

Bioinspiration & Biomimetics



PAPER

Synchronisation through learning for two self-propelled swimmers

RECEIVED
13 October 2016

REVISED
5 January 2017

ACCEPTED FOR PUBLICATION
27 February 2017

PUBLISHED
29 March 2017

Guido Novati^{1,5}, Siddhartha Verma^{1,5}, Dmitry Alexeev¹, Diego Rossinelli¹, Wim M van Rees^{1,2}
and Petros Koumoutsakos^{1,3,4}

¹ Computational Science and Engineering Laboratory, Clausiusstrasse 33, ETH Zürich, CH-8092, Switzerland

² School of Engineering and Applied Sciences, Harvard University, MA, United States of America

³ Radcliffe Institute of Advanced Study, Harvard University, MA, United States of America

⁴ Wallace Visiting Professor, Massachusetts Institute of Technology, MA, United States of America

⁵ These authors contributed equally to this work.

E-mail: petros@ethz.ch

Keywords: swimming, synchronisation, learning, hydrodynamics

Supplementary material for this article is available [online](#)

Abstract

The coordinated motion by multiple swimmers is a fundamental component in fish schooling. The flow field induced by the motion of each self-propelled swimmer implies non-linear hydrodynamic interactions among the members of a group. How do swimmers compensate for such hydrodynamic interactions in coordinated patterns? We provide an answer to this riddle through simulations of two, self-propelled, fish-like bodies that employ a learning algorithm to synchronise their swimming patterns. We distinguish between learned motion patterns and the commonly used a-priori specified movements, that are imposed on the swimmers without feedback from their hydrodynamic interactions. First, we demonstrate that two rigid bodies executing pre-specified motions, with an alternating leader and follower, can result in substantial drag-reduction and intermittent thrust generation. In turn, we study two self-propelled swimmers arranged in a leader-follower configuration, with a-priori specified body-deformations. These two self-propelled swimmers do not sustain their tandem configuration. The follower experiences either an increase or decrease in swimming speed, depending on the initial conditions, while the swimming of the leader remains largely unaffected. This indicates that a-priori specified patterns are not sufficient to sustain synchronised swimming. We then examine a tandem of swimmers where the leader has a steady gait and the follower learns to synchronize its motion, to overcome the forces induced by the leader's vortex wake. The follower employs reinforcement learning to adapt its swimming-kinematics so as to minimize its lateral deviations from the leader's path. Swimming in such a sustained synchronised tandem yields up to 30% reduction in energy expenditure for the follower, in addition to a 20% increase in its swimming-efficiency. The present results show that two self-propelled swimmers can be synchronised by adapting their motion patterns to compensate for flow-structure interactions. Moreover, swimmers can exploit the vortical structures of their flow field so that synchronised swimming is energetically beneficial.

1. Introduction

The coordinated motion of fish is thought to provide an energetic advantage to individuals, as well as to their group, in terms of increased swimming range, endurance and chances of survival. Schooling has also been credited with serving diverse biological functions including defence from predators [8, 9], enhanced feeding and reproductive opportunities [28]. At present, there is no consensus regarding the evolutionary purpose of schooling behaviour [27]. However, there is growing evidence supporting the

hypothesis that fluid dynamics affects swimming patterns in fish schools [35] and related experimental model configurations [3] as well as flying patterns of birds [29].

Experiments that have investigated swimming of fish groups indicate a reduction in energy expenditure, based on respirometer readings and reduced tail-beat frequency [1, 18, 20, 24, 33]. Importantly, there is evidence to suggest that reduction in energy expenditure is not distributed uniformly throughout a schooling group. Several studies [18, 20, 33] have observed that the tail-beat frequency of trailing fish was lower than

that of fish at the front of the school. Moreover, Killen *et al* [20] note that fish with inherently lower aerobic scope prefer to stay towards the rear of a group. Studies investigating the response of solitary fish to unsteady flow [22] found that trout swimming behind obstacles exerted reduced effort for station-keeping. The trout adopted a gait which allowed them to ‘slalom’ through the oncoming vortex street. The ensuing reduction in muscle activity was confirmed using neuromuscular measurements [22] and respirometer readings [34]. These experimental studies suggest that fish can detect and exploit vortices present in their surroundings [35].

There is a well documented hypothesis [7, 39, 40] that flow patterns which emerge as a consequence of schooling, can be exploited by individual swimmers. This hypothesis was first quantified [39, 40] using inviscid point-vortices as models of the fish wake-vortices. It was postulated that large groups of fish could gain a propulsive advantage by swimming in a ‘diamond’ configuration, with opposing tail-beat phase. The energetic gain was attributed to two distinct mechanisms: drag-reduction resulting from decreased relative velocity in the vicinity of specific vortices; and a forward ‘push’ originating from a ‘channelling effect’ between lateral neighbors. Weihs noted that a rigid geometrical arrangement, and perfectly synchronized anti-phase swimming among lateral neighbours, were unlikely to occur in nature. Nonetheless, he postulated that given the immense potential for energy savings, even intermittent utilization of the proposed arrangement could lead to a tangible benefit [40]. Such simplified models of hydrodynamics in fish schools have even inspired the optimal design of wind turbine farms [41]. We note also that synchronised motion has been obtained for viscous simulations of two cylinders performing rotary oscillations [38]. The role of hydrodynamics in fish-schooling was later questioned [26], based on empirical observations of fish-schools which rarely displayed diamond formations. However, a later study based on aerial photographs of hunting tuna schools [25] provided evidence for such diamond-like formations. We believe that these studies highlight the difficulties of maintaining fixed patterns in the dynamically evolving environment of schooling fish. These difficulties are also reflected in simulations studies. Bergmann and Iollo [4] presented short-time simulation of two and three self-propelled swimmers with prescribed kinematics that may eventually lead to the collision or divergence of the swimmers. Long-time simulations of multiple swimmers, have either pre-specified their spatial distribution and kinematics [10, 17] or they have employed potential flow models [15, 36]. We note that simulations of single swimmers performing optimised swimming motions have been performed over the last ten years [5, 11, 16, 19]. To the best of our knowledge, no viscous simulations have ever been presented for multiple self-propelled swimmers with sustained synchronised motions.

Here we present two-dimensional simulations of viscous, incompressible flows of self-propelled

swimmers in coordinated swimming patterns that can dynamically adapt their motion. We focus on two swimmers in a sustained leader-follower configuration, a biologically relevant schooling pattern characterized by the follower interacting with the unsteady flow in the leader’s wake. We investigate the hydrodynamic interactions of the swimmers in various scenarios including pre-specified coordinated motions and initial distances, as well as the dynamic adaptation of the follower’s motion using a reinforcement learning algorithm, so as to remain within a specific region in the leader’s wake. We investigate the impact of the leader’s wake on the follower’s motion and identify the mechanisms that lead to energy savings. The paper is organised as follows: we outline the numerical methods for the simulations of self-propelled swimmers in section 2, and the reinforcement learning algorithm is discussed in section 2.2. We present results and discussion of the three synchronised swimming scenarios in section 3, followed by concluding remarks in section 4.

2. Methods

2.1. Swimmers and Numerical Methods

We solve the two-dimensional, viscous, incompressible Navier–Stokes equations in velocity–vorticity form using remeshed vortex methods on wavelet-adapted grids [30], and a divergence-free penalisation technique to enforce the no-slip boundary condition [13]. The wavelet adaptivity and the computational efficiency of the solver are critical aspects for this work as they enable the utilisation of the costly reinforcement learning algorithms. The momentum equation is expressed in velocity–pressure form as:

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = \frac{-\nabla P}{\rho} + \nu \nabla^2 \mathbf{u} + \lambda (\chi(\mathbf{u}_s - \mathbf{u})) \quad (1)$$

Here, λ is the penalization parameter and $\chi(\mathbf{x})$ is the level set function representing the discretized solid on the computational grid, while \mathbf{u}_s denotes the velocity vector of the solid object, and accounts for rotation, translation, and deformation of the object.

