

# Optimal Parallelization Strategies for Active Flow Control in Deep Reinforcement Learning-Based Computational Fluid Dynamics

Wang Jia<sup>1</sup> and Hang Xu<sup>1</sup>

State Key Lab of Ocean Engineering, School of Naval Architecture, Ocean and Civil Engineering, Shanghai Jiao Tong University, Shanghai, 200240, China

(\*Electronic mail: hangxu@sjtu.edu.cn)

(Dated: 30 April 2024)

Deep Reinforcement Learning (DRL) has emerged as a promising approach for handling highly dynamic and nonlinear Active Flow Control (AFC) problems. However, the computational cost associated with training DRL models presents a significant performance bottleneck. To address this challenge and enable efficient scaling on high-performance computing architectures, this study focuses on optimizing DRL-based algorithms in parallel settings. We validate an existing state-of-the-art DRL framework used for AFC problems and discuss its efficiency bottlenecks. Subsequently, by deconstructing the overall framework and conducting extensive scalability benchmarks for individual components, we investigate various hybrid parallelization configurations and propose efficient parallelization strategies. Moreover, we refine input/output (I/O) operations in multi-environment DRL training to tackle critical overhead associated with data movement. Finally, we demonstrate the optimized framework for a typical AFC problem where near-linear scaling can be obtained for the overall framework. We achieve a significant boost in parallel efficiency from around 49% to approximately 78%, and the training process is accelerated by approximately 47 times using 60 central processing unit (CPU) cores. These findings are expected to provide valuable insights for further advancements in DRL-based AFC studies. Consequently, it continues to be a prominent and actively studied problem of significant interest.

## I. INTRODUCTION

Active flow control (AFC)<sup>1,2</sup> holds significant practical value in delaying transition<sup>3</sup>, reducing separation<sup>4</sup>, enhancing performance<sup>5</sup>, lowering resistance<sup>6</sup>, intensifying turbulence<sup>7-9</sup>, and suppressing noise<sup>10</sup>. Efficient flow control systems can improve vehicle performance and fuel efficiency while promoting economically efficient, environmentally friendly, and competitive industrial production<sup>11,12</sup>, thereby making flow control technology a central focus of attention in the academic and engineering communities in fields like fluid mechanics for the past several decades<sup>6,13,14</sup>. While AFC has made significant advancements in practice, it still faces various challenges, including the design and optimization of control strategies, real-time requirements, computational complexity, uncertainties, robustness issues, and the effects of multi-physics coupling<sup>15-17</sup>. Overcoming the challenges in AFC requires further research in efficient mathematical modeling, robust control strategies, and algorithm optimization to enhance system performance<sup>18-21</sup>. This involves developing accurate and efficient mathematical models that capture the complex dynamics of flow control systems, designing control strategies capable of handling uncertainties and variations, and improving the computational efficiency of control algorithms<sup>22,23</sup>. By focusing on these areas, it is expected to overcome challenges, improve performance, reduce energy consumption, and enhance system functionality across various industries.

The rapid advancement of artificial intelligence has introduced novel approaches and implementation pathways for closed-loop AFC<sup>6</sup>. Particularly, deep reinforcement learning (DRL) methods, which facilitate autonomous learning and decision-making through interactions with the environment<sup>24-27</sup>, align well with the requirements of closed-loop AFC<sup>28</sup>. The fundamental concept of DRL involves continuous trial-and-error learning as the intelligent agent interacts with the environment, progressively enhancing its decision-making capabilities. It utilizes deep neural networks to approximate value or policy functions, thereby mapping observations to actions. This approach, through ongoing interaction with the environment, data collection, and optimization, enables the intelligent agent to acquire the optimal decision strategy, thereby harnessing artificial intelligence algorithms to optimize and make decisions in flow control systems. In literature, Gazzola, Hejazialhosseini, and Koumoutsakos<sup>29</sup>, Novati *et al.*<sup>30</sup>, Verma, Novati, and Koumoutsakos<sup>31</sup>, and Yan *et al.*<sup>32</sup> successively conducted a series of studies that combined numerical simulations with DRL algorithms to investigate how fish can enhance their propulsion efficiency by utilizing the wake vortices generated by other swimmers. These promising findings provided an initial glimpse into the capability of DRL algorithms to address complex physical problems.

Many researchers have acknowledged that computational efficiency plays a crucial role in effectively combining numerical methods with learning algorithms, such as DRL, in terms of feasibility and practicality. Ma *et al.* are recognized as pioneers in integrating reinforcement learning with flow control<sup>33</sup>. Their groundbreaking study on coupled control system, utilizing DRL for interactions between fluid and rigid bodies in a 2D setting, has revolutionized flow control and paved the way for innovative implementations of AFC. Rabault *et al.* achieved remarkable results in AFC by training an artificial neural network using a DRL agent, successfully stabilizing vortex shedding and reducing resistance by approximately 8%<sup>34</sup>. Rabault and Kuhnle further addressed the limitation of training speed caused by the computational fluid dynamics (CFD) component in training DRL

agents by proposing a parallelization solution that combines parallelization techniques for both the numerical simulation and the DRL algorithm, resulting in significant speedup<sup>35</sup>. Thanks to the open-source code released by Rabault *et al.* and his team<sup>34</sup>, there has been a remarkable surge in scholarly publications utilizing DRL algorithms to address fluid dynamics challenges, especially since 2020. Until now, DRL has demonstrated successful applications in various AFC tasks, encompassing the 1D Kuramoto–Sivashinsky equation<sup>36</sup>, 1D falling fluid flow<sup>37</sup>, 2D Rayleigh–Bénard convection<sup>38</sup>, 2D convection control in confined space<sup>39</sup>, 3D model cavity with one heat source<sup>40</sup>, 2D flow over a cylinder oscillating around its axis<sup>41</sup> and the interaction between rotating small cylinders and the wake of a main cylinder<sup>42</sup>.

Despite the proclaimed advantages of DRL-based methods, various challenges persist and prevent widespread implementation of these methods. One major bottleneck is the excessive training cost associated with the DRL-based framework. As reported by Rabault *et al.*, their codebase allows to train the artificial neural network to perform AFC on a 2D cylinder in about 24 hours, using a modern CPU running on a single core<sup>34</sup>. Koizumi, Tsutsumi, and Shima utilized the DRL framework for a similar control problem where they spent 48 hours on training, using a 32-core processor<sup>43</sup>. Evidently, to tackle more complex and demanding AFC problems, the cost of training will quickly become prohibitively expensive and parallelization techniques are necessary to scale the framework for larger problems. Unfortunately, the scaling performance of DRL-based CFD problems remains a largely unexplored research field. Rabault and Kuhnle discovered that the DRL agents can learn from multiple environments launched in parallel with similar rate of convergence, but without any scaling analysis, it remains unclear whether multi-environment training translates into a proportional reduction in the actual training time<sup>35</sup>. The closest performance analysis on this topic in literature is the work by Kurz *et al.*, who introduced an Reinforcement Learning-augmented CFD solver, showing 30% parallel efficiency on 16 nodes in a HPC system<sup>44</sup>. Nonetheless, to the best of our knowledge, there is no existing work that comprehensively analyzes the optimal form of parallelization for maximum efficiency in a DRL-based AFC problem. In this work, we want to provide a thorough analysis for different forms of parallelization, which is supported by scaling performance benchmarks, in order to identify performance bottlenecks and find ways to further optimize the current state of the art. Ultimately, through tackling performance issues, we wish to shed some light on the best practices for optimal DRL-based training in the context of AFC problems.

The outline of this paper is as follows. In Section II, the problem description is presented, followed by an in-depth overview of DRL and its formulation in addressing the AFC challenge. The available parallelization strategies in DRL-based AFC problems are also described, along with a discussion of their software implementations. In Section III, a validation study is conducted, followed by a detailed scaling analysis of both CFD simulations and the DRL framework within multi-environment settings. This analysis is then extended to the examine hybrid parallelization techniques that strive to balance CFD parallelization and DRL parallelization for optimal efficiency. In particular, the potential for further performance enhancement through the refinement of input/output (I/O) operations in Section III D is also investigated. Finally, Section IV concludes the paper by summarizing the key findings, discussing their implications, and suggesting potential directions for future research.

## II. PROBLEM DESCRIPTION AND METHODOLOGY

Here, we investigate the application of a DRL-based AFC technique to control flow separation over a circular cylinder, including the numerical setup, DRL algorithm, formulation of the AFC problem, hybrid parallelization strategy, and the software platform integrating CFD and DRL.

### A. Numerical simulation

The configuration adopted in this work for simulating the flow around a circular cylinder closely follows the classical benchmark by Schäfer *et al.*<sup>45</sup> We consider the two-dimensional, unsteady, incompressible and viscous flow around a circular cylinder. The corresponding governing equations are the Navier-Stokes equations (in their dimensionless form)

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot (\nabla \mathbf{u}) = -\nabla p + \frac{1}{Re} \Delta \mathbf{u}, \quad (1)$$

$$\nabla \cdot \mathbf{u} = 0. \quad (2)$$

Here,  $\mathbf{u} = (U, V)$  represents the fluid velocity, where  $U$  and  $V$  are the velocity components along the  $x$ - and  $y$ -axes, respectively.  $t$  denotes the dimensionless time unit.  $p$  is the thermodynamic pressure.  $Re = \frac{\bar{U}D}{\nu}$  is the Reynolds number, where  $\bar{U}$  is the mean velocity at the inlet,  $D$  is the diameter of the cylinder and  $\nu$  is the kinematic viscosity.

