on Lagrangian interpolation points, leads to a unique mapping from the parameter space $(a_1, a_2, T^+, W^+_{max})$ to the resulting polynomial $P_1(t)$ (or equivalently to the spanwise wall velocity signal $w_w(t)$). This proven injectivity validates the robustness of the parameter space definition and underpins the reliability of the optimisation process. The constraint $\sum_{i=1}^{5} y_i^1 = 1$ is fundamental to ensuring this property, as it prevents degenerate cases where different parameter combinations could produce identical waveforms through proportional scaling. This mathematical guarantee ensures that the PBO algorithm explores a well-defined, non-redundant parameter space during optimisation.

## Appendix B. Validation of Lagrangian parameterisation against DNS benchmarks

Before employing the Lagrangian interpolation method for waveform optimisation, it was essential to verify that the parameterisation introduces negligible approximation errors when generating control signals. The validation strategy consisted in testing the method's ability to reproduce a signal whose performance characteristics are precisely known from high-fidelity DNS. A standard sinusoidal waveform was selected as the test case, with parameters $T^+ = 125$ and $W^+ = 18$ corresponding to well-documented drag-reduction conditions (Quadrio & Ricco 2004). The validation procedure involved three distinct simulations for comparison: (i) DNS results from the literature providing the reference benchmark (ground truth), (ii) a true sinusoidal signal implemented directly in the LES framework ($S_1$ configuration) and (iii) a Lagrangian-approximated sinusoid ($S_{1,\mathrm{lag}}$) generated through the parameterisation method, with waveform parameters $(a_1, a_2)$ optimised to minimise the $L^2$ deviation from a pure sine wave. The primary validation criterion was whether the polynomial approximation could reproduce the performance metrics ($P_{sav}$, $P_{req}$, $P_{net}$) of the true sinusoid within acceptable tolerances.

Table 10 presents the comprehensive comparison across all three configurations. The Lagrangian-approximated signal achieved $P_{sav} = 39.27$ %, in near-perfect agreement with the DNS reference value of 39.3 % reported by Quadrio & Ricco (2004). Strong consistency was observed across all performance metrics, with the $S_{1,\mathrm{lag}}$ results closely matching both the DNS benchmark and the true sinusoidal LES. Minor discrepancies between the required power ($P_{req}$) and the net power ($P_{net}$), each remaining below 2 %, are attributable to the mathematical constraint that a fourth-order polynomial cannot analytically reproduce an exact sinusoidal function, resulting in local deviations in wall shear stress. This close correspondence establishes confidence in the reliability of the Lagrangian parameterisation for two important aspects of the subsequent optimisation work. First, it demonstrates that the method does not introduce systematic biases that could mislead the optimisation algorithm towards spurious optima. Second, it confirms that the parameterisation possesses sufficient expressiveness to accurately represent smooth,



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

| Configuration | $P_{sav}$ ( %) (2.7) | $P_{req}$ ( %) (2.8) | $P_{net}$ ( %) (2.9) |
|---|---|---|---|
| DNS (Quadrio & Ricco 2004) | 39.3 | −161.2 | −121.9 |
| $S_1$ (true sinusoid) | 41.3 | −162.1 | −120.8 |
| $S_{1,\mathrm{lag}}$ (Lagrangian approx.) | 39.27 | −158.1 | −118.83 |

Table 10. Validation of Lagrangian parameterisation method. Comparison of drag reduction ($P_{sav}$), required power ($P_{req}$) and net energy gain ($P_{net}$) between DNS reference (ground truth), true sinusoidal LES ($S_1$) and Lagrangian-approximated sinusoid ($S_{1,\mathrm{lag}}$) at oscillation parameters $T^+ = 125$, $W^+ = 18$.

continuous waveforms across the parameter space. Having validated the signal-generation method on a known benchmark, the framework can be applied with confidence to discover optimal waveforms where ground-truth solutions are not yet established.

## Appendix C. Maximising heat-transfer enhancement with control-cost penalty

The ability of SWO to achieve significant DHT has been demonstrated in the main text. However, the optimisations presented did not incorporate the energetic cost of control within their reward functions — an essential consideration for practical applications.

This appendix presents preliminary investigations aimed at achieving heat-transfer enhancement ($Nu_r > 1$) whilst ensuring a positive global relative gain in energy. The relative gain in heat transfer is expressed as $Nu_r - 1$, whilst the energetic gains due to drag variations are represented by $P_{net}$, accounting for control costs through $P_{req}$ as defined in (2.9).