The self-propelled swimmers used in the simulations are based on a simplified physical model of zebrafish as described in [13]. Undulations of the swimmer’s body are generated by imposing a spatially and temporally varying body curvature ($k(s, t)$), which passes down from the head to the tail as a sinusoidal travelling wave:

$$k(s, t) = A(s) \sin \left[2\pi \left(\frac{t}{T_p} - \frac{s}{L} \right) + \phi \right] \quad (2)$$

Here, L is the length of the swimmer, $T_p = 1$ is the tail-beat period, and ϕ is a phase-shift. The curvature amplitude $A(s)$ varies linearly from the head to the tail, with $A(0) = 0.82$ and $A(L) = 5.7$, thus reducing head motion and amplifying tail-beat amplitude. We

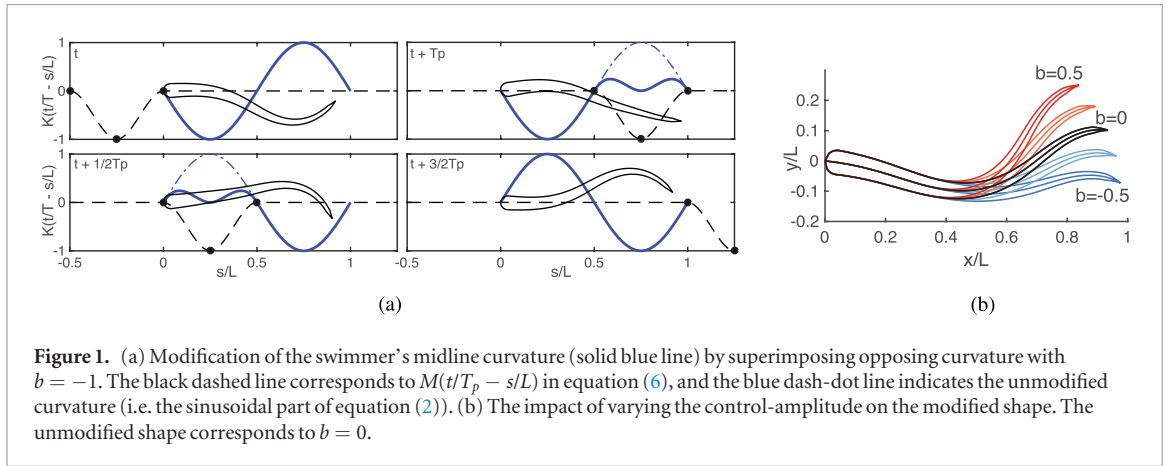


Figure 1. (a) Modification of the swimmer's midline curvature (solid blue line) by superimposing opposing curvature with $b = -1$. The black dashed line corresponds to $M(t/T_p - s/L)$ in equation (6), and the blue dash-dot line indicates the unmodified curvature (i.e. the sinusoidal part of equation (2)). (b) The impact of varying the control-amplitude on the modified shape. The unmodified shape corresponds to $b = 0$.

estimate the swimming-efficiency using a modified form of the Froude efficiency proposed by [37]:

$$\eta = \frac{P_{\text{thrust}}}{P_{\text{thrust}} + \max(P_{\text{def}}, 0)} = \frac{Tu}{Tu + \max(-\int \int \mathbf{u}_{\text{def}} \cdot d\mathbf{F}, 0)} \quad (3)$$

P_{thrust} and P_{def} represent the power output related to thrust generated by the body, and the power exerted in deforming the swimmer's body against fluid-induced forces. The max operator effectively clips the maximum of η to 1. This is necessary to avoid undefined values of η , which can occur when fluid-induced surface-forces and the deformational velocity (\mathbf{u}_{def}) point in the same direction, giving rise to negative P_{def} . The thrust is computed as $T = \int \int (\mathbf{u} \cdot d\mathbf{F} + |\mathbf{u} \cdot d\mathbf{F}|)/(2|\mathbf{u}|)$, where $d\mathbf{F} = d\mathbf{F}_p + d\mathbf{F}_v$ is comprised of the viscous- and pressure-based forces acting on the swimmer:

$$d\mathbf{F}_v = 2\mu\mathbf{D} \cdot \mathbf{n} dS \quad \text{and} \quad d\mathbf{F}_p = -P\mathbf{n} dS \quad (4)$$

Here, $\mathbf{D} = (\nabla\mathbf{u} + \nabla\mathbf{u}^T)/2$ is the strain-rate tensor, P is the surface-pressure, and μ is the dynamic viscosity. \mathbf{n} represents the surface-normal and dS denotes the corresponding infinitesimal surface area. The pressure field is computed by assuming neutrally buoyant swimmers so that taking the divergence of equation (1) we obtain a Poisson's equation $\nabla^2 P = -\rho(\nabla\mathbf{u}^T : \nabla\mathbf{u}) + \rho\lambda\nabla \cdot (\chi(\mathbf{u}_s - \mathbf{u}))$.

2.2. Reinforcement learning

Swimming dynamics of multiple, self-propelled swimmers imply complex vortex-body interactions. In the case of two swimmers the follower experiences the vorticity field of the leader and has to react to it so as to maintain the tandem arrangement. We adapt the motions of the follower, to overcome the vortex street of the leader, by employing a Reinforcement Learning (RL) [32] algorithm. RL is well suited to flow control problems [12, 15] as it provides a feedback controller by maximising a numeric reward signal without requiring any prior knowledge about the underlying dynamics. In RL, the swimmers receive information about their *State* and choose *Actions* to maximise a cumulative future *Reward* in an unsupervised manner.

The swimmer learns to estimate the action-value $Q(s, a)$ which is defined as the expected sum of the discounted future rewards r , for each action a being performed in each state s . The reward is obtained by starting in s , performing a to end up in a new state s' , and thereafter following the policy $\pi(s')$:

$$Q_\pi(s, a) = \mathbb{E}[r + \gamma Q(s', \pi(s'))] \quad (5)$$

The discount factor $\gamma \in [0, 1]$ (set to $\gamma = 0.8$ for all present results) determines the trade-off between immediate and future rewards. The learning process terminates upon convergence of $Q(s, a)$, and the swimmer can make optimal decisions by following a 'greedy' policy ($\pi(s) = \arg \max_a Q_\pi(s, a)$). An important aspect of this work is that the swimmer learns a parameterized approximation of $Q(s, a)$ by training a Neural Network with experience replay [23]. This involves storing all observed transitions $\{s, a, s', r\}$ and iteratively drawing samples to train the value function.

Actions taken by a swimmer involve manipulating its body curvature in a manner which allows it to execute turns and to control its speed. This is achieved by introducing a linear superposition to the travelling wave described in equation (2):

$$k_{\text{learner}}(s, t) = k(s, t) + k'(s, t) = A(s) \left(\sin \left[2\pi \left(\frac{t}{T_p} - \frac{s}{L} \right) + \phi \right] + M \left(\frac{t}{T_p} - \frac{s}{L} \right) \right) \quad (6)$$

$M \left(\frac{t}{T_p} - \frac{s}{L} \right)$ defines a travelling natural cubic spline, computed using three evenly-spaced nodes separated by a distance of $L/4$ (figure 1(a)):

$M \left(\frac{t}{T_p} - \frac{s}{L} \right) = \sum_k b_k \cdot m \left(\frac{t-t_k}{T_p} - \frac{s}{L} \right)$. Here b_k is the control amplitude associated with the action taken at time t_k and m is composed of two cubic splines such that $m(0) = m'(0) = m(\frac{1}{2}) = m'(\frac{1}{2}) = m'(\frac{1}{4}) = 0$, $m(\frac{1}{4}) = 1$. The swimmers act in every half tail-beat period, to either increase or reduce the undulation amplitude (figures 1(a) and (b)). In practice, at most three previous actions affect the deformation of the swimmer at any time.