As shown in Fig. 1(a), the cylinder is submerged in a rectangular domain with dimensions of  $22D$  (along the  $x$ -axis) and  $4.1D$  (along the  $y$ -axis), with the origin of the Cartesian coordinate system established at the center of the cylinder. The cylinder is positioned slightly off-centered in the  $y$ -direction in order to trigger vortex shedding. The boundaries of the computational




FIG. 1: Description of the numerical setup. (a) Computational domain ( $22D \times 4.1D$ ) of flow around a cylinder. (b) Boundary conditions for the computational domain; (c) Details of the boundary conditions on the cylinder. Jets are located at  $\Omega = 90^\circ$  and  $\Omega = 270^\circ$  on the cylinder, with a jet width of  $\omega = 10^\circ$ . Parabolic velocity distribution is used for each jet.

domain are divided into an inlet  $\Gamma_{in}$ , an outlet  $\Gamma_{out}$ , no-slip walls  $\Gamma_w$ , and two separate jets on the cylinder  $\Gamma_i$  ( $i = 1, 2$ ), as shown in Fig. 1(b). At the inlet  $\Gamma_{in}$ , the inflow velocity along  $x$ -axis is prescribed by a parabolic velocity profile in the form,

$$U_{inlet}(y) = U_m \frac{(H - 2y)(H + 2y)}{H^2}, \quad (3)$$

and that along  $y$ -axis is prescribed as,

$$V_{inlet}(y) = 0. \quad (4)$$

$U_m$  is the maximum velocity magnitude of the parabolic profile, and  $H = 4.1D$  represents the total height of the rectangular domain. The average inlet velocity  $\bar{U}$ , is related to the parabolic velocity profile  $U_{inlet}(y)$  through the expression:

$$\bar{U} = \frac{1}{H} \int_{-H/2}^{H/2} U_{inlet}(y) dy = \frac{2}{3} U_m. \quad (5)$$

At the outlet  $\Gamma_{out}$ , an outflow boundary condition is applied where the velocity is to be extrapolated. No-slip condition is applied at both the upper and lower walls  $\Gamma_w$ . On the surface of the cylinder, we implement AFC by introducing synthetic jets at the highest and lowest points of the cylinder, as shown in Fig. 1(c). These jets  $\Gamma_i$  ( $i = 1, 2$ ) can be used to inject or suction fluid for flow control. Each jet has a width of  $\omega = 10^\circ$  and is directly perpendicular to the outer surface of the cylindrical wall. The velocity distribution within each jet follows a parabolic profile. The jet velocities can be positive or negative, corresponding to blowing or suction, respectively. It is crucial to maintain a balance in the net mass flow rates of all the jets to ensure that the net flow into or out of the system is zero, i.e.,  $V_{\Gamma_1} = -V_{\Gamma_2}$ . This constraint guarantees overall conservation of mass within the system. No-slip boundary condition is applied on the surface of the cylinder where the jets are not located.

The aforementioned problem is solved using the open-source CFD software OpenFOAM, written in C/C++<sup>46</sup>. Specifically, the built-in solver, PimpleFoam which combines the Pressure-Implicit with Splitting of Operators (PISO) algorithm and the Semi-Implicit Method for Pressure-Linked Equations (SIMPLE) algorithm, is employed to solve this incompressible fluid flow

problem<sup>47–49</sup>. The computational domain is discretized into 16,200 grid elements using unstructured triangular mesh units. The mean velocity magnitude is selected as  $\bar{U} = 1$ , and the Reynolds number is set to  $Re = 100$ . To ensure numerical stability, the time step is chosen as  $\Delta t = 0.0005$ . The simulation results for this numerical setup were verified against the classical benchmark in a previous study by Wang *et al.*<sup>50</sup>. In this work, we are interested in comparing the lift coefficient ( $C_L$ ) and drag coefficient ( $C_D$ ), which are defined as

$$C_L = \frac{F_L}{0.5\rho\bar{U}^2D}, \quad C_D = \frac{F_D}{0.5\rho\bar{U}^2D}. \quad (6)$$

Here,  $F_L$  and  $F_D$  represent the lift and drag forces integrated on the surface of the cylinder, respectively, and  $\rho$  is the fluid density.

## B. Overview of DRL

This subsection provides a brief overview of DRL. However, this is by no means an exhaustive summary and readers are referred to Arulkumaran *et al.*, Ladosz *et al.*, Shakya, Pillai, and Chakrabarty<sup>51–53</sup> for a more thorough discussion. The reinforcement learning problem is typically modeled as a Markov decision process (MDP), which is a mathematical framework that formalizes sequential decision-making problems involving an agent interacting with an environment<sup>54</sup>. Fig. 2 illustrates the typical interaction between an agent and an environment. At any time-step  $t$ , the environment is at some state  $s_t$ . Based on this observing state, the agent's policy  $\pi_\theta(a_t|s_t)$  determines the agent's action  $a_t$ , which in turn causes the environment to transit into a new state  $s_{t+1}$ <sup>54,55</sup>. At the same time, a reward  $r_t$  is awarded to the agent. The ultimate goal is to find the optimal policy that maximizes the agent's cumulative reward over the long run. By actively collecting and learning from this data, the agent can enhance its performance over time by improving its policy optimization and decision-making<sup>24,56,57</sup>.




FIG. 2: The agent and the environment are fundamental components of reinforcement learning. The environment is the entity with which the agent interacts. At any given time step  $t$ , the agent first observes the current state  $s_t$  of the environment, along with the corresponding reward value  $r_t$ . Based on these state and reward information, the agent decides how to take action  $a_t$ . The agent receives feedback from the environment, obtaining the next time step's state  $s_{t+1}$  and reward  $r_{t+1}$ .

The Proximal Policy Optimization (PPO) algorithm<sup>58–61</sup>, which is a specific implementation based on the Policy Gradient approach for optimizing an agent's policy function in reinforcement learning, is used for training in this study. The PPO algorithm enhances training stability by introducing a *proximal optimization* mechanism<sup>24,26,62</sup> and effectively addresses challenges in continuous action spaces and high-dimensional state spaces, leveraging the power of deep learning for efficient processing and enabling deep decision-making models to excel in complex tasks<sup>24,26</sup>.

For a given trajectory  $\tau = \{(s_0, a_0), (s_1, a_1, r_1), \dots, (s_n, a_n, r_n)\}$  that describes a sequence of states from the initial state, through various states and actions, the end state, the expected value of the discounted return  $G(\tau)$  is obtained by maximizing the following objective function:<sup>63</sup>

$$J(\theta) = \mathbb{E}_{\tau \sim \pi_\theta} [G(\tau)], \quad (7)$$

which quantifies the expected cumulative reward that an agent can achieve by following the policy  $\pi_\theta$ . Here  $\mathbb{E}$  denotes the expectation over the trajectory  $\tau$  sampled from the policy  $\pi_\theta$ , and  $G(\tau)$  is the discounted return function at a specific time step  $t$ , defined as:

$$G(t) = \sum_{k=0}^T \gamma^k R_{t+k+1}, \quad (8)$$

where  $\gamma \in [0, 1]$  is the discount factor that determines the relative importance of future rewards compared to immediate rewards, and  $T$  is the total length of the game. The optimal parameter  $\theta^*$  is obtained as:<sup>63</sup>

$$\theta^* = \operatorname{argmax}_{\theta} \mathbb{E}_{\tau \sim \pi_\theta} [G(\tau)]. \quad (9)$$

This equation represents the search for  $\theta$  that maximizes the expected value of the discounted return by sampling trajectories  $\tau$  from the policy  $\pi_\theta$ . For the purpose of iterative updates, the objective function of the PPO algorithm is typically defined as a surrogate objective that balances exploration and exploitation,<sup>58</sup> as follows:

$$L^{CLIP}(\theta) = \hat{\mathbb{E}}_t [\min(r_t(\theta)) \hat{A}_t, \text{clip}(r_t(\theta), 1 - \varepsilon, 1 + \varepsilon) \hat{A}_t)], \quad (10)$$

where  $L^{CLIP}(\theta)$  represents the objective function that is to be maximized or minimized.  $\theta$  represents the parameters of the policy being optimized.  $\hat{\mathbb{E}}_t$  represents the expectation operator, computing the average value of the clipped surrogate objective term over sampled trajectories with respect to the policy's probability distribution.  $r_t(\theta)$  is the ratio between the probabilities of selected actions under the new and old policies.  $\varepsilon$  is a hyperparameter.  $\hat{A}_t$  denotes the advantage estimate, which measures the disparity between the expected value and the value of a specific action in a given state. The  $\text{clip}(\cdot)$  function controls the scale of updates during training by constraining the ratio between the probabilities of the new and old policies within a specified range  $[1 - \varepsilon, 1 + \varepsilon]$ . Note that in this equation, the first term within the min function is employed in calculating the objective function for policy updates, and the second term in the equation,  $\text{clip}(r_t(\theta), 1 - \varepsilon, 1 + \varepsilon) \hat{A}_t$ , modifies the surrogate objective by constraining the probability ratio<sup>58</sup>.

### C. Formulation of the AFC problem using DRL

The relationship between the DRL component and the CFD component is outlined in this subsection. We follow the coupling strategy proposed by Wang *et al.*<sup>50</sup>. For a more comprehensive understanding, we encourage readers to consult their work and the references therein.

The overarching goal of the AFC problem is to achieve drag reduction on the cylinder by suppressing vortex shedding in the von Kármán vortex street. To address this problem using DRL, the CFD numerical simulation serves as the environment with which the agent interacts. Specifically, the agent interacts with the environment through three components: state, action and reward.




FIG. 3: Instantaneous velocity field with black dots representing the positions of the 149 probes.

- At each moment in time, the state consists of the instantaneous flow field data gathered from specific positions called probes within the computational domain. A comprehensive collection of 149 probes is strategically placed around the cylinder and in the wake region, and their positions are shown in Fig. 3. These probes are carefully positioned to capture crucial flow characteristics while maintaining a sparse distribution. For validation purposes, the probe layout in this work is adopted from that proposed by Wang *et al.*<sup>50</sup>. Other designs of the probe layout can also be found in literature<sup>34,64–66</sup>.
- Based on the observed state, the agent's policy determines the action. The agent's policy is parameterized by a two-layer artificial neural network with 512 neurons in each layer, as advised by Rabault *et al.*<sup>34</sup>. The resulting action value, denoted as  $a$ , is responsible for adjusting the velocities of the synthetic jets. This enables precise control over the flow dynamics. To

ensure stability and avoid non-physical instabilities in the flow, a smoothing function is applied, which holds the following relationship:

$$V_{\Gamma_1, T_i} = V_{\Gamma_1, T_i} + \beta (a - V_{\Gamma_1, T_{i-1}}), \quad (11)$$

where  $V_{\Gamma_1, T_i}$  represents the jet velocity at the end of the  $i$ th actuation control period  $T_i$ , and  $\beta$  is a numerical parameter used for smoothing. The value of  $\beta$  is determined through trial and error, and in this case, it is set to  $\beta = 0.4$ . Each control period  $T_i$  consists of 50 time steps, denoted as  $T_i = (50i)\Delta t$ . To prevent excessive energy input in the jets, an upper limit constraint is imposed, ensuring that  $V_{\Gamma_1} \leq U_m$ .

- At the end of each control period, a reward function is employed to evaluate the agent's performance. With the primary aim of achieving maximum drag reduction, the reward function is defined as

$$r_{T_i} = C_{D,0} - (C_D)_{T_i} - \omega |(C_L)_{T_i}|, \quad (12)$$

where  $C_D$  and  $C_L$  are the drag and lift coefficients, respectively. The term  $C_{D,0}$  corresponds to the mean drag coefficient of the circular cylinder without any flow control, specifically set to  $C_{D,0} = 3.205$ . The notation  $(\cdot)_{T_i}$  indicates an averaging operation over the actuation period  $T_i$ . The parameter  $\omega = 0.1$  serves as a weighting factor that emphasizes the contributions of lift fluctuations in the reward calculation. The inclusion of the lift term is intended to prevent unnecessary increases in lift fluctuations during the optimization process.