An initial objective is to satisfy

$$(Nu_r - 1) + P_{net} > 0, \tag{C1}$$

which ensures that the control strategy yields a net energetic benefit. However, this approach has limitations, as it directly compares thermal-energy gains with momentum energy losses due to friction and control costs. In some applications, the importance of thermal-energy gains may outweigh that of momentum losses, leading to overall efficiency improvements even when (C1) is not achieved. To avoid ambiguity, an ideal control strategy would result in

$$Nu_r > 1 \quad \text{and} \quad P_{net} > 0, \tag{C2}$$

ensuring simultaneous improvements in both thermal and momentum efficiency.

Two PBO investigations were conducted to explore these objectives, focusing on achieving at least the goal defined in (C1). The optimisations were performed on a sinusoidal signal, over 10 generations with 8 individuals each, which was sufficient to reach a plateau in the reward function and assess the feasibility of achieving efficient DHT with control-cost constraints. The parameter space was constrained to $T^+ \in [10, 750]$ and $W^+ \in [3, 30]$, using the PBO configuration detailed in § 2.2. The lower bound $W^+ > 3$ is imposed to prevent premature convergence to the unactuated case, particularly in the case where (C1) should prove challenging to achieve.

### *Investigation 1: combined thermal and momentum performance optimisation*

The first optimisation process aimed to maximise $(Nu_r - 1) + P_{net}$ by adopting it as the reward function. The best rewards plateaued at −0.019, indicating that a net energetic gain is not achievable under the current framework, considering both temperature and momentum equally. This also implies that the objective defined in (C2) remains



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

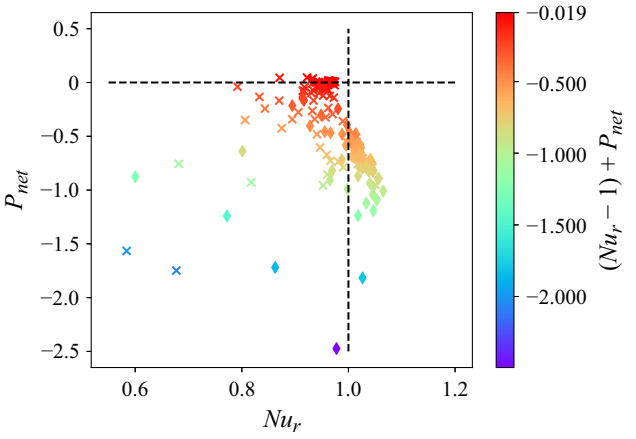



Figure 12. Summary of all individuals obtained during the first (crosses; minimisation of (C1)) and second (diamonds; minimisation of (C3)) PBO investigations, focusing on optimal DHT with a control-cost penalty. The colour bar represents the combined relative gains in thermal and momentum energy, including control costs, whilst the black dashed lines indicate the thresholds for positive gains in $P_{net}$ and $Nu_r$.

unattainable, as no individual simultaneously satisfies $P_{net} > 0$ and $Nu_r > 1$. The best-performing individuals exhibited $P_{net} \in [0, 0.04]$ and $Nu_r \in [0.935, 0.98]$, suggesting that the optimisation prioritises favourable $P_{net}$ values at the expense of heat-transfer reduction (see figure 12). This highlights the significant energetic penalty associated with heat-transfer enhancement. The best-performing individuals also exhibit $\overline{A_n}$ values within the range [0.996, 1.004], reinforcing the observation that the optimisation prioritises $P_{net}$ over dissimilarity. Furthermore, the agent consistently favours longer periods ($T^+ > 400$) and converges toward the lower boundary of $W^+ = 3$, strongly suggesting that the unactuated case may indeed be optimal under the given constraints.

*Investigation 2: prioritising heat-transfer enhancement using sigmoid reward*

The second optimisation focused on achieving Nusselt number enhancement whilst still incorporating a control-cost penalty, which may be particularly relevant in certain industrial contexts. The reward function for this investigation is defined as

$$(Nu_r - 1) + P_{net} + \frac{1}{1 + e^{-100(Nu_r - 1)}}. \tag{C3}$$

The sigmoid term acts as a soft constraint, creating a strong incentive for the agent to push $Nu_r$ above a threshold. As illustrated in figure 13, this term features a steep slope for $Nu_r$ values between 0.95 and 1.05. This sharp transition encourages the agent to achieve a Nusselt number enhancement of at least 5 %, ensuring this objective is prioritised over the maximisation of $P_{net}$.