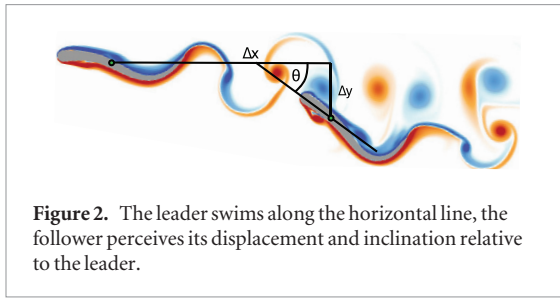


Figure 2. The leader swims along the horizontal line, the follower perceives its displacement and inclination relative to the leader.

In the present two fish swimming (leader-follower tandem) problem, learning is performed only by the follower. The follower defines its current *State* using its displacement (Δx and Δy) and orientation (θ) relative to the leader (figure 2). Moreover, the follower considers as part of its state whether it is taking an action in the first or second half of its tail-beat period (modulo(t, T_p)/ T_p). This is necessary because the same action can either increase or decrease the mid-line curvature, depending on whether it is taken at T_p or at $T_p/2$ (figure 1(a)). Furthermore, the swimmer performs two actions every period, as described earlier. The effect of each action travels along the swimmer's body, affecting its interaction with the flow over the next swimming period. For this reason, the state of the swimmer also includes the two actions taken over the previous tail-beat period.

Here, the *Reward* used to provide feedback to the follower regarding its performance is defined as: $\mathcal{R}_{\Delta y} = 1 - \frac{|\Delta y|}{0.5L}$. This function penalizes the follower when it deviates laterally from the path of the leader. The follower learns the policy $\pi_{\Delta y}$, which maximizes $\mathcal{R}_{\Delta y}$. The state space is restricted to $1 \leq \Delta x/L \leq 3$, $|\Delta y/L| \leq 0.5$ and $|\theta| \leq \pi/2$. When an action leads the follower to exceed these thresholds, the learner transitions to a terminal state with reward $\mathcal{R}_{\Delta y} = -1$, and the simulation is terminated.

3. Results and discussion

We distinguish externally imposed motions on the swimmers to those that are achieved by the deformation of the body of self-propelled swimmers. In this section we discuss results concerning three distinct scenarios, namely, two rigid airfoils executing pre-specified motion, two self-propelled swimmers interacting without control, and a 'smart' follower utilizing adaptive control to interact with the leader's wake. An effective resolution of 2048^2 grid points was used in all the simulations, with a domain size of $[0, 10L] \times [0, 10L]$, where L is the typical length of the solid objects. The Lagrangian CFL number was set to 0.1, with the resulting time-step size ranging from $1e-4$ to $1e-3$.

3.1. Rigid objects with pre-specified motion

We investigate the potential energy-savings of interacting swimmers, starting with two rigid, airfoil-shaped bodies (shape identical to swimmers) with

a-priori specified motion. We mimic a swimming pattern, often observed in schooling, which involves exchanging the positions of the leader and the follower. We drag the two objects along prescribed intersecting sinusoidal paths (figure 3(a)), with an acceleration $a_x = (u_{\max} - u_{\min})/T_s$ to periodically exchange their position as leader and follower (see supplementary movie 1 stacks.iop.org/BB/12/036001/mmedia). Here, $u_{\max} = 4.5L/T_s$, $u_{\min} = 1.5L/T_s$, and T_s represents the time-period with which the bodies exchange their position as leader and follower. The vertical displacement of the center-of-mass is determined as $y(\Delta x, L) = L/5 \cos(\pi \Delta x/L)$, where Δx is the horizontal distance traversed. The orientation of the airfoils is aligned with the tangents of their respective trajectories. Both the airfoils start their motion at the same x -location; one of the objects is initialized at a crest with u_{\max} and undergoes steady deceleration (a_x), whereas the other object starts with u_{\min} on a trough and is subjected to constant acceleration ($-a_x$). This arrangement of positions and velocities alternates between the two airfoils every time-period T_s . The resulting Strouhal number ranges from 0.11 to 0.33, based on the period of oscillation, the airfoil length, and the minimum and maximum speed.

A snapshot of the vorticity field, along with the sinusoidal path followed by the two airfoils, is shown in figure 3. The flow pattern that emerges influences the net drag acting on the two objects. Despite the flow separation caused by the large angle of attack of the prescribed motion (see figure 3(a)), we observe that the follower experiences a dramatic reduction in drag (see figure 3(b) at $t \approx 3.9$). This can be attributed to a decrease in relative velocity, due to the presence of the positive vortex highlighted in figure 3(a). The drag-reduction at this time instance is greater than 100%, which corresponds to a net thrust being generated due to the interaction of the follower's motion with the wake. Moreover, figure 3(b) indicates that both the leader and the follower may experience a reduction in drag as a result of mutual interaction. To examine the drag-reducing mechanism in detail, we inspect the vorticity field, and the resulting fluid-induced force-distribution on the airfoils in figure 4. The distribution on the lower surface of the trailing airfoil (figure 4(c)) indicates a weakening of deceleration-producing forces (i.e. force-vectors that point away from the direction of motion), due to interactions with the vortex entrained from the leader's wake. A similar effect is not observed in the absence of a leading airfoil (figure 4(d)). The mechanism at work may be understood by considering the rate of change of linear impulse of a body due to the 2D vortical flow field:

$$F_w = -\rho \frac{d}{dt} \int \mathbf{x} \times \boldsymbol{\omega} dV \quad (7)$$

Approximating the vortex field by a linear superposition of vortical structures with strength Γ_i at locations (x_i, y_i) [21] the horizontal force acting on the object due to the vorticity in the flow is expressed as:

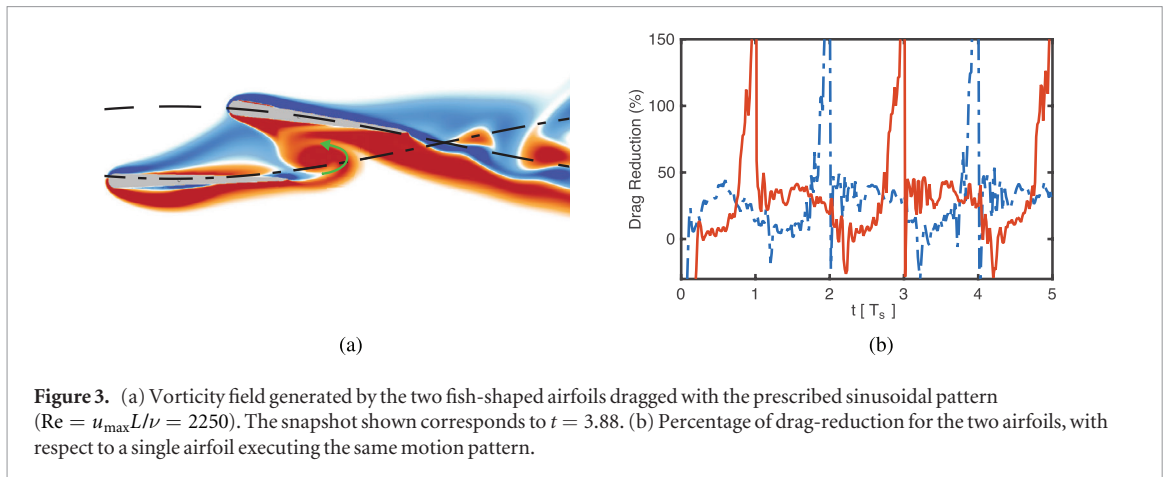


Figure 3. (a) Vorticity field generated by the two fish-shaped airfoils dragged with the prescribed sinusoidal pattern ($Re = u_{\max}L/\nu = 2250$). The snapshot shown corresponds to $t = 3.88$. (b) Percentage of drag-reduction for the two airfoils, with respect to a single airfoil executing the same motion pattern.