In the training process, each episode consists of a complete CFD simulation that spans 5000 time steps. This duration corresponds to 100 actuation periods. Therefore, the maximum total time for each episode is  $T_{\max} = T_{100} = 2.5$  non-dimensional time units. The length of the episode is intentionally selected to encompass multiple vortex shedding cycles. This allows the agent to effectively learn the control algorithm by observing and adapting to the dynamics of the flow over several cycles.

#### D. Parallelization of DRL-based AFC problem

In a previous study, it was demonstrated that approximately 99% of the computational time is dedicated to CFD simulations. With increasingly complex simulations, there is an intrinsic demand for utilizing parallelization in DRL training to reduce the overall runtime<sup>35</sup>.

In a general DRL-based framework for AFC, there are two available parallelization strategies: parallelizing the CFD simulation and parallelizing the DRL training process through multiple environments. The CFD component can allow parallelization through parallel programming softwares, such as Message Passing Interface (MPI) where one CFD instance is distributed across several processes. On the other hand, the data collection process in DRL can be parallelized by creating multiple environments simultaneously. Each environment represents a separate instance of the AFC problem, allowing the DRL agent to interact and learn from multiple environments concurrently.

The hybrid parallelization approach used for DRL parallel training is illustrated below. Each environment consists of multiple concurrent processes (or MPI ranks) to launch CFD simulations. Typically, the number of processes or MPI ranks ( $N_{\text{ranks}}$ ) is equal to the number of CPUs ( $N_{\text{CPUs}}$ ) dedicated to a single CFD simulation. This setup accelerates the simulation time for each actuation period, enabling faster transitions between states according to the agent's actions and reducing the time required to collect  $(s, a, r)$ -tuples along each trajectory. Furthermore, training is conducted simultaneously in multiple environments, which expedites data collection and processing, thus speeding up the overall training process. At any given time, multiple environments with multiple CFD instances can coexist, each having different flow states. Each environment operates independently and in parallel, with states influenced solely by the history of actions within that environment. This parallelization allows for concurrent state transitions and reward calculations, providing a larger number of samples and experiences for the DRL algorithm within a given timeframe. Once all environments complete one training episode, data from multiple trajectories are batched together in mini-batches to update weights and parameters in the DRL algorithm. In a hybrid setup with  $N_{\text{envs}}$  environments, the total number of CPUs for the overall framework is the product of the number of ranks per CFD simulation ( $N_{\text{ranks}}$ ) and the number of environments ( $N_{\text{envs}}$ ), i.e.,  $N_{\text{total CPUs}} = N_{\text{envs}} \times N_{\text{ranks}}$ . Figure 4 presents an example of hybrid parallelization using a total of 20 CPUs, with 4 environments ( $N_{\text{envs}} = 4$ ) and 5 MPI ranks ( $N_{\text{ranks}} = 5$ ). The framework allows flexibility in choosing the number of CPUs and environments for either form of parallelization, without any specific restrictions.

#### E. Software platform

We build our work upon the codebase `DRLinFluids` developed by Wang *et al.*<sup>50</sup>. As mentioned in Section II A, the CFD component is implemented using the open-source software `OpenFOAM`<sup>46</sup>. The DRL component is implemented using another open-source Python framework `TensorForce`<sup>67</sup>. Both components allow for parallelization as discussed in Section II D.

## DRL in Parallel




FIG. 4: Illustration of the process and allocation of computational resources for DRL parallel training. Here, four training environments are utilized, each employing 5 MPI ranks for parallel CFD computations. This configuration requires a total of 20 cores. During the training, the agent continuously interacts with the CFD environments, generating a series of tuples  $(s_i^m, a_i^m, r_i^m)$ . Here, the superscript  $m$  denotes the specific environment, and the subscript  $i$  represents the timestep when the agent interacts with that environment. The trajectory  $\tau^m$  resulting from the agent-environment interaction is then used to calculate the gradient  $\nabla_{\theta} J(\tau^m)$  and update the neural network parameters  $\theta$ . DRL parallelizes the training across the four environments, meaning that in each training step, the four environment instances operate simultaneously. This allows for the concurrent collection of multiple samples. Importantly, each environment instance operates independently, creating an isolated setting for training.

To establish the communication between the CFD simulation software and the DRL package, Wang *et al.* proposed the utilization of file I/O streams as the interface for these two components within the framework of DRLinFluids<sup>50,68</sup>. At the end of each actuation period, the observed states, histories of aerodynamics coefficients (which are necessary for reward computation), and resulting flow fields are saved from memory to the hard disk. Subsequently, the state and history files are parsed into Python scripts to be further processed by the DRL package.

The action values generated by the DRL algorithm are subsequently extracted and integrated back into the OpenFoam configuration files using Regex (Regular expression) operations<sup>69</sup>. These updated configuration files, along with the most recent flow field data, are used to initiate the subsequent instance of OpenFoam for simulating the next actuation period. Therefore, each episode of training with 100 actuation periods consist of 100 instances of OpenFoam processes interleaved with I/O to/from the Python package TensorFlow<sup>67,70</sup>. This iterative process ensures continuous interaction and synchronization between the DRL algorithm and the CFD simulation environment. Wang *et al.*<sup>50</sup> argued that the adoption of an I/O-based interface is suitable, since the primary performance bottleneck lies in the computation in numerical simulations rather than the data movements in I/O. In Section III C, we will evaluate the validity of this assumption rigorously, highlight its shortcomings, and propose improvements to further optimize these interfaces.

## III. RESULTS

The results' analysis is as follows: In Section III A, the control results achieved by Wang *et al.*<sup>50</sup> are initially replicated, with a particular emphasis on training efficiency and performance. Subsequently, the entire framework is deconstructed, and the parallel efficiency of each individual component is meticulously analyzed. In Section III B, the MPI-based parallelization of the



CFD solver is delved into, while Section III C focuses on multi-environment parallel training within the DRL framework. To further optimize the framework and overcome bottlenecks in I/O operations, refinements are explored in Section III D.

Throughout the analysis, computational hardware equipped with an Intel(R) Xeon(R) Platinum 8358 CPU @ 2.60GHz, featuring a total of 64 cores (32 cores per socket across 2 sockets), is employed. The software suite utilized includes TensorForce version 0.6.0 as the reinforcement learning framework and OpenFOAM version 8 as the CFD platform.