As anticipated, none of the individuals achieved the objective defined in (C1) (see figure 12). However, the best-performing individuals did attain moderate heat-transfer enhancement ($Nu_r \in [1.014, 1.047]$ and $\overline{A_n} \in [1.022, 1.036]$), albeit with corresponding $P_{net}$ values in the range $[-0.83, -0.46]$. Whilst this approach yields lower combined values of $P_{net} + (Nu_r - 1)$ compared with the first case (with the best score being $-0.45$ among the top-performing individuals and $-0.22$ globally), it ensures that the control strategy successfully achieves the targeted heat-transfer enhancement, along with moderate DHT. Notably, the agent converges toward the upper boundary of $T^+ = 750$,



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

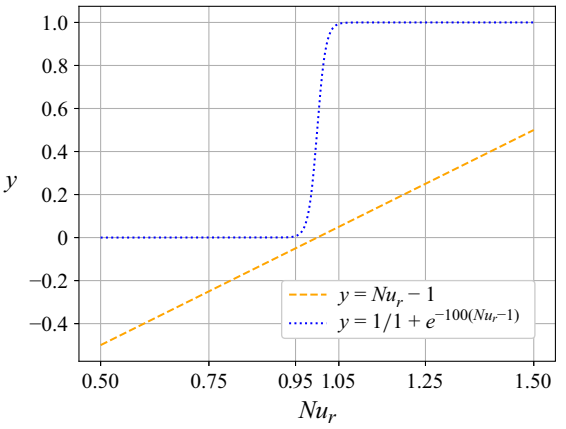



Figure 13. Illustration of the $Nu_r$ terms in (C3) reward function. The sigmoid term creates a strong incentive to push $Nu_r$ above unity.

indicating that longer periods may improve outcomes, and stabilises around $W^+ \approx 17$, reflecting a reduction in the $P_{net}$ penalty due to the reward function design.

The results, illustrated in figure 12, confirm that neither (C1) nor (C2) was achieved in either optimisation process, underscoring the limitations of SWO in delivering effective heat-transfer enhancement when control costs and drag penalties are considered. The figure also reveals that even modest heat-transfer enhancements incur substantial energetic costs, with $P_{net}$ dropping below $-0.5$. Larger enhancements above 5 % further reduce momentum net energetic gains, reaching as low as $P_{net} \approx -1.1$. In summary, this appendix has explored the feasibility of achieving DHT enhancement whilst accounting for the energetic cost of control—an essential consideration for practical applications. The investigations revealed that, under the current optimisation framework, simultaneously satisfying both heat-transfer enhancement ($Nu_r > 1$) and a positive net energetic gain ($P_{net} > 0$) remains challenging. Additionally, achieving the easier objective $(Nu_r - 1) + P_{net} > 0$ also proved unattainable. The results underscore the inherent trade-off between thermal and momentum energy gains, with even modest improvements in heat transfer often accompanied by significant energy costs.

However, this study remains preliminary and the inclusion of waveform parameter optimisation or the consideration of larger oscillation periods may pave the way for more favourable outcomes. Future work could incorporate these changes or even consider closed-loop control implementations of SWO optimised with deep reinforcement learning methods. Whilst the energetic penalty associated with enhanced heat transfer can be substantial, such enhancements may still hold considerable value in specific applications where thermal performance is prioritised. This highlights the importance of tailoring control strategies to the unique demands of each application through reward shaping, balancing the relative significance of thermal and momentum considerations.

## References

Agostini, L., Touber, E. & Leschziner, M. 2015 The turbulence vorticity as a window to the physics of friction-drag reduction by oscillatory wall motion. *Intl J. Heat Fluid Flow* **51**, 3–15.

Bartholomew, P., Deskos, G., Frantz, R.A., Schuch, F.N., Lamballais, E. & Laizet, S. 2020 Xcompact3D: an open-source framework for solving turbulence problems on a Cartesian mesh. *SoftwareX* **12**, 100550.

Blesbois, O., Chernyshenko, S.I., Touber, E. & Leschziner, M.A. 2013 Pattern prediction by linear analysis of turbulent flow with drag reduction by wall oscillation. *J. Fluid Mech.* **724**, 607–641.



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press

CAVALLAZZI, G.M., GUASTONI, L., VINUESA, R. & PINELLI, A. 2024 Deep reinforcement learning for the management of the wall regeneration cycle in wall-bounded turbulent flows. *Flow Turbul. Combust.* **115** (3), 1291–1317.

CIMARELLI, A., FROHNAPFEL, B., HASEGAWA, Y., DE ANGELIS, E. & QUADRIO, M. 2013 Prediction of turbulence control for arbitrary periodic spanwise wall movement. *Phys. Fluids* **25** (7), 075102.

DAVID, M., TOUTANT, A. & BATAILLE, F. 2023 Study of asymmetrically heated flows passing through gas-pressurized solar receivers using Direct Numerical Simulations. *Intl J. Heat Mass Transfer* **201**, 123577.

FUKAGATA, K., IWAMOTO, K. & KASAGI, N. 2002 Contribution of Reynolds stress distribution to the skin friction in wall-bounded flows. *Phys. Fluids* **14** (11), L73-L76.

GATTI, D. & QUADRIO, M. 2016 Reynolds-number dependence of turbulent skin-friction drag reduction induced by spanwise forcing. *J. Fluid Mech.* **802**, 553–582.