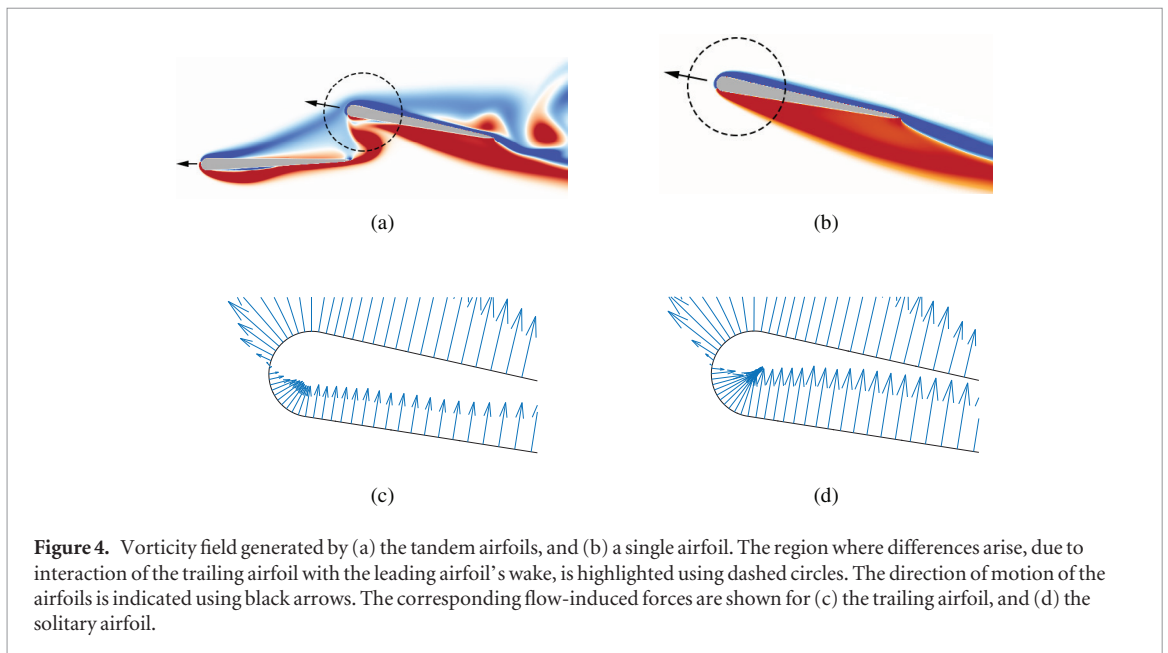


Figure 4. Vorticity field generated by (a) the tandem airfoils, and (b) a single airfoil. The region where differences arise, due to interaction of the trailing airfoil with the leading airfoil's wake, is highlighted using dashed circles. The direction of motion of the airfoils is indicated using black arrows. The corresponding flow-induced forces are shown for (c) the trailing airfoil, and (d) the solitary airfoil.

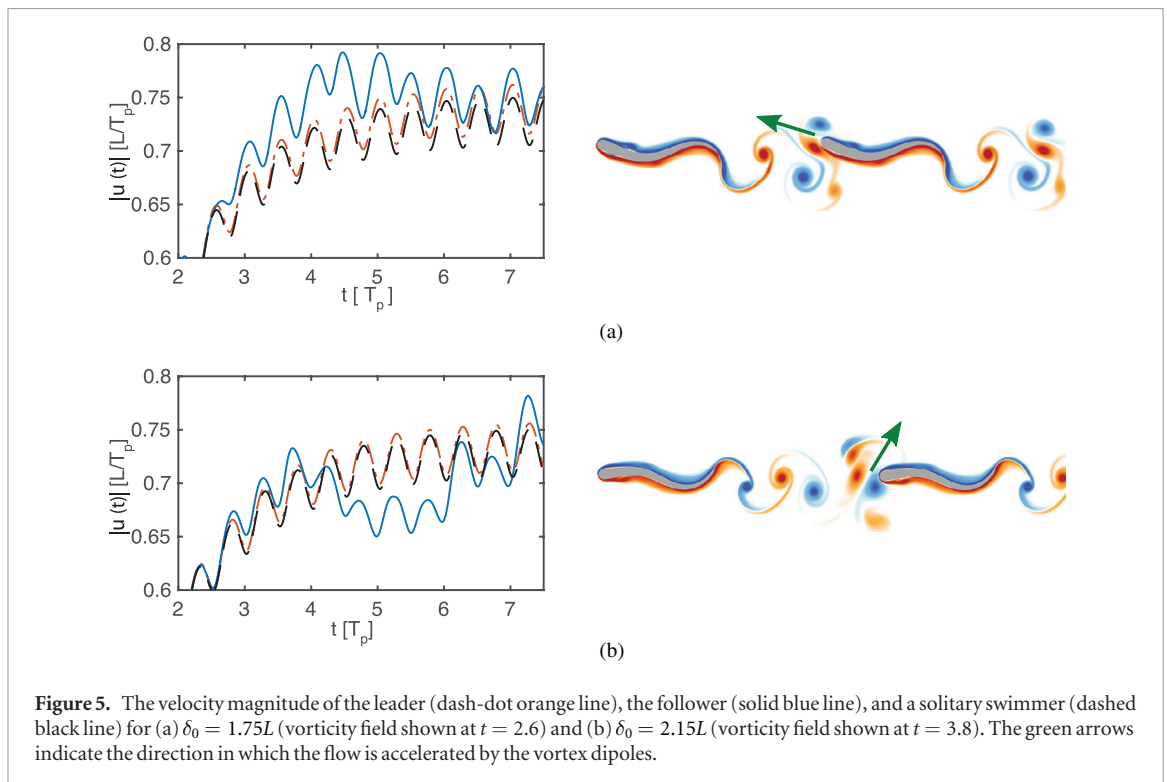
$$F_{\omega}^x = -\rho \frac{d}{dt} \left(\sum_i \gamma_i \Gamma_i \right) \quad (8)$$

This indicates that a drag-reduction (i.e. here a reduction in F_{ω}^x) can be obtained either by decreasing the circulation of the vortices (for example by producing dipole like structures) or by the advection of a positive/negative vortex respectively in the positive/negative y -direction. The first of these two possible scenarios is observed in the region highlighted in figure 4(a), where interaction with the wake-vortex decreases vorticity at the lower surface of the trailing airfoil.

The results suggest that hydrodynamic interactions between solid objects executing specific motion patterns can give rise to substantial drag-reduction, and even intermittent thrust production for the follower, while the leader remains largely unaffected. The forces experienced by the leader and follower in this setting are consistent with experimental and computational studies of tandem arrangements of cylinders in free flow [31, 42], where it has been observed that the follower can experience a substantial drag-reduction, and the leader is mostly unaffected by the presence of the follower.

3.2. Tandem of two-self propelled swimmers: no control

We examine the behaviour of a self-propelled swimmer placed initially in a tandem configuration with a leader. The two swimmers are positioned in a straight line, one directly behind the other, with both of them swimming in the same direction initially. The crucial difference from the configuration studied in section 3.1 is that both swimmers have a-priori defined sinusoidal body-deformations (section 2.1), and their trajectories are not imposed, but emerge from their interaction with the flow. The sinusoidal body undulations produce vorticity and impart momentum to the viscous fluid, which in turn modifies the pressure and viscous stress on the body. The resultant force-distribution on the surface of the object (similar to figures 4(c) and (d)) gives rise to the swimmer's rotational and translational motion. An extensive discussion of the relevant numerical algorithms and formulas used for this purpose may be found in [13]. The kinematics imposed for body-undulations are identical for both the leader and the follower (equation (2)), and correspond to a Reynolds number of $Re = L^2/Tp\nu = 5000$.