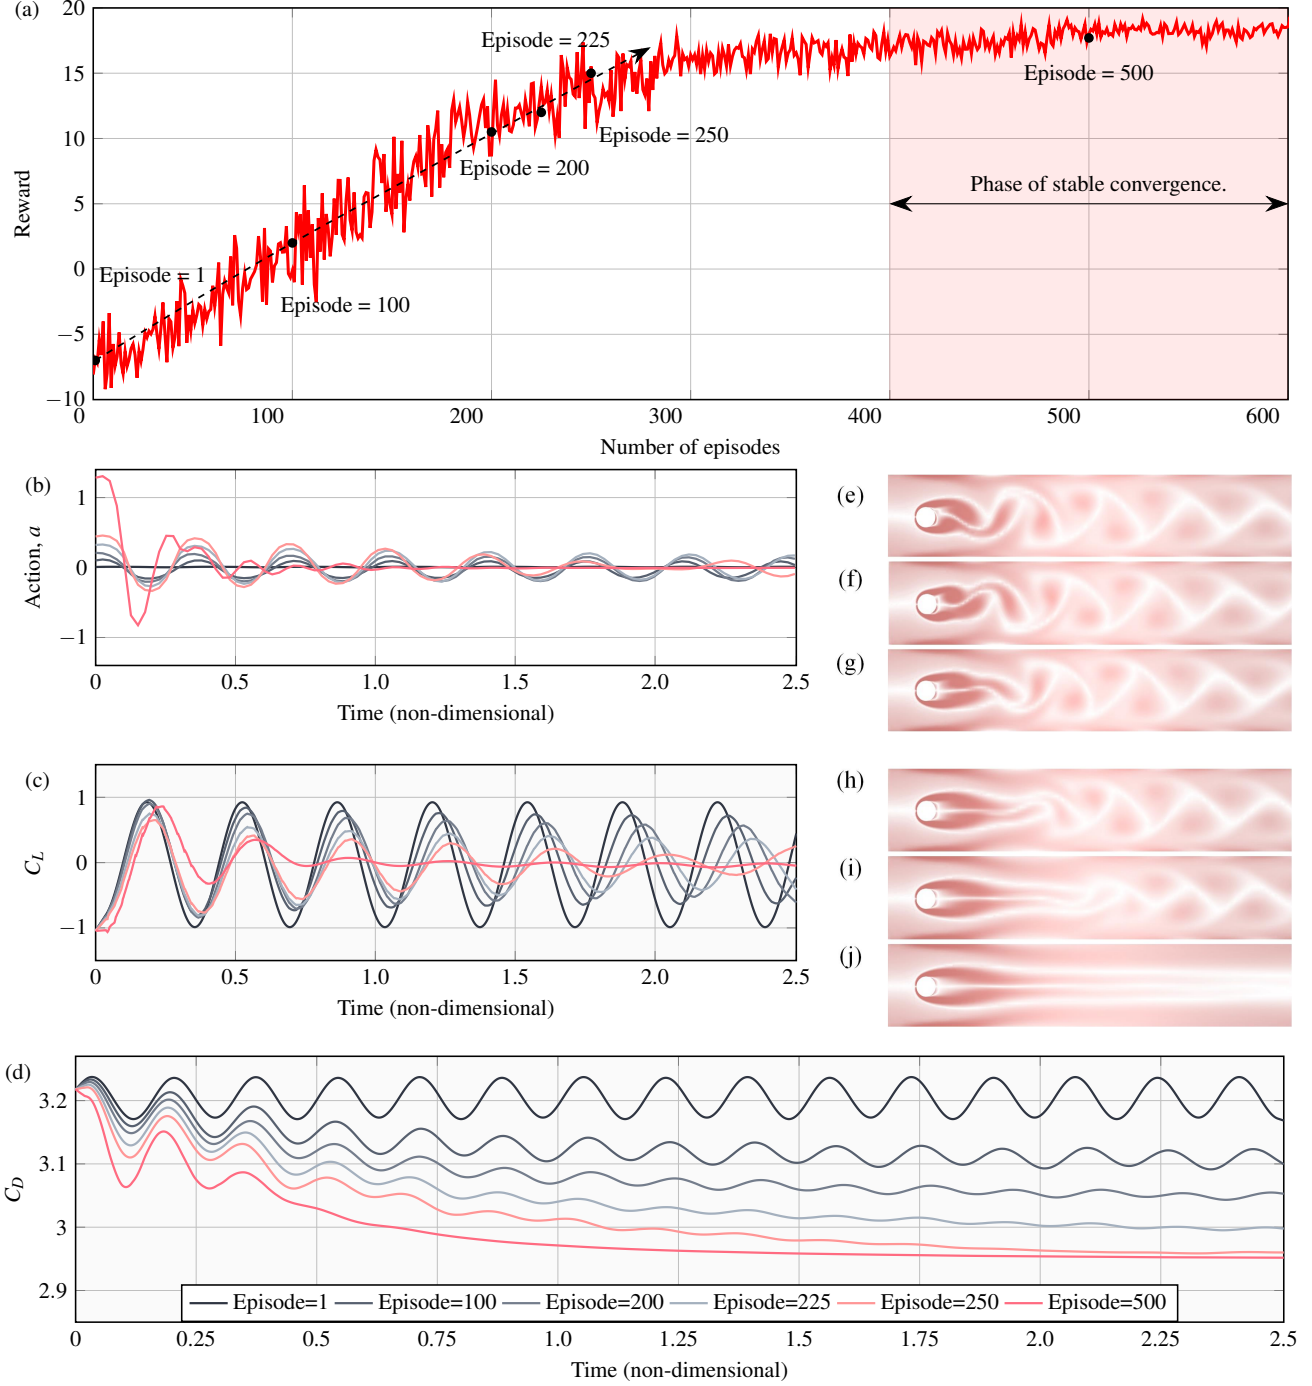


FIG. 5: Results of DRL training. (a) displays the cumulative reward at each episode during the training process. (b), (c), and (d) show the changes in action  $a$ , lift coefficient  $C_L$  and drag coefficient  $C_D$ , respectively. Histories of each parameter at selected episodes are shown to illustrate the convergence. (e) to (j) correspond to vorticity contours at the end of each selected episode.



### A. Validation study

The DRLinFluids framework proposed by Wang *et al.*<sup>50</sup> is validated here. We apply the DRL algorithm to learn AFC strategies, with a specific focus on the application of synthetic jet excitation in the presence of flow around a circular cylinder. The results are summarized in Fig. 5. Our findings demonstrate the effectiveness of the DRL algorithm in reducing flow resistance around a circular cylinder through AFC. These results serve as evidence of the capability of the DRLinFluids framework in achieving the desired control objectives.

During training, the PPO algorithm iteratively optimizes the parametrized neural network-based policy to maximize cumulative rewards at each episode. Fig. 5(a) displays the learning curve, which depicts the development trend of the return for each round. Starting from an initial reward value, the agent's return gradually increases and tends to converge around 300 episodes. This tortuous rise in the learning curve is a direct consequence of the reinforcement learning process, where the agent continuously experiments, makes mistakes, accumulates experience, and gradually strengthens the reward function.

To illustrate the learning process, we focus on selected episodes and analyze the action value ( $a$ ), lift coefficient ( $C_L$ ), and drag coefficient ( $C_D$ ) as shown in Fig. 5(b) to Fig. 5(d). In the first episode, the agent starts with small action values, leading to minimal changes in the flow. Consequently, the lift and drag coefficients remain similar to those in the uncontrolled flow. As the number of episodes increases, the agent learns to explore the action space to optimize the objective. By the 100th episode, there are indications that the agent is learning to introduce flow changes periodically, aligning with the shedding frequency of vortices. This control strategy successfully reduces drag and yields increased reward. By the 200th episode, the agent further improves the reward through better drag reduction policies. Notably, at the 225th and 250th episodes, significant optimization of drag reduction is observed within a few training episodes. This suggests that the agent capitalizes on learned knowledge from previous episodes and focuses more on exploitation than exploration to find the optimal policy. After 400 episodes, the reward converges to a steady value. At the 500th episode, the drag and lift coefficients ( $C_D$  and  $C_L$ ) demonstrate the successful suppression of the vortex shedding phenomenon. The vorticity contours in Fig. 5(b) to Fig. 5(d) further illustrate the complete suppression or control of vortex shedding behind the circular cylinder. Overall, the implementation of DRL control results in a significant decrease in the drag coefficient ( $C_D$ ), reducing it from 3.205 to approximately 2.95. This corresponds to a remarkable drag reduction rate of 8%. Additionally, our proposed control strategy effectively mitigates the periodic fluctuation of the lift coefficient ( $C_L$ ).

In the context of run-time performance, our focus is on the DRL-based AFC framework. In this validation study, we utilized a single reinforcement learning environment, with CFD simulations running on a single core. The training process consisted of 600 episodes, resulting in a total training time of 45.1 hours, with each episode taking approximately 4.5 minutes. Comparatively, Wang *et al.*<sup>50</sup> also employed a single environment training for the same AFC problem. They reported an average time of 10 minutes per episode on a 22-core CPU, resulting in a total training time of 100 hours for 600 episodes. While the total time consumption may vary due to differences in computer hardware, it is evident that both studies require a significant amount of time to achieve stable convergence results. Based on the previous work Rabault *et al.*<sup>34</sup> and our experiences (as mentioned in Section III D), it is observed that the computational time of CFD simulations constitutes the majority (over 95%) of the total time. For complex flow phenomena, the computational demands of CFD simulations can become prohibitively expensive for single-core computation. Therefore, leveraging parallel computing architectures is crucial to optimize the run-time performance and make DRL feasible for complex AFC problems.

In a previous study<sup>35</sup>, parallelization through multi-environment training has demonstrated potential in expediting the training duration. Given that the convergence of rewards appears to be largely unaffected by the quantity of environments, initiating numerous training environments could feasibly diminish the time required to accumulate learning samples essential for the agent's training. In Fig. 6, we verify this claim using our framework and observe that the convergence rate of rewards is indeed consistent across various scenarios. Drag and lift coefficients are controlled successfully in all cases. Under the assumption that training duration for each episode remains the same, this could imply a reduction in total training time proportional to the increase in the number of environments. However, it is crucial to acknowledge that the duration of each training episode may not remain consistent in a parallel setting due to communication overhead and additional complexities inherent in multi-environment training. Furthermore, the prospect of parallelizing each training episode through the application of multi-core CFD simulations has not been thoroughly explored in existing literature.

Therefore, in a parallel setting, the question of how best to balance the parallelization of CFD simulations with the use of multiple DRL environments for optimal performance is still open to exploration. To the best of our knowledge, no existing work provides a comprehensive analysis for the scaling performance of DRL frameworks when applied to AFC problems. In the subsequent subsections, we shall analyze various parallelization strategies, computing resource allocation in different scenarios, and other bottlenecks that hinder performance. Our objective is to identify strategies that effectively utilize computing resources, minimize training time, and provide valuable insights for optimizing the design of parallel training and computation in AFC applications.




FIG. 6: (a) shows the reward function curve when the number of parallel environments  $N_{\text{envs}} = 1, 4, 8, 10, 20$ . (b) shows the final results of training when different numbers of environments are run in parallel. Drag coefficient  $C_D$ , lift coefficient  $C_L$ , and action values are consistent with the single-environment training in Fig. 5.

## B. Scaling analysis of CFD simulations

In this subsection, we shift our attention to the scaling efficacy of CFD computations with the objective of identifying the most suitable core allocation for parallelized simulations. In an effort to isolate the CFD component from potential DRL-related disturbances, this set of numerical experiments are conducted where only a single environment is used for training, i.e.,  $N_{\text{total CPUs}} = N_{\text{ranks}}$ . It should be noted that the parallel efficiency of OpenFoam,<sup>71–75</sup> along with CFD algorithms at large,<sup>76–78</sup> has been extensively studied in literature. Despite this, to date, there lacks an investigation specifically targeted at CFD parallelization within the framework of DRL applications. For instance, given that DRL can be parallelized through multi-environment training, no clear guidelines exist on how to balance the hybrid parallelization between CFD computations and DRL training. Moreover, integrating CFD solvers within a DRL framework introduces distinct characteristics, such as frequent interactions with the DRL module, which could theoretically impact the scaling performance.

Taking all these factors into consideration, we benchmark the parallel efficiency of CFD computations to establish a foundation for discussing the overall scalability of the framework. To achieve this, we employ a multi-core computing approach and make adjustments to the core count or number of MPI ranks,  $N_{\text{ranks}}$ , for CFD computations. This allows us to evaluate the performance of the CFD solver under different parallelization configurations and provide insights into the scalability of the framework. We systematically record the computation time for each  $N_{\text{rank}}$  and calculate the relationship between speedup and the number of ranks. This allows us to gain insights into the effectiveness of parallel computing of the CFD component and determine the number of  $N_{\text{rank}}$  ranks that yields optimal performance. For investigating the effect of interaction with DRL, we present two sets of scaling results. The first set with  $T_1 = 0.025$  denotes scaling performance of a single instance of OpenFoam. This evaluates the parallelism for the CFD solver by itself. The other set with  $T_{\text{max}} = T_{100} = 2.5$  denotes scaling performance where we consider the entire training episode with 100 instances of OpenFoam and intermediate data movement with the DRL component. This evaluates the parallel performance of the CFD solver in a realistic setting when coupled with DRL frameworks.

In Fig. 7, we present the influence of the number of MPI ranks ( $N_{\text{ranks}}$ ) on the parallel speedup ratio and efficiency for CFD computations. In an ideal scenario, the speedup ratio would exhibit a linear relationship with  $N_{\text{ranks}}$ , where the speedup ratio equals the number of cores employed. However, as shown in Fig. 7(a), the current CFD parallel computations demonstrate a lower speed ratio compared to the ideal value. Even when two  $N_{\text{ranks}}$  are utilized in parallel, the speedup drops significantly and fails to reach 2. Further analysis reveals that as  $N_{\text{ranks}}$  increases, the incremental gain diminishes and eventually saturates. In addition, Fig. 7(b) demonstrates that the parallel efficiency decreases rapidly as the number of MPI ranks increases. With a 2-core parallel scheme, the utilization of computational resources can only reach 90%. As  $N_{\text{ranks}}$  further increases, the parallel efficiency drops below 20% when  $N_{\text{ranks}} = 16$ . This drastic decrease in parallel efficiency indicates that allocating resources to parallelize individual OpenFOAM instances has limited performance benefits. It is worth noting that both sets of benchmarks show similar performance, suggesting that the interaction with the DRL component does not introduce significant overhead under single-environment training. The discussion regarding multi-environment training will be addressed in Section III D.