GHRAIEB, H., VIQUERAT, J., LARCHER, A., MELIGA, P. & HACHEM, E. 2021 Single-step deep reinforcement learning for open-loop control of laminar and turbulent flows. *Phys. Rev. Fluids* **6** (5), 053902.

GUÉRIN, L. 2025 Improving solar-receiver efficiency by manipulating turbulence in heat-collecting, fluid-conveying channels: from physical understanding to machine-learning-based optimization. PhD thesis, Université de Poitiers. Available at: https://sdrive.cnrs.fr/s/TqJtbAj5GFWad6A.

GUÉRIN, L., CORDIER, L., FLAGEUL, C., GRIEU, S. & AGOSTINI, L. 2025 Policy-based optimization for drag reduction via spanwise wall oscillations. *Neural Comput. Appl.* **37** (16), 9451–9479.

GUÉRIN, L., FLAGEUL, C., CORDIER, L., GRIEU, S. & AGOSTINI, L. 2024 Preferential enhancement of convective heat transfer over drag via near-wall turbulence manipulation using spanwise wall oscillations. *Intl J. Heat Fluid Flow* **110**, 109564.

HAMILTON, J.M., KIM, J. & WALEFFE, F. 1995 Regeneration mechanisms of near-wall turbulence structures. *J. Fluid Mech.* **287**, 317–348.

HANSEN, N. 2023 The CMA evolution strategy: a tutorial. arXiv: 1604.00772.

JIMÉNEZ, J. 2013 Near-wall turbulence. *Phys. Fluids* **25** (10), 101302.

KADER, B.A. & YAGLOM, A.M. 1981 Temperature and concentration profiles in fully turbulent boundary layers. *Intl J. Heat Mass Transfer* **24** (9), 1541–1544.

KASAGI, N., TOMITA, Y. & KURODA, A. 1992 Direct numerical simulation of Passive scalar Field in a turbulent channel flow. *J. Heat Transfer* **114** (3), 598–606.

KINGMA, D.P. & BA, J. 2014 Adam: a method for stochastic optimization. arXiv: 1412.6980.

LAIZET, S. & LAMBALLAIS, E. 2009 High-order compact schemes for incompressible flows: a simple approach. *J. Comput. Phys.* **228** (16), 5989–6015.

LAMBALLAIS, E., FORTUNÉ, V. & LAIZET, S. 2011 Straightforward high-order numerical dissipation via the viscous term for direct and large eddy simulation. *J. Comput. Phys.* **230** (9), 3270–3275.

LIGRANI, P.M., OLIVEIRA, M.M. & BLASKOVICH, T. 2003 Comparison of heat transfer augmentation techniques. *AIAA J.* **41** (3), 337–362.

QUADRIO, M. & RICCO, P. 2004 Critical assessment of turbulent drag reduction through spanwise wall oscillations. *J. Fluid Mech.* **521**, 251–271.

QUADRIO, M., RICCO, P. & VIOTTI, C. 2009 Streamwise-travelling waves of spanwise wall velocity for turbulent drag reduction. *J. Fluid Mech.* **627**, 161–178.

RICCO, P., SKOTE, M. & LESCHZINER, M.A. 2021 A review of turbulent skin-friction drag reduction by near-wall transverse forcing. *Prog. Aerosp. Sci.* **123**, 100713.

TOUBER, E. & LESCHZINER, M. 2012 Near-wall streak modification by spanwise oscillatory wall motion and drag-reduction mechanisms. *J. Fluid Mech.* **693**, 150–200.

UCHINO, K., MAMORI, H. & FUKAGATA, K. 2017 Heat transfer in fully developed turbulent channel flow with streamwise traveling wave-like wall deformation. *J. Therm. Sci. Tech-JPN.* **12** (1), JTST0003–JTST0003.

VIOTTI, C., QUADRIO, M. & LUCHINI, P. 2009 Streamwise oscillation of spanwise velocity at the wall of a channel for turbulent drag reduction. *Phys. Fluids* **21** (11), 115109.

VIQUERAT, J., RABAULT, J., KUHNLE, A., GHRAIEB, H., LARCHER, A. & HACHEM, E. 2023 Policy-based optimization: single-step policy gradient method seen as an evolution strategy. *Neural Comput. Appl.* **35** (1), 523–537.

WALEFFE, F. 1997 On a self-sustaining process in shear flows. *Phys. Fluids* **9** (4), 883–900.

YAMAMOTO, A., HASEGAWA, Y. & KASAGI, N. 2013 Optimal control of dissimilar heat and momentum transfer in a fully developed turbulent channel flow. *J. Fluid Mech.* **733**, 189–220.



https://doi.org/10.1017/jfm.2026.11660 Published online by Cambridge University Press