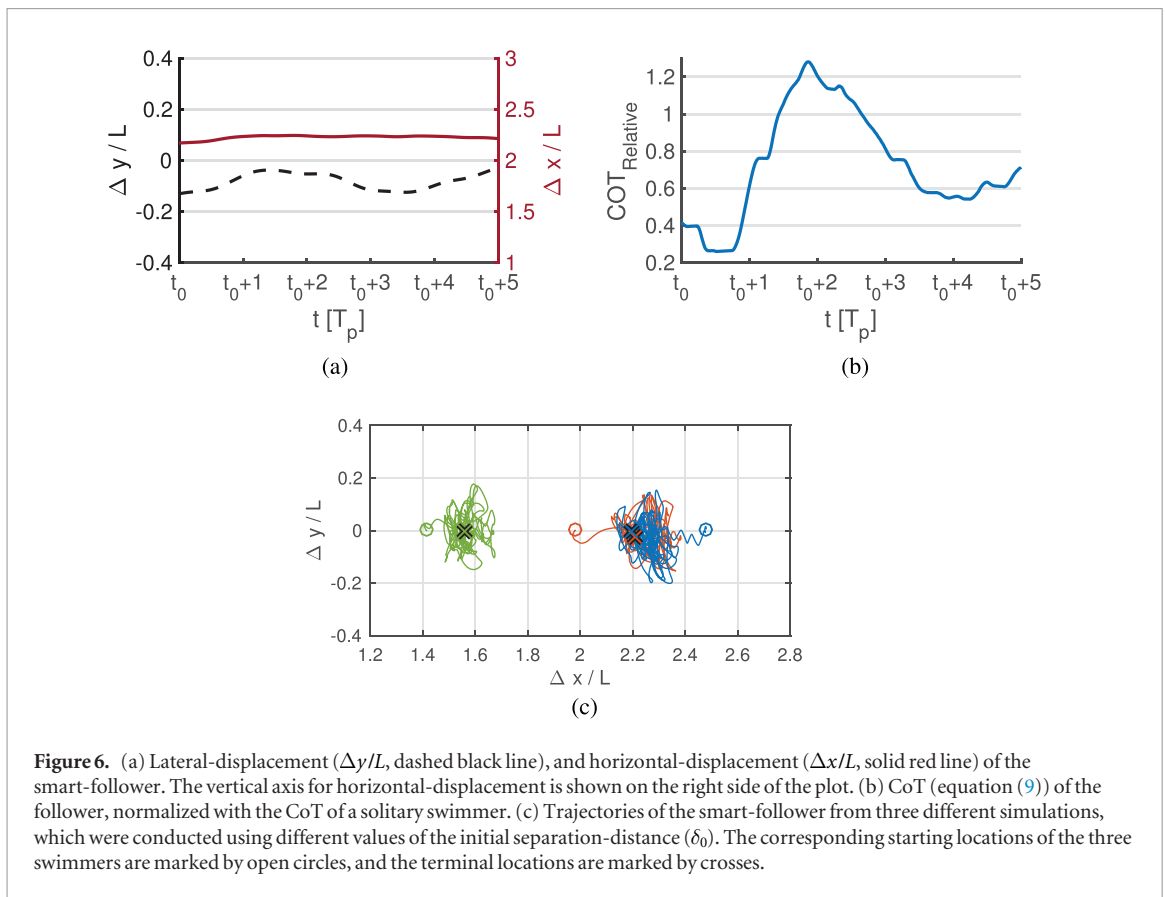
We consider two different cases, with the leader and the follower starting from rest at a separation distance of $\delta_0 = 1.75L$ with $\phi = 0$ (figure 5(a)), and at $\delta_0 = 2.15L$ with $\phi = \pi/2$ (figure 5(b)). The vorticity fields shown in both figures 5(a) and (b) correspond to instances when the follower first encounters the leader's wake. In the first case, the follower intercepts a pair of vortices which accelerate the flow in the direction of the swimmer's motion (marked by a green arrow in figure 5(a)). The reduction in relative velocity of the flow provides a drag-reduction resulting in a 9.5% increase of the follower's maximum speed (figure 5(a), $t \approx 4.4$). In the second case, the follower intercepts a pair of vortices which increases the relative velocity of the flow, and causes a large lateral deviation of the follower, with a subsequent speed-change of up to -9% (figure 5(b), $t \approx 4.9$). This vortex-induced acceleration or deceleration is equivalent, in a self-propelled setting, to the reduction in drag discussed in the previous section for the prescribed motion (please see figure 4). The Strouhal number for the fast follower is 0.23, and that for the slow follower is 0.31, based on the tail-beat period, maximum tail-beat amplitude, and the maximum/minimum speeds.

These results suggest that unsteady vortical structures in a leader's wake can have both a beneficial, as well as a detrimental impact on the performance of a follower. Furthermore, in both cases, the follower's trajectory starts deviating laterally as soon as it encounters the wake, and the follower is completely clear of the wake after approximately 4 to 6 tail-beat periods (supplementary movie 2 stacks.iop.org/BB/12/036001/mmedia). This suggests the need for active modulation of the trailing swimmer's actions when navigating a leader's wake, in order to maintain a tandem configuration.

3.3. Tandem of two-self propelled swimmers: adaptive control

In this section, we discuss the swimming-efficiency of a follower that adapts its motion using a RL algorithm, in response to velocity fluctuations in the leader's wake. The steady gait of the leader corresponds to a Reynolds number of $Re = L^2/T_p\nu = 5000$. In order to ensure adequate exploration of the state space, the follower initially performs random actions with a 50% probability, which is gradually reduced to 10%. After training, RL provides a control policy $\pi_{\Delta y}$, which is a deterministic mapping between the current state of the agent and the optimal action to correct its trajectory towards the center of the leader's wake ($\Delta y = 0$). For the present case, $\pi_{\Delta y}$ was determined using approximately 100 000 state-action-reward sets obtained during training. In the resulting simulations of the trained smart-swimmer, we do not observe a perfectly periodic sequence of actions, since the smart-swimmer can only choose among a finite set of corrective actions, and therefore tends to over/undershoot its goal. It maintains the schooling arrangement by course correcting towards $\Delta y = 0$, as can be observed from the time-evolution of Δy in figure 6(a), and supplementary movie 3 (stacks.iop.org/BB/12/036001/mmedia).

The time evolution of the horizontal distance, Δx in figure 6(a), indicates that the smart-swimmer prefers to hold position at $\Delta x/L \approx 2.2$. This seems to be a stable point, since the horizontal distance usually converges to this value, irrespective of the initial conditions used for the simulations (figure 6(c)). Moreover, the periodicity of the shed vortices (the wake-vortices are separated by a distance of $\pm 0.7L$) leads to an equivalent stable position for $\Delta x/L = 1.5$, as can be observed in figure 6(c). At these positions, the velocity of the follower's head