We acknowledge that it is possible to optimize the parallel performance of the CFD solver through algorithmic and software development. However, such optimization is highly dependent on various factors, including the numerical scheme and mesh

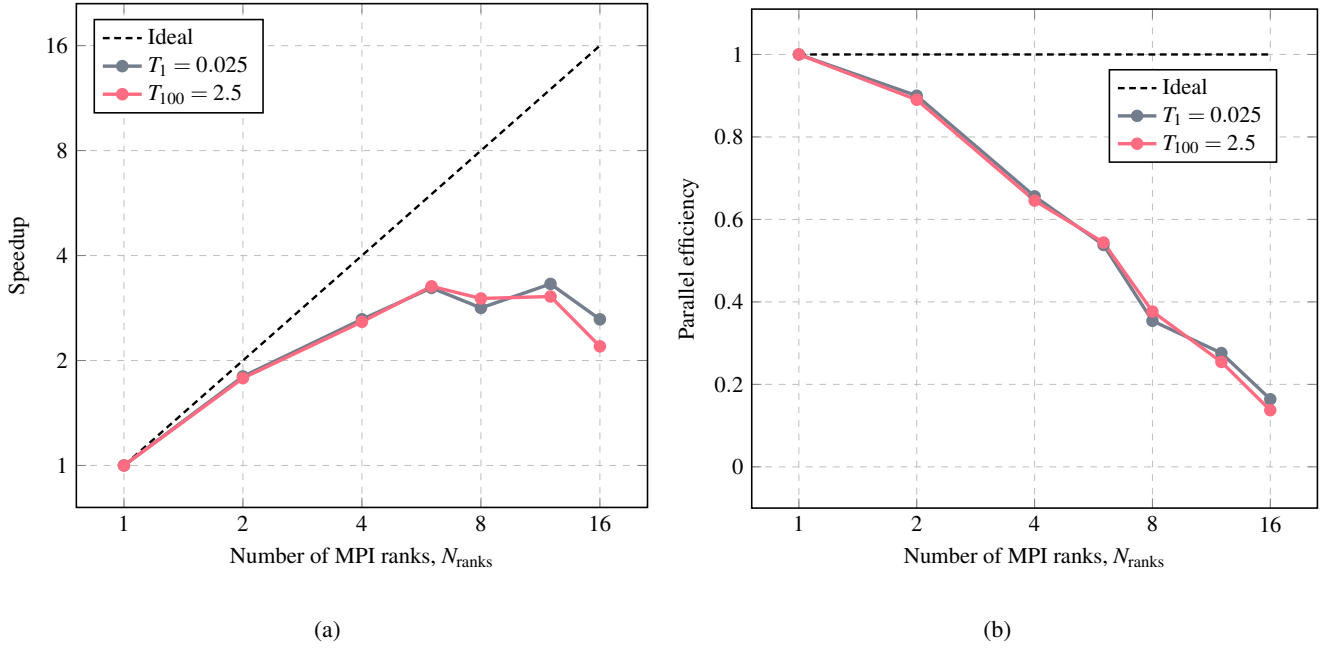


FIG. 7: Scaling performance of the CFD solver: (a) speedup ratio, (b) parallel efficiency of the CFD instances.  $N_{\text{ranks}} = 1$  in each set of data is taken as the respective reference point.  $T_1$  denotes one instance of the CFD solver.  $T_{100}$  denotes 100 instances of the CFD solver interleaved with data transfer with the DRL framework, which is equivalent to one episode of single-environment training.

topology, and it remains an active research field. Therefore, the performance optimization of the CFD solver for the application of DRL to generic AFC problems is considered beyond the scope of the current work. The results presented above are characteristic and provide a sufficient foundation for the upcoming discussions.

### C. Scaling analysis of the DRL framework with multi-environments

Following the analysis of the CFD component in the previous section, this segment focuses on the parallel efficiency of the DRL component. Our objective is to comprehensively evaluate the impact of varying the number of environments on the performance of parallel training and explore the relationship between the number of environments in DRL training and the speedup ratio. To achieve this, we employ three different scenarios for the parallelization of CFD, utilizing one, two, and five MPI ranks, respectively. This enables us to investigate the influence of different computing resources on the parallel efficiency of the DRL framework.

#### 1. Multi-environment DRL parallelization

In the first set of tests, we allocate five MPI ranks for each environment to conduct the flow simulations using CFD. This configuration is to analyze the parallel efficiency of DRL when CFD parallelization is not optimal. All tests are executed on a 64-core CPU, where we concurrently utilize up to 12 environments to parallelize the DRL component. (Four cores are not used in this case.) For each test, encompassing a training volume of  $N_{\text{episodes}} = 3000$  episodes, we document the duration required for the training process and evaluate its parallel efficiency. Results of these tests are consolidated and presented in Table. I. As observed in Table. I, increasing the number of environments from 1 to 12 while keeping the number of MPI ranks constant at 5 per environment results in a decrease in training duration from 305.8 hours to 32.4 hours. This trend is indicative of successful parallel scaling up to a certain point. However, the parallel efficiency exhibits a slight decrement from 100% down to 78.6% as the number of environments increases.

It is noteworthy that the parallel efficiency remains around 80% even as the number of environments increases to 12, which suggests that the DRL component scales relatively well with the addition of computational resources. The notable efficiency in the multi-environment setup is largely due to the independent execution of CFD simulations across different environments, which run concurrently without data dependency. Interactions between the CFD simulations and the DRL framework within

TABLE I: Statistics of parallel multi-environment training with  $N_{\text{ranks}} = 1, 2, 5$ .

$N_{\text{episodes}}$	$N_{\text{envs}}$	$N_{\text{ranks}}$	$N_{\text{total CPUs}}$	Total duration (h)	Speedup	Parallel efficiency (%)
<b>CFD MPI Ranks, <math>N_{\text{ranks}} = 5</math></b>						
3000	1	5	5	305.8	1.0	100.0
3000	2	5	10	170.8	1.8	89.5
3000	4	5	20	88.5	3.5	86.4
3000	6	5	30	59.7	5.1	85.3
3000	8	5	40	47.3	6.5	80.8
3000	10	5	50	38.3	8.0	79.8
3000	12	5	60	32.4	9.4	78.6
<b>CFD MPI Ranks, <math>N_{\text{ranks}} = 2</math></b>						
3000	1	2	2	289.6	1.0	100.0
3000	2	2	4	156.3	1.9	92.6
3000	4	2	8	80.0	3.6	90.4
3000	6	2	12	53.4	5.4	90.3
3000	8	2	16	40.8	7.1	88.7
3000	10	2	20	33.2	8.7	87.1
3000	20	2	40	17.7	16.4	81.9
3000	30	2	60	12.4	23.3	77.5
<b>CFD MPI Ranks, <math>N_{\text{ranks}} = 1</math></b>						
3000	1	1	1	225.2	1.0	100.0
3000	2	1	2	123.7	1.8	91.1
3000	4	1	4	64.6	3.5	87.1
3000	6	1	6	44.4	5.1	84.5
3000	8	1	8	33.9	6.6	83.0
3000	10	1	10	26.3	8.6	85.5
3000	20	1	20	14.2	15.9	79.3
3000	30	1	30	9.6	23.5	78.4
3000	40	1	40	9.0	25.0	62.4
3000	50	1	50	8.1	27.7	55.4
3000	60	1	60	7.6	29.6	49.3

each trajectory incur minimal communication costs since extracting policy actions from the DRL framework is computationally inexpensive compared to CFD simulations. Post-trajectory updates in the DRL framework, while slightly more communication-intensive, are still less demanding than the CFD simulations themselves. Additional factors such as the management of multi-threading processes and file I/O operations may contribute marginally to the overall computational overhead. Nevertheless, The dominant CFD simulations are almost embarrassingly parallel with respect to the number of environments, leading to remarkable scaling performance.

We emphasize that the scaling analyses detailed above utilized a multi-core configuration for CFD with  $N_{\text{ranks}} = 5$ . As outlined in Section III B, such a configuration does not represent an optimal allocation of computational resources, attributed to the limited parallel efficiency of the CFD software. Nonetheless, we observe good scaling performance for the DRL component as we increase the number of environments. In Table. I, we present two further sets of experimental outcomes wherein  $N_{\text{ranks}}$  are configured to 1 and 2, respectively. The speedup ratio for all tests are visualized in Fig. 8. As explained in Section III B, simulations conducted on a single-core with  $N_{\text{ranks}} = 1$  are the most efficient, whereas dual-core simulations with  $N_{\text{ranks}} = 2$  are sub-optimal. Across all configurations, we observe a consistently robust scaling performance. The scaling performance for  $N_{\text{ranks}} = 1$  drops slightly for very large-scale runs. We investigate this performance regression in Section III D. Nevertheless, the parallel efficiency is significantly better than those presented in Section III B. The implications of these findings are twofold: firstly, the parallel efficiency of multi-environment training appears to be invariant to the parallelization of the underlying CFD simulations; secondly, prioritizing the parallelization of the DRL component over that of CFD simulations is a more resource-efficient strategy.

## 2. Resource allocation for hybrid parallelization

Here, we evaluate resource allocation strategies with the intent to ascertain the most efficient hybrid parallel configuration. The objective is to provide a road-map for leveraging hybrid parallelism in DRL-infused CFD simulations. This requires a reassessment of the configurations involving one, two, and five MPI ranks for CFD parallelization, in conjunction with a variable

number of environments, to analyze their scaling performance against the total number of cores employed. A standardized reference point is essential for a consistent comparison on resource utilization efficiency. For this purpose, we utilize the test case with a single-core CFD and single environment ( $N_{\text{rank}} = 1, N_{\text{envs}} = 1$ ) as the reference baseline for calculating the parallel speedup and efficiency across the entire framework.




FIG. 8: Speedup ratio of the multi-environment DRL training. Three different configurations for the CFD solver are used:  $N_{\text{ranks}} = 1, 2, 5$ .  $N_{\text{envs}} = 1$  in each set of data is taken as the respective reference point.

The corresponding results are presented in Fig. 9. Specifically, in the speedup plot in Fig. 9(a), the speedup ratios increase as the number of CPU cores increases. The gradients of each curve illustrate the efficiency of parallelizing DRL through multiple environments and agree with those in Fig. 8. The consistency in the gradient provides proof that multi-environment parallelization is largely unaffected by the underlying CFD parallelization configuration due to the embarrassingly parallel nature of these environments. The offsets among different curves indicate the efficiency of CFD parallelization. With the same total number of CPUs, using multiple ranks for CFD lowers the speedup ratio as the CFD component with OpenFoam has poor parallel performance. For a more direct comparison, Fig. 9(b) details the total parallel efficiency for all cases. When CFD utilizes five-rank MPI parallel computation, the parallel efficiency remains consistently below 20%. When CFD adopts two-rank MPI parallel computation, the parallel efficiency decreases from 50% to around 38%. However, when CFD relies on a single-core computation, the parallel efficiency declines from 100% to approximately 49%. Despite these declining trends, single-core configuration still exhibits superior parallel efficiency compared to other parallelization schemes.

The empirical evidence suggests that the most optimal parallelization strategy involves DRL multi-environment parallel training coupled with CFD single-core computation. Our experiments utilizing 60 cores demonstrate that the duration for training can be reduced from 225.2 hours to a mere 7.6 hours. This equates to a speedup of approximately 30-fold, substantiating the effectiveness of this hybrid parallelization approach.

#### D. Further optimization through refining I/O operations

In the preceding subsection, we notice that when CFD employs single-core computation, utilizing multi-environment parallel training in DRL achieves the optimal parallel acceleration. However, as the number of parallel environments in DRL is substantially increased ( $N_{\text{total CPUs}} > 30$ ), there is a drastic decrease in parallel efficiency. Such a decrease is not evident when employing  $N_{\text{ranks}} = 2$  or 5. This pattern implies that the regression in performance may be linked to the expansive number of environments. To further explore this issue, we conducted a profiling of the computational time required for a single training episode, with the results shown in Fig. 10. The data indicates that the CFD simulation time predominates the training duration. However, there is

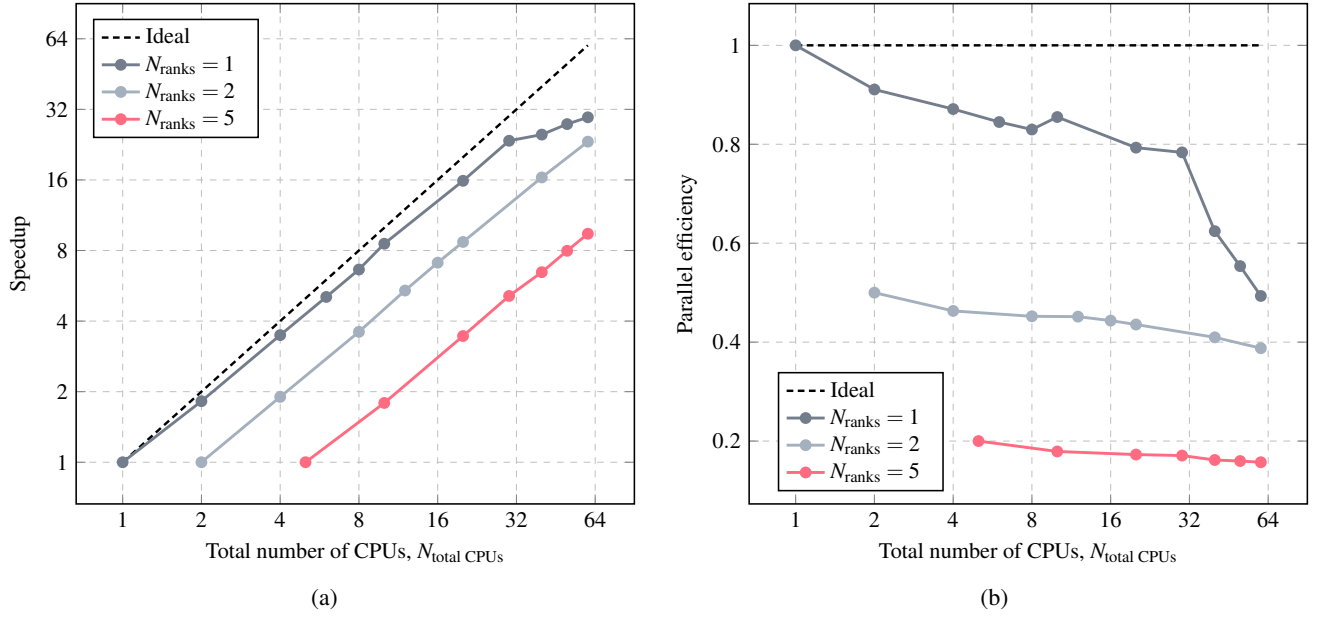


FIG. 9: Scaling performance of the multi-environment DRL training: (a) speedup ratio (b) parallel efficiency. Three different configurations for the CFD solver are used:  $N_{\text{ranks}} = 1, 2, 5$ .  $\{N_{\text{ranks}}=1, N_{\text{envs}} = 1\}$  is taken as the reference point for *all data points*.

an unexpected and significant rise in the time taken for CFD simulations when the number of environments is very large, even though these times should remain fairly consistent due to the absence of data dependencies among the environments. We postulate that this downturn may be attributed to the substantial involvement of I/O operations associated with multi-environment parallel training.




FIG. 10: Time breakdown for one episode of training. CFD simulation time predominates in all cases, but increases rapidly after  $N_{\text{envs}} > 30$ . The DRL training cost, which consists of policy update and interfacing with the CFD solver, is minimal compared to CFD simulation time.

To verify this hypothesis, we design three sets of experiments, namely *Baseline*, *I/O-Disabled*, and *Optimized*, for comparative analysis. All experiments use single-core CFD computation. Results obtained in the previous subsection are used as the baseline (*Baseline*), where the training time reduces from 225.2 hours for a single environment to 7.6 hours for 60 parallel environments. It is worth mentioning that multiple files with a total size of 5.0 MB are generated at the end of each instance of CFD simulation within each environment. For 60 parallel environments with 100 actuation periods within an episode, this is equivalent to 300



MB of CFD-related I/O at the end of each actuation period and at least 30 GB of data transfer within each episode.

In the second experimental set (*I/O-Disabled*), we determine the theoretical upper bound of performance by completely disabling all I/O operations within the workflow. Although this configuration deviates from a realistic operational setting – foregoing the essential data exchange between distinct CFD simulations and the DRL framework – it provides an exploratory benchmark indicative of the maximum attainable performance in the absence of I/O overhead. It is critical to clarify that the control outcomes from this experimental condition are unrealistic due to the lack of data transfer, but the runtime performance is sufficient to offer insight into the extent of performance gains achievable.

The third set of experiments (*Optimized*) is directed towards achieving an optimized state of I/O operations in the overall framework. The optimization measures implemented include, but are not limited to, the removal of unnecessary I/O of flow field data to reduce the number of files, and the use of binary file formats to reduce file sizes. The crux of this approach is to pare down the I/O to its functional essence, ensuring that only the necessary data exchanges that are crucial for the simulation and learning process are preserved, thereby maintaining the integrity and accuracy of the DRL and CFD simulations. With the applied I/O optimization, we are able to reduce the file size from 5.0 MB to 1.2 MB for each CFD simulation, which is a significant 76% reduction in data transfer volume.

TABLE II: Statistics of parallel multi-environment training for different I/O strategies: *Baseline*, *I/O-Disabled*, and *Optimized*.

$N_{\text{episodes}}$	$N_{\text{envs}}$	$N_{\text{ranks}}$	$N_{\text{total CPUs}}$	$T_{\text{baseline}}$ (h)	$T_{\text{I/O-disabled}}$ (h)	$T_{\text{optimized}}$ (h)
3000	1	1	1	225.2	193.1 (14%)	200.0 (11%)
3000	2	1	2	123.7	104.7 (15%)	103.8 (16%)
3000	4	1	4	64.6	53.4 (17%)	52.1 (19%)
3000	6	1	6	44.4	35.5 (20%)	35.7 (20%)
3000	8	1	8	33.9	26.3 (22%)	26.7 (21%)
3000	10	1	10	26.3	21.3 (19%)	21.5 (18%)
3000	20	1	20	14.2	11.3 (20%)	11.3 (20%)
3000	30	1	30	9.6	7.9 (18%)	8.3 (14%)
3000	40	1	40	9.0	6.4 (29%)	6.3 (30%)
3000	50	1	50	8.1	5.5 (32%)	5.3 (35%)
3000	60	1	60	7.6	4.8 (37%)	4.8 (37%)

The corresponding results are shown in Table. II, where  $T_{\text{baseline}}$ ,  $T_{\text{I/O-disabled}}$  and  $T_{\text{optimized}}$  denote the total training time for the *Baseline*, *I/O-Disabled*, and *Optimized* tests, respectively. The relative speedup of *I/O-Disabled* and *Optimized* with respect to *Baseline* is computed and included in Table. II as percentages. A comparative analysis reveals that the suspension of I/O operations yields a consistent reduction in training duration across various number of environments ( $N_{\text{envs}}$ ), indicative of the significant cost attributable to I/O processes in the baseline scenario. The degree of speedup realized through disabling I/O increases from 14% to 37%, highlighting the I/O-related bottlenecks that become more pronounced as the parallelization scales.

By optimizing I/O operations, the training time ( $T_{\text{optimized}}$ ) exhibits a reduction that closely mirrors the *I/O-Disabled* state. Notably, the *Optimized* configuration could reach the theoretical maximum speedup observed in most of the *I/O-Disabled* experiments, which underscores the efficacy of the optimization techniques applied. We observed that I/O optimization appears to be particularly effective when  $N_{\text{envs}}$  is very high. This is because the amount of I/O is directly related to the number of environments since each environment would require one set of files for data exchange with the DRL agent. With few environments, the amount of data movement is insufficient to saturate the disk I/O bandwidth and I/O costs are negligible. For large-scale parallelization, I/O bottlenecks emerge due to large amount of data transfer for multiple environments. The optimization techniques employed — such as minimizing file numbers and transitioning to binary formats — significantly alleviate the I/O overhead, thereby enhancing the overall computational efficiency.

The scaling performance of these experiments are visualized in Fig. 11 and Fig. 12. The speedup ratio and parallel efficiency for each set of experiment is computed using the respective single-environment data of each set as the reference. In Fig. 11, the *I/O-Disabled* strategy, which hypothetically eliminates all I/O constraints, closely tracks the perfect scaling line, particularly at lower environment counts. This close alignment diminishes slightly with a higher count of environments, suggesting that factors other than I/O, possibly computational overhead or inter-process communication, begin to influence performance at scale. The *Optimized* curve showcases the practical application of I/O optimization techniques, which exhibits a substantial speedup that indicates the effectiveness of the optimizations in reducing I/O overhead. The *Optimized* set of results appear above the *I/O-Disabled* curve due to differences in the reference used for computing speedup. In Fig. 12, it is evident that both strategies greatly improve the parallel efficiency. The *Optimized* strategy is close to the ideal efficiency. Similar to the speedup computation, the parallel efficiency of *Optimized* strategy is higher than the that of *I/O-Disabled* because of the choice of reference. Nonetheless, the scaling trend and achievable parallel efficiency underscores the potential of I/O optimization in large-scale, parallel DRL and CFD computations to achieve near-ideal scaling efficiency.