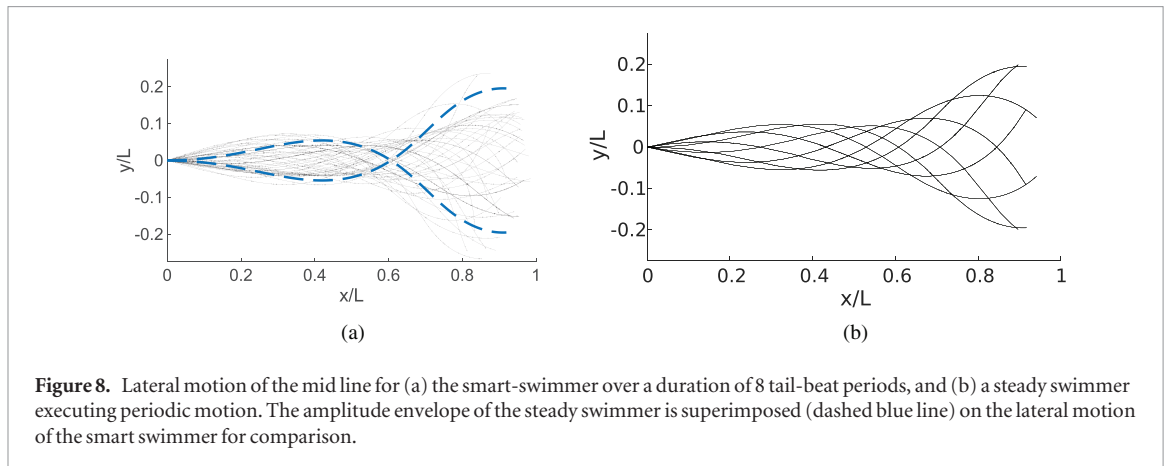


is synchronized with the fluid-velocity induced by the wake-vortices, which yields an energetic benefit by reducing the power required for body-deformation. Moreover, this motion pattern allows the swimmer to bring positive vorticity on the top and negative vorticity on the bottom of its body, thereby reducing drag in accordance with equation (8).

The wake-structures generated by the smart-follower and an uncontrolled fast swimmer are compared in figure 7. The primary difference observable is the existence of sustained flow pattern and alternating lateral jets in the case of the smart-follower, which arise

due to its systematic interception of the wake-vortices. There is no sustained wake pattern observed in the case of the uncontrolled follower and vortical structures are shed from different parts of the swimmer's body, depending on its instantaneous interactions with the leader's wake.

The Strouhal number for the smart-follower is 0.34, based on the fact that the tail-beat frequency is fixed, and the average swimming-speed remains approximately steady for maintaining the tandem configuration with the leader. The lateral-motion of the midline for the smart-follower is shown in figure 8, and is compared



to the midline kinematics of a steady-swimmer. Both these motions are comparable to the body-deformation observed in sub-carangiform swimmers [2, 6]. We observe that for the most part, the lateral motion of the smart-swimmer remains bounded by the maximum amplitude of the steady swimmer, except for a few occasions when the smart-swimmer executes a strong turn.

In order to make reliable conclusions on the energetic savings, we report quantities measured over at least 40 tail-beat periods, while showing a portion of this dataset (5 tail-beat periods from t_0 , which corresponds to the first frame of supplementary movie 3) in relevant figures, for the sake of clarity. To examine the impact that minimizing $|\Delta y|$ has on the energy consumption of the follower, we compute the ‘Cost of Transport’ (CoT) as follows:

$$\text{CoT}(t) = \frac{\int_{t-T_p}^t \max(P_{\text{def}}, 0) dt}{\int_{t-T_p}^t \|\mathbf{u}\| dt} \quad (9)$$

The CoT indicates the energy spent per unit distance travelled⁵. Figure 6(b) shows the CoT for the active trailing fish, normalized with respect to the CoT of an isolated swimmer (which remains approximately constant during steady-state swimming). The relative CoT tends to be smaller than 1 for the most part, which is indicative of the follower spending less energy in traversing a unit distance compared to an isolated swimmer. There are instances when the relative CoT exceeds 1, which corresponds to the follower exerting additional effort to execute corrective actions. Actively minimizing the lateral distance from the leader results in substantial energy savings per unit distance travelled; the smart-swimmer following $\pi_{\Delta y}$ requires 29.3% less energy than a solitary swimmer in an unperturbed flow. These results showcase the significant energetic benefits

⁵In this work we do not take into account the elastic response of the body to the fluid forces. Therefore, the \max operator in equation (9) precludes negative values of P_{def} , and accounts for the fact that the trailing swimmer may not elastically ‘store’ energy from the flow. Furthermore, taking only the positive part of P_{def} gives us a more conservative estimate of efficiency.

that can be obtained by a follower when exploiting a leader’s wake.

A reduction in power spent for deforming the body does not necessarily guarantee an increase in efficiency, as it may be accompanied by a reduction in thrust-power generated by the swimmer. Thus, we examine the efficiency η (equation (3)), the thrust-related power P_{thrust} , and power consumed by body-deformation P_{def} for a solitary swimmer adopting a steady gait, and the follower acting according to $\pi_{\Delta y}$ (figure 9). The active follower experiences an increase in swimming-efficiency (figure 9(a)). This points to an ability to extract energy from the oncoming vortices, and a consequential reduction in effort exerted by the swimmer. The net increase in average efficiency is approximately 19.4%. These gains do not arise due to an increase in P_{thrust} (figure 9(b)), but rather due to a reduction in P_{def} (figure 9(c)). The time-averaged P_{thrust} varies only moderately for the active swimmer compared to the solitary swimmer (−1.2%), but the reduction in P_{def} is substantial (−36.6%).

Vorticity contours corresponding to instances of maxima and minima in η are shown in figures 9(d) and (e), in order to examine the impact of the flow-field on instantaneous swimming-efficiency. At the instance of maximum swimming-efficiency (i.e. at $t = t_0 + 4.1$), the undulatory motion (i.e. deformation-velocity) of the swimmer is synchronized with the velocity induced by the wake-vortices, which minimizes the power required for deforming the body (negative P_{def} in figure 9(c)). At the instance of minimum swimming efficiency ($t = t_0 + 4.4$), the deformation-velocity of the swimmer is mostly out of phase with the flow-induced velocity, which results in large power requirement for body-deformation and decreases efficiency, even though the thrust-power is reasonably high.

The thrust- and deformation-power for the follower show a noticeable variation in amplitude, compared to those for a solitary swimmer (figure 9(b)). The fluctuating power-output is related to the fact that the distance-based reward $\mathcal{R}_{\Delta y}$ aims solely to minimize Δy . This can affect efficiency adversely for relatively short durations, as observed at times $(t_0 + 0.5)$, $(t_0 + 3.5)$ and $(t_0 + 4.5)$ in figure 9(a). Nonetheless, the smart-follower is more