In short, optimizing I/O processes within DRL-based CFD simulations is crucial as it significantly enhances computational efficiency, particularly in large-scale parallel computing environments where file I/O becomes a serious bottleneck. Such opti-




FIG. 11: Speedup ratio for three I/O strategies under single-core CFD ( $N_{\text{ranks}} = 1$ ) and multi-environment training: *Baseline*, *I/O-Disabled*, and *Optimized*.  $N_{\text{envs}} = 1$  in each set of data is taken as the respective reference point.




FIG. 12: Parallel efficiency for three I/O strategies under single-core CFD ( $N_{\text{ranks}} = 1$ ) and multi-environment training: *Baseline*, *I/O-Disabled*, and *Optimized*.  $N_{\text{envs}} = 1$  in each set of data is taken as the respective reference point.

mizations directly contribute to reducing training times and improving the scalability of simulations. In our tests, refining I/O operations improves the parallel efficiency significantly, from  $\sim 49\%$  to  $\sim 76\%$  on 60 cores. For larger problems with more significant disk footprint, the runtime performance regression associated with I/O operations will be more pronounced, since the bandwidth for hard disk is significant lower than that for cache or main memory. The imperative to implement I/O optimizations is markedly more critical for efficient scalability.

#### IV. CONCLUSIONS

In this work, we have conducted a comprehensive investigation into the parallel scaling performance of DRL-based frameworks for AFC problems, shedding light on key factors influencing computational efficiency in high-performance computing environments.

- We validated a state-of-the-art DRL framework for a typical AFC problem of flow past a circular cylinder and showcased its capability to achieve drag reduction within a few hundreds of episodes. In addition, we verified that parallelization through multiple environments gives rise to similar reward convergence, indicating the potential of speedup through parallelization.
- We deconstructed the framework and analyzed parallel efficiency of the CFD component. The parallel efficiency of CFD drops rapidly in multi-core settings, reaching less than 20% for 16 MPI ranks.
- On the other hand, the parallel efficiency of multi-environment training is much higher, where 60 single-core environments could reach 49% efficiency. By extensive benchmarking of various hybrid parallelization configurations, we identified an optimal hybrid parallelization strategy that strongly favors parallelizing DRL across multiple environments over parallelizing CFD simulations. With this parallelization strategy, the training time can be reduced from 225.2 hours ( $N_{\text{ranks}} = 1, N_{\text{envs}} = 1$ ) to 7.6 hours ( $N_{\text{ranks}} = 1, N_{\text{envs}} = 60$ ), leading to a speedup of 30 times.
- To further optimize the parallelization performance, this study brought to the forefront the critical issue of I/O operations, a previously under-emphasized aspect in such simulations. By implementing targeted solutions to mitigate I/O costs, we observed a substantial enhancement in scaling efficiency from 49% to 78%. The final optimal training time is reduced to 4.8 hours ( $N_{\text{ranks}} = 1, N_{\text{envs}} = 60$  with *I/O Optimized*), which is equivalent to a remarkable speedup of approximately 47 times in runtime performance.

Overall, we were able to achieve near-linear scaling performance and significantly boost parallel efficiency, underscoring the effectiveness of the optimizations introduced.

It is important to acknowledge that there are still many unknowns and uncertainties when it comes to applying the frameworks and methodologies discussed here to different DRL-based AFC problems. Nonetheless, the theoretical trend uncovered here - that parallelizing DRL through multiple environments is generally advantageous due to reduced data dependencies - could potentially be extrapolated to other frameworks. This is particularly relevant unless the CFD component itself can be parallelized with very high efficiency comparable to that of multi-environment parallelism. Additionally, this research highlights the importance of not overlooking I/O issues, especially in situations where there is substantial data transfer between the DRL and CFD components. We believe these insights are significant for the field of DRL-based CFD simulations and provide a foundation for future directions, such as asynchronous reinforcement learning training in AFC problems.

#### DECLARATION OF INTERESTS

The authors report no conflict of interest.

#### REFERENCES

- <sup>1</sup>L. N. Cattafesta and M. Sheplak, "Actuators for active flow control," *Annual Review of Fluid Mechanics* **43**, 247–272 (2011).
- <sup>2</sup>Y. Li, J. Chang, C. Kong, and W. Bao, "Recent progress of machine learning in flow modeling and active flow control," *Chinese Journal of Aeronautics* **35**, 14–44 (2022).
- <sup>3</sup>H. Liepmann and D. Nosenchuck, "Active control of laminar-turbulent transition," *Journal of Fluid Mechanics* **118**, 201–204 (1982).
- <sup>4</sup>M. Gad-el Hak and D. M. Bushnell, "Separation control," *Journal of Fluids Engineering* **113**, 5–30 (1991).
- <sup>5</sup>M. Z. Akhter and F. K. Omar, "Review of flow-control devices for wind-turbine performance enhancement," *Energies* **14**, 1268 (2021).
- <sup>6</sup>S. L. Brunton and B. R. Noack, "Closed-loop turbulence control: Progress and challenges," *Applied Mechanics Reviews* **67**, 050801 (2015).
- <sup>7</sup>M. Gad-el Hak, "Modern developments in flow control," *Applied Mechanics Reviews* **49**, 365–379 (1996).
- <sup>8</sup>O. Coceal, A. Dobre, T. Thomas, and S. Belcher, "Structure of turbulent flow over regular arrays of cubical roughness," *Journal of Fluid Mechanics* **589**, 375–409 (2007).
- <sup>9</sup>F. Pino, L. Schena, J. Rabault, and M. A. Mendez, "Comparative analysis of machine learning methods for active flow control," *Journal of Fluid Mechanics* **958**, A39 (2023).
- <sup>10</sup>D. Shi, B. Lam, W.-S. Gan, J. Cheer, and S. J. Elliott, "Active noise control in the new century: The role and prospect of signal processing," in *INTER-NOISE and NOISE-CON Congress and Conference Proceedings*, InterNoise23, Vol. 268(3) (Institute of Noise Control Engineering, Chiba, Japan, 2023) pp. 5141–5151.
- <sup>11</sup>A. M. Omer, "Energy, environment and sustainable development," *Renewable and Sustainable Energy Reviews* **12**, 2265–2300 (2008).
- <sup>12</sup>M. Jahanmiri, "Active flow control: a review," *Flow Measurement and Instrumentation* **21**, 7–28 (2010).

- <sup>13</sup>D. P. Rizzetta, M. R. Visbal, and M. J. Stanek, "Numerical investigation of synthetic-jet flowfields," *AIAA Journal* **37**, 919–927 (1999).
- <sup>14</sup>P. Garnier, J. Viquerat, J. Rabault, A. Larcher, A. Kuhnle, and E. Hachem, "A review on deep reinforcement learning for fluid mechanics," *Computers & Fluids* **225**, 104973 (2021).
- <sup>15</sup>J. C. Lin, M. Y. Andino, M. G. Alexander, E. A. Whalen, M. A. Spoor, J. T. Tran, and I. J. Wagnanski, "An overview of active flow control enhanced vertical tail technology development," in *54th AIAA Aerospace Sciences Meeting* (2016).
- <sup>16</sup>S. S. Collis and R. D. Joslin, "Issues in active flow control: theory, control, simulation, and experiment," *Progress in aerospace sciences* **40**, 237–289 (2004).
- <sup>17</sup>W. He, Z. bing Luo, X. Deng, C. Peng, Q. Liu, T. xiang Gao, P. Cheng, Y. Zhou, and W. qiang Peng, "Numerical study on the atomization mechanism and energy characteristics of synthetic jet/dual synthetic jets," *Applied Energy* **346**, 121376 (2023).
- <sup>18</sup>L. D. Kral, "Active flow control technology," *ASME Fluids Engineering Technical Brief*, 1–28 (2000).
- <sup>19</sup>R. King, *Active Flow Control*, 1st ed. (Springer Berlin, Heidelberg, 2007) pp. XIII, 442.
- <sup>20</sup>T. R. Bewley, "Flow control: new challenges for a new renaissance," *Progress in Aerospace Sciences* **37**, 21–58 (2001).
- <sup>21</sup>S. L. Brunton and B. R. Noack, "Closed-Loop Turbulence Control: Progress and Challenges," *Applied Mechanics Reviews* **67**, 050801 (2015).
- <sup>22</sup>S. Low and D. Lapsley, "Optimization flow control. i. basic algorithm and convergence," *IEEE/ACM Transactions on Networking* **7**, 861–874 (1999).
- <sup>23</sup>M. Gerla and L. Kleinrock, "Flow control: A comparative survey," *IEEE Transactions on Communications* **28**, 553–574 (1980).
- <sup>24</sup>K. Arulkumar, M. P. Deisenroth, M. Brundage, and A. A. Bharath, "Deep reinforcement learning: A brief survey," *IEEE Signal Processing Magazine* **34**, 26–38 (2017).
- <sup>25</sup>V. François-Lavet, P. Henderson, R. Islam, M. G. Bellemare, J. Pineau, *et al.*, "An introduction to deep reinforcement learning," *Foundations and Trends® in Machine Learning* **11**, 219–354 (2018).
- <sup>26</sup>X. Wang, S. Wang, X. Liang, D. Zhao, J. Huang, X. Xu, B. Dai, and Q. Miao, "Deep reinforcement learning: A survey," *IEEE Transactions on Neural Networks and Learning Systems*, 1–15 (2022).
- <sup>27</sup>V. François-Lavet, P. Henderson, R. Islam, M. G. Bellemare, J. Pineau, *et al.*, "An introduction to deep reinforcement learning," *Foundations and Trends® in Machine Learning* **11**, 219–354 (2018).
- <sup>28</sup>N. Gautier, J.-L. Aider, T. Duriez, B. R. Noack, M. Segond, and M. Abel, "Closed-loop separation control using machine learning," *Journal of Fluid Mechanics* **770**, 442–457 (2015).
- <sup>29</sup>M. Gazzola, B. Hejiazhosseini, and P. Koumoutsakos, "Reinforcement learning and wavelet adapted vortex methods for simulations of self-propelled swimmers," *SIAM Journal on Scientific Computing* **36**, B622–B639 (2014).
- <sup>30</sup>G. Novati, S. Verma, D. Alexeev, D. Rossinelli, W. M. Van Rees, and P. Koumoutsakos, "Synchronisation through learning for two self-propelled swimmers," *Bioinspiration & Biomimetics* **12**, 036001 (2017).
- <sup>31</sup>S. Verma, G. Novati, and P. Koumoutsakos, "Efficient collective swimming by harnessing vortices through deep reinforcement learning," *Proceedings of the National Academy of Sciences* **115**, 5849–5854 (2018).
- <sup>32</sup>L. Yan, X. Chang, R. Tian, N. Wang, L. Zhang, and W. Liu, "A numerical simulation method for bionic fish self-propelled swimming under control based on deep reinforcement learning," *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science* **234**, 3397–3415 (2020).
- <sup>33</sup>P. Ma, Y. Tian, Z. Pan, B. Ren, and D. Manocha, "Fluid directed rigid body control using deep reinforcement learning," *ACM Transactions on Graphics* **37**, 1–11 (2018).
- <sup>34</sup>J. Rabault, M. Kuchta, A. Jensen, U. Réglade, and N. Cerardi, "Artificial neural networks trained through deep reinforcement learning discover control strategies for active flow control," *Journal of fluid mechanics* **865**, 281–302 (2019).
- <sup>35</sup>J. Rabault and A. Kuhnle, "Accelerating deep reinforcement learning strategies of flow control through a multi-environment approach," *Physics of Fluids* **31**, 094105 (2019).
- <sup>36</sup>D. Xu and M. Zhang, "Reinforcement-learning-based control of convectively-unstable flows," *Journal of Fluid Mechanics* **954**, A37 (2023).
- <sup>37</sup>V. Belus, J. Rabault, J. Viquerat, Z. Che, E. Hachem, and U. Reglade, "Exploiting locality and translational invariance to design effective deep reinforcement learning control of the 1-dimensional unstable falling liquid film," *AIP Advances* **9**, 125014 (2019).
- <sup>38</sup>G. Beintema, A. Corbetta, L. Biferale, and F. Toschi, "Controlling Rayleigh–Bénard convection via reinforcement learning," *Journal of Turbulence* **21**, 585–605 (2020).
- <sup>39</sup>Y.-Z. Wang, X.-J. He, Y. Hua, Z.-H. Chen, W.-T. Wu, and Z.-F. Zhou, "Closed-loop forced heat convection control using deep reinforcement learning," *International Journal of Heat and Mass Transfer* **202**, 123655 (2023).
- <sup>40</sup>Y.-Z. Wang, J.-Z. Peng, N. Aubry, Y.-B. Li, Z.-H. Chen, and W.-T. Wu, "Control policy transfer of deep reinforcement learning based intelligent forced heat convection control," *International Journal of Thermal Sciences* **195**, 108618 (2024).
- <sup>41</sup>M. Tokarev, E. Palkin, and R. Mullyadzhannov, "Deep reinforcement learning control of cylinder flow using rotary oscillations at low reynolds number," *Energies* **13**, 5920 (2020).
- <sup>42</sup>H. Xu, W. Zhang, J. Deng, and J. Rabault, "Active flow control with rotating cylinders by an artificial neural network trained by deep reinforcement learning," *Journal of Hydrodynamics* **32**, 254–258 (2020).
- <sup>43</sup>H. Koizumi, S. Tsutsumi, and E. Shima, "Feedback control of karman vortex shedding from a cylinder using deep reinforcement learning," in *2018 Flow Control Conference* (2018).
- <sup>44</sup>M. Kurz, P. Offenhäuser, D. Viola, O. Shcherbakov, M. Resch, and A. Beck, "Deep reinforcement learning for computational fluid dynamics on HPC systems," *Journal of Computational Science* **65**, 101884 (2022).
- <sup>45</sup>M. Schäfer, S. Turek, F. Durst, E. Krause, and R. Rannacher, "Benchmark computations of laminar flow around a cylinder," in *Flow Simulation with High-Performance Computers II*, Notes on Numerical Fluid Mechanics (NNFM), Vol. 48, edited by E. H. Hirschel (Vieweg+Teubner Verlag, 1996) pp. 547–566.
- <sup>46</sup>H. Jasak, A. Jemcov, Z. Tukovic, *et al.*, "OpenFOAM: : A C++ library for complex physics simulations," in *International Workshop on Coupled Methods in Numerical Dynamics*, Vol. 1000 (2007) pp. 1–20.
- <sup>47</sup>R. I. Issa, "Solution of the implicitly discretised fluid flow equations by operator-splitting," *Journal of Computational Physics* **62**, 40–65 (1986).
- <sup>48</sup>S. Patankar, *Numerical Heat Transfer and Fluid Flow* (Hemisphere Publishing Corporation, 1980) provided by the SAO/NASA Astrophysics Data System.
- <sup>49</sup>D. Jang, R. Jetli, and S. Acharya, "Comparison of the piso, simpler, and simplec algorithms for the treatment of the pressure-velocity coupling in steady flow problems," *Numerical Heat Transfer* **10**, 209–228 (1986).
- <sup>50</sup>Q. Wang, L. Yan, G. Hu, C. Li, Y. Xiao, H. Xiong, J. Rabault, and B. R. Noack, "DRLinFluids: An open-source Python platform of coupling deep reinforcement learning and OpenFOAM," *Physics of Fluids* **34**, 081801 (2022).
- <sup>51</sup>K. Arulkumar, M. P. Deisenroth, M. Brundage, and A. A. Bharath, "A Brief Survey of Deep Reinforcement Learning," *IEEE Signal Processing Magazine* **34**, 26–38 (2017).
- <sup>52</sup>P. Ladosz, L. Weng, M. Kim, and H. Oh, "Exploration in Deep Reinforcement Learning: A Survey," *Information Fusion* **85**, 1–22 (2022), 2205.00824 [cs].
- <sup>53</sup>A. K. Shakya, G. Pillai, and S. Chakrabarty, "Reinforcement learning algorithms: A brief survey," *Expert Systems with Applications* **231**, 120495 (2023).

- <sup>54</sup>R. Bellman, "A markovian decision process," *Journal of Mathematics and Mechanics* **6**, 679–684 (1957).
- <sup>55</sup>M. L. Puterman, "Markov decision processes," *Handbooks in operations research and management science* **2**, 331–434 (1990).
- <sup>56</sup>M. van Otterlo and M. Wiering, "Reinforcement learning and markov decision processes," in *Reinforcement Learning: State-of-the-Art*, Vol. 2012 (Springer Berlin Heidelberg, Berlin, Heidelberg, 2012) pp. 3–42.
- <sup>57</sup>K. Arulkumaran, M. P. Deisenroth, M. Brundage, and A. A. Bharath, "Deep reinforcement learning: A brief survey," *IEEE Signal Processing Magazine* **34**, 26–38 (2017).
- <sup>58</sup>J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," (2017), arXiv:1707.06347 [cs.LG].
- <sup>59</sup>J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz, "Trust region policy optimization," in *Proceedings of the 32nd International Conference on Machine Learning*, Proceedings of Machine Learning Research, Vol. 37 (2015) pp. 1889–1897.
- <sup>60</sup>N. Heess, D. Tb, S. Sriram, J. Lemmon, J. Merel, G. Wayne, Y. Tassa, T. Erez, Z. Wang, S. Eslami, *et al.*, "Emergence of locomotion behaviours in rich environments," (2017), arXiv:1707.02286 [cs.AI].
- <sup>61</sup>J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz, "Trust region policy optimization," (2017), arXiv:1502.05477 [cs.LG].
- <sup>62</sup>N. Parikh, S. Boyd, *et al.*, "Proximal algorithms," *Foundations and Trends® in Optimization* **1**, 127–239 (2014).
- <sup>63</sup>C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, "Policy gradient methods for reinforcement learning with function approximation," in *Advances in Neural Information Processing Systems*, Vol. 12 (1999).
- <sup>64</sup>R. Paris, S. Beneddine, and J. Dandois, "Robust flow control and optimal sensor placement using deep reinforcement learning," *Journal of Fluid Mechanics* **913**, A25 (2021).
- <sup>65</sup>J. Li and M. Zhang, "Reinforcement-learning-based control of confined cylinder wakes with stability analyses," *Journal of Fluid Mechanics* **932**, A44 (2022).
- <sup>66</sup>R. Castellanos, G. Y. Cornejo Maceda, I. de la Fuente, B. R. Noack, A. Ianiro, and S. Discetti, "Machine-learning flow control with few sensor feedback and measurement noise," *Physics of Fluids* **34**, 047118 (2022).
- <sup>67</sup>A. Kuhnle and M. Schaarschmidt, "Tensorforce: a tensorflow library for applied reinforcement learning," Web page (2017).
- <sup>68</sup>D. M. Ritchie, "The unix system: A stream input-output system," *AT&T Bell Laboratories Technical Journal* **63**, 1897–1910 (1984).
- <sup>69</sup>K. Thompson, "Programming techniques: Regular expression search algorithm," *Commun. ACM* **11**, 419–422 (1968).
- <sup>70</sup>M. Schaarschmidt, A. Kuhnle, B. Ellis, K. Fricke, F. Gessert, and E. Yoneki, "Lift: Reinforcement learning in computer systems by learning from demonstrations," (2018), arXiv:1808.07903 [cs.LG].
- <sup>71</sup>V. Starikovičius, R. Čiegis, and A. Bugajev, "On Efficiency Analysis of the OpenFOAM-Based Parallel Solver for Simulation of Heat Transfer in and Around the Electrical Power Cables," *Informatica* **27**, 161–178 (2016).
- <sup>72</sup>M. Towara, M. Schanen, and U. Naumann, "MPI-Parallel Discrete Adjoint OpenFOAM," *Procedia Computer Science* **51**, 19–28 (2015).
- <sup>73</sup>M. Culp, "Current bottlenecks in the scalability of openfoam on massively parallel clusters," PRACE White Paper (2011).
- <sup>74</sup>G. Janeš, A. Sikirica, L. Grbčić, and L. Kranjčević, "MPI Associated Scalability of Open-Source CFD Codes for Oil Spill Assessment," *Journal of Maritime & Transportation Science Special edition* **4**, 67–75 (2022).
- <sup>75</sup>W. Liu, W. Wang, A. Skillen, S. Longshaw, C. Moulinec, and D. Emerson, "A Parallel Partitioned Approach On Fluid-Structure Interaction Simulations Using the Multiscale Universal Interface Coupling Library," in *14th WCCM-ECCOMAS Congress (CIMNE)*, (2021).
- <sup>76</sup>F. Capuano, A. Palumbo, and L. De Luca, "Comparative study of spectral-element and finite-volume solvers for direct numerical simulation of synthetic jets," *Computers & Fluids* **179**, 228–237 (2019).
- <sup>77</sup>F. Witherden, A. Farrington, and P. Vincent, "Pyfr: An open source framework for solving advection–diffusion type problems on streaming architectures using the flux reconstruction approach," *Computer Physics Communications* **185**, 3028–3040 (2014).
- <sup>78</sup>C. Cantwell, D. Moxey, A. Comerford, A. Bolis, G. Rocco, G. Mengaldo, D. De Grazia, S. Yakovlev, J.-E. Lombard, D. Ekelschot, B. Jordi, H. Xu, Y. Mohamied, C. Eskilsson, B. Nelson, P. Vos, C. Biotto, R. Kirby, and S. Sherwin, "Nektar++: An open-source spectral/hp element framework," *Computer Physics Communications* **192**, 205–219 (2015).