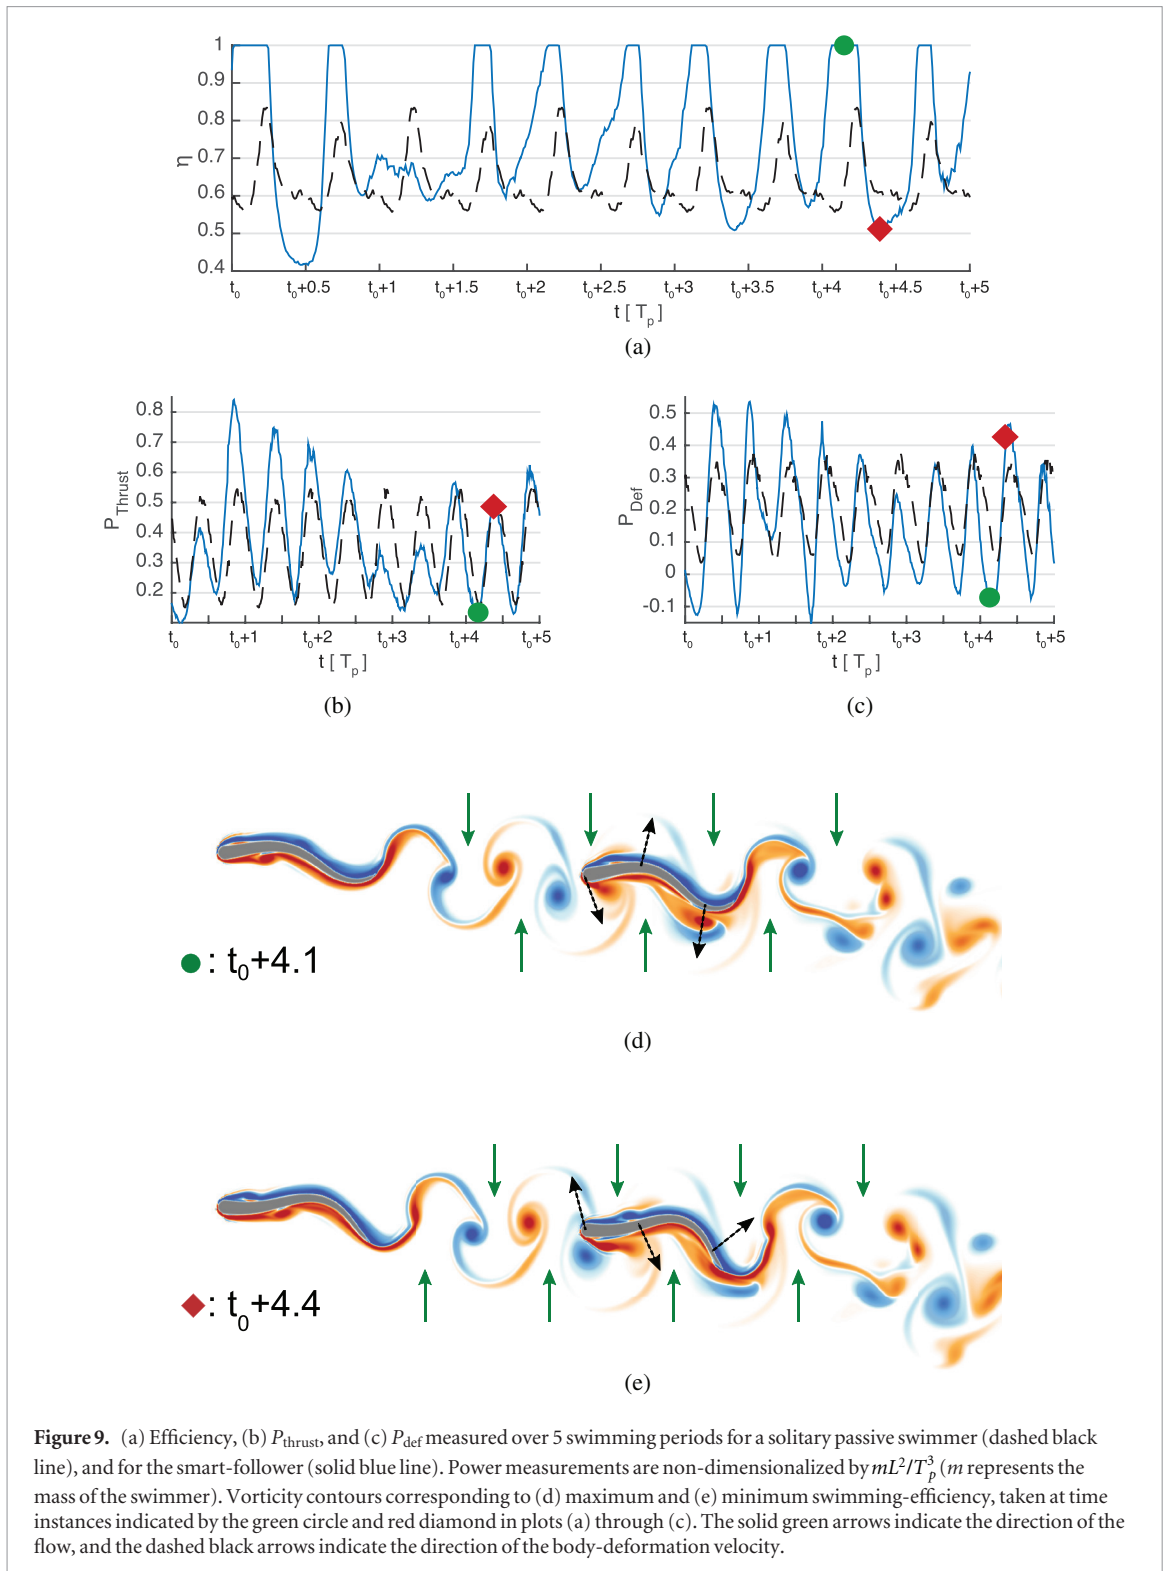


Figure 9. (a) Efficiency, (b) P_{thrust} , and (c) P_{def} measured over 5 swimming periods for a solitary passive swimmer (dashed black line), and for the smart-follower (solid blue line). Power measurements are non-dimensionalized by mL^2/T_p^3 (m represents the mass of the swimmer). Vorticity contours corresponding to (d) maximum and (e) minimum swimming-efficiency, taken at time instances indicated by the green circle and red diamond in plots (a) through (c). The solid green arrows indicate the direction of the flow, and the dashed black arrows indicate the direction of the body-deformation velocity.

efficient on average than a solitary swimmer (19.4% higher efficiency) while inducing only a small penalty to the swimming efficiency of the leader (3.3% lower efficiency). This indicates that the substantial reduction in energy consumption (29.3% drop in relative CoT) does not come at the expense of decreased efficiency.

The current simulations were conducted at relatively moderate Reynolds numbers, and are relevant for smaller fish species which exhibit a strong tendency for swimming in tandem (e.g. zebrafish). We now consider whether the present study is relevant for flows with

multiple swimmers, as is the situation encountered in large fish schools. We note that in fish schools, the vortical structures would be far more complex than for the case of two swimmers. However, we expect that at close distances (of approximately 2–3 body lengths), the vorticity shed by a leader will be the dominant feature encountered by a follower. For example, the work of Daghoobi & Boarjani [10] studies an ‘infinite school’ by considering the flow in a domain with periodic boundary conditions, such that the ‘leader’ encounters its own wake as a ‘follower’. In that study,

the vortical structures encountered by the swimmer are noisier than the ones observed for the two swimmers considered in the present study. Nonetheless, a coherent wake is observed near the swimmer even in that case (please see figure 6 in [10]). The existence of a coherent wake-structure has also been observed in other studies investigating multiple self-propelled swimmers (please see figures 18 and 20 in [13]). However, the coherence of wake-vortices may be disrupted more readily in three-dimensional cases as well as at higher Reynolds numbers, due to the occurrence of vortex-stretching. This will be a challenge for Reinforcement Learning, and is the subject of currently ongoing investigations.

4. Conclusion

In this paper, we demonstrate the energetic benefits of coordinated swimming for two swimmers in a leader-follower configuration, through a series of simulations. First, an arrangement of rigid airfoil-shaped swimmers, executing pre-specified motion, is observed to give rise to substantial drag-reduction. This simplified scenario demonstrates that interacting swimmers may see a benefit arising from vortices present in the flow even in the presence of large flow separation. Following this, we investigate self-propelled fish shapes, with both the leader and the follower employing identical kinematics. Without any active adaptation, the follower's interactions with the leader's wake can be either energetically beneficial or detrimental, depending on the initial condition. Furthermore, the follower tends to diverge from the leader's wake, which points to the need for active modulation of the follower's actions to maintain a stable tandem configuration. Finally, we examine the case where the leader swims with a steady gait and the follower adapts its behaviour dynamically to account for the effects of the wake encountered. The actions of the follower are selected autonomously from an optimal policy determined via reinforcement learning, and allow the swimmer to maximize a specified long-term reward. The results indicate that swimming in tandem can lead to measurable energy savings for the follower. We measure about 30% reduction in energy spent per unit distance, compared to a solitary swimmer, even when the goal of the follower is to minimize lateral distance from the leader. The results demonstrate that for two fish, swimming in a synchronised tandem configuration can give rise to substantial energetic benefits. We note that the training process of Reinforcement Learning is computationally expensive thus requiring large scale computational resources. We believe that two-dimensional simulations and low order models are essential in order to extend this work to three-dimensional flows. We also anticipate the possibility that swimming policies may be obtained that will lead to an open loop flow control for the swimmers. We envision that the findings of the present work will be relevant for the design of multiple robotic swimmers, that need to account for strong

hydrodynamic interactions in order to be energetically efficient.

Acknowledgment

This work utilized computational resources granted by the Swiss National Supercomputing Centre (CSCS) under project ID 's436'. We gratefully acknowledge support from the European Research Council (ERC) Advanced Investigator Award (No. 2-73985-14), and the Swiss National Science Foundation Sinergia Award (CRSII3 147675). WvR wishes to thank the Swiss National Science Foundation (SNF) for financial support.

References

- [1] Abrahams M V and Colgan P W 1985 Risk of predation, hydrodynamic efficiency and their influence on school structure *Environ. Biol. Fish.* **13** 195–202
- [2] Akanyeti O and Liao J C 2013 A kinematic model of kármán gaiting in rainbow trout *J. Exp. Biol.* **216** 4666–77
- [3] Becker A D, Masoud H, Newbolt J W, Shelley M and Ristroph L 2015 Hydrodynamic schooling of flapping swimmers *Nat. Commun.* **6** 8514
- [4] Bergmann M and Iollo A 2011 Modeling and simulation of fish-like swimming *J. Comput. Phys.* **230** 329–48
- [5] Borazjani I and Sotiropoulos F 2008 Numerical investigation of the hydrodynamics of carangiform swimming in the transitional and inertial flow regimes *J. Exp. Biol.* **211** 1541–58
- [6] Borazjani I and Sotiropoulos F 2009 On the role of form and kinematics on the hydrodynamics of self-propelled body/caudal fin swimming *J. Exp. Biol.* **213** 89–107
- [7] Breder C M 1965 Vortices and fish schools *Zoologica* **50** 97–114
- [8] Brock V E and Riffenburgh R H 1960 Fish schooling: a possible factor in reducing predation *ICES J. Mar. Sci.* **25** 307–17
- [9] Cushing D H and Harden Jones F R 1968 Why do fish school? *Nature* **218** 918–20
- [10] Daghooghi M and Borazjani I 2015 The hydrodynamic advantages of synchronized swimming in a rectangular pattern *Bioinspir. Biomim.* **10** 056018
- [11] Daghooghi M and Borazjani I 2016 Self-propelled swimming simulations of bio-inspired smart structures *Bioinspir. Biomim.* **11** 056001
- [12] Gautier N, Aider J-L, Duriez T, Noack B R, Segond M and Abel M 2015 Closed-loop separation control using machine learning *J. Fluid Mech.* **770** 442–57
- [13] Gazzola M, Chatelain P, van Rees W M and Koumoutsakos P 2011 Simulations of single and multiple swimmers with non-divergence free deforming geometries *J. Comput. Phys.* **230** 7093–114
- [14] Gazzola M, Hejazialhosseini B and Koumoutsakos P 2014 Reinforcement learning and wavelet adapted vortex methods for simulations of self-propelled swimmers *SIAM J. Sci. Comput.* **36** B622–B639
- [15] Gazzola M, Tchieu A A, Alexeev D, de Brauer A and Koumoutsakos P 2016 Learning to school in the presence of hydrodynamic interactions *J. Fluid Mech.* **789** 726–49
- [16] Gilmanov A and Sotiropoulos F 2005 A hybrid cartesian/immersed boundary method for simulating flows with 3d, geometrically complex, moving bodies *J. Comput. Phys.* **207** 457–92
- [17] Hemelrijk C K, Reid D A P, Hildenbrandt H and Padding J T 2015 The increased efficiency of fish swimming in a school *Fish Fish.* **16** 511–21
- [18] Henskin J and Steffensen J F 1998 Energy savings in sea bass swimming in a school: measurements of tail beat frequency and oxygen consumption at different swimming speeds *J. Fish Biol.* **53** 366–76

- [19] Kern S and Koumoutsakos P 2006 Simulations of optimized anguilliform swimming. *J. Exp. Biol.* **209** 4841–57
- [20] Killen S S, Marras S, Steffensen J F and McKenzie D J 2011 Aerobic capacity influences the spatial position of individuals within fish schools *Proc. R. Soc. B* **279** 357–64
- [21] Koumoutsakos P and Leonard A 1995 High-resolution simulations of the flow around an impulsively started cylinder using vortex methods *J. Fluid Mech.* **296** 1–38
- [22] Liao J C, Beal D N, Lauder G V and Triantafyllou M S 2003 Fish exploiting vortices decrease muscle activity *Science* **302** 1566–9
- [23] Mnih V et al 2015 Human-level control through deep reinforcement learning *Nature* **518** 529–33
- [24] Parker F R 1973 Reduced metabolic rates in fishes as a result of induced schooling *Trans. Am. Fish. Soc.* **102** 125–31
- [25] Partridge B L, Johansson J and Kalish J 1983 The structure of schools of giant bluefin tuna in Cape Cod Bay *Environ. Biol. Fish.* **9** 253–62
- [26] Partridge B L and Pitcher T J 1973 Evidence against a hydrodynamic function for fish schools *Nature* **279** 418–9
- [27] Pavlov D S and Kasumyan A O 2000 Patterns and mechanisms of schooling behavior in fish: a review *J. Ichthyol.* **40** 163–231
- [28] Pitcher T J, Magurran A E and Winfield I J 1982 Fish in larger shoals find food faster *Behav. Ecol. Sociobiol.* **10** 149–51
- [29] Portugal S J, Hubel T Y, Fritz J, Heese S, Trobe D, Voelkl B, Hailes S, Wilson A M and Usherwood J R 2014 Upwash exploitation and downwash avoidance by flap phasing in ibis formation flight *Nature* **505** 399–402
- [30] Rossinelli D, Hejazialhosseini B, van Rees W M, Gazzola M, Bergdorf M and Koumoutsakos P 2015 MRAG-I2D: Multi-resolution adapted grids for remeshed vortex methods on multicore architectures *J. Comput. Phys.* **288** 1–18
- [31] Sumner D 2010 Two circular cylinders in cross-flow: a review *J. Fluids Struct.* **26** 849–99
- [32] Sutton R S and Barto A G 1998 *Reinforcement Learning: an Introduction* (Cambridge, MA: MIT Press)
- [33] Svendsen J C, Skov J, Bildsoe M and Steffensen J F 2003 Intra-school positional preference and reduced tail beat frequency in trailing positions in schooling roach under experimental conditions *J. Fish Biol.* **62** 834–46
- [34] Taguchi M and Liao J C 2011 Rainbow trout consume less oxygen in turbulence: the energetics of swimming behaviors at different speeds *J. Exp. Biol.* **214** 1428–36
- [35] Triantafyllou M S, Weymouth G D and Miao J M 2016 Biomimetic survival hydrodynamics and flow sensing *Annu. Rev. Fluid Mech.* **48** 1–24
- [36] Tsang A C H and Kanso E 2013 Dipole interactions in doubly periodic domains *J. Nonlinear Sci.* **23** 971–91
- [37] Tytell E D and Lauder G V 2004 The hydrodynamics of eel swimming *J. Exp. Biol.* **207** 1825–41
- [38] van Rees W M, Novati G and Koumoutsakos P 2015 Self-propulsion of a counter-rotating cylinder pair in a viscous fluid *Phys. Fluids* **27** 063102
- [39] Weihs D 1973 *Nature* **241** 290–1
- [40] Weihs D 1975 Some hydrodynamical aspects of fish schooling *Swimming and Flying in Nature* vol 2 (Boston, MA: Springer) pp 703–18
- [41] Whittlesey R W, Liska S and Dabiri J O 2010 Fish schooling as a basis for vertical axis wind turbine farm design *Bioinspir. Biomim.* **5** 035005
- [42] Zdravkovich M M 1977 Review of flow interference between two circular cylinders in various arrangements *ASME Trans. J. Fluids Eng.* **99** 618–33