

Synchronised Swimming of Two Fish

Guido Novati¹, Siddhartha Verma¹, Dmitry Alexeev¹, Diego Rossinelli¹, Wim M. van Rees^{1,2} and Petros Koumoutsakos^{1†}

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

²School of Engineering and Applied Sciences, Harvard University, USA

We study the fluid dynamics of two fish-like bodies with synchronised swimming patterns. Our studies are based on two-dimensional simulations of viscous incompressible flows. We distinguish between motion patterns that are externally imposed on the swimmers and self-propelled swimmers that learn manoeuvres to achieve certain goals. Simulations of two rigid bodies executing pre-specified motion indicate that flow-mediated interactions can lead to substantial drag reduction and may even generate thrust intermittently. In turn we examine two self-propelled swimmers arranged in a leader-follower configuration, with a-priori specified body-deformations. We find that the swimming of the leader remains largely unaffected, while the follower experiences either an increase or decrease in swimming speed, depending on the initial conditions. Finally, we consider a follower that synchronises its motion so as to minimize its lateral deviations from the leader's path. The leader employs a steady gait while the follower uses a reinforcement learning algorithm to adapt its swimming-kinematics. We find that swimming in a synchronised tandem can yield up to about 30% reduction in energy expenditure for the follower, in addition to a 20% increase in its swimming-efficiency. The present results indicate that synchronised swimming of two fish can be energetically beneficial.

Key words: Swimming/flying, Propulsion, Vortex interactions, Wakes

1. Introduction

The coordinated motion of fish is thought to provide an energetic advantage to individuals, as well as to the group as a whole, in terms of increased swimming range, endurance and chances of survival. Schooling has been credited with serving various other biological functions including defence from predators (Brock & Riffenburgh 1960; Cushing & Harden Jones 1968), enhanced feeding and reproductive opportunities (Shaw 1978; Pitcher *et al.* 1982). Amid a lack of clear agreement regarding the evolutionary purpose of schooling behaviour (Pavlov & Kasumyan 2000), there is growing evidence supporting the role of hydrodynamic mediation in facilitating propulsion. Experiments investigating fish swimming in groups point to a reduction in energy expenditure, based on respirometer readings and reduced tail-beat frequency (Parker 1973; Abrahams & Colgan 1985; Herskin & Steffensen 1998; Svendsen *et al.* 2003; Killen *et al.* 2011). Importantly, there is evidence to suggest that reduction in energy expenditure is not distributed uniformly throughout a schooling group. Herskin & Steffensen (1998); Svendsen *et al.* (2003); Killen *et al.* (2011) observed that the tail-beat frequency of trailing fish was lower than that of

† Email address for correspondence: petros@ethz.ch

fish at the front of the school. Moreover, Killen *et al.* (2011) 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 (Liao *et al.* 2003) found that trout swimming behind obstacles exerted reduced effort for station-keeping. The trout adopted a modified gait, which allowed them to ‘slalom’ through the oncoming vortex street. The ensuing reduction in muscle activity was confirmed using neuromuscular measurements (Liao *et al.* 2003) and respirometer readings (Taguchi & Liao 2011). These experimental studies suggest that fish can detect and exploit vortices present in their surroundings.

There is a well documented hypothesis (Breder 1965; Weihs 1973, 1975; Shaw 1978) that flow patterns which emerge as a consequence of schooling, can be exploited by individual swimmers. This hypothesis was first quantified (Weihs 1973, 1975) by 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 (Weihs 1975). The role of hydrodynamics in fish-schooling was later questioned (Partridge & Pitcher 1979), based on empirical observations of fish-schools which rarely displayed diamond formations. However, a later study based on aerial photographs of hunting tuna schools (Partridge *et al.* 1983) 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. Theoretical and numerical studies of schooling have employed potential flow models (Tsang & Kanso 2013; Gazzola *et al.* 2016), or they have pre-specified the spatial distribution and motion of the swimmers (Hemelrijk *et al.* 2015; Daghooghi & Borazjani 2015).

Here we present two-dimensional simulations of viscous, incompressible flows of multiple self-propelled swimmers that can dynamically adapt their motion. We focus in particular on two swimmers arranged in a leader-follower configuration. 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 of the three synchronised swimming scenarios in Section 3, followed by concluding remarks in Section 4.

2. Simulation details

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 (Rossinelli *et al.* 2015), and a divergence-free penalisation technique to enforce

the no-slip boundary condition (Gazzola *et al.* 2011). 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 self-propelled swimmers used in the simulations are based on a simplified physical model of zebrafish as described in Gazzola *et al.* (2011). 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:

$$k(s, t) = A(s) \sin \left[2\pi \left(\frac{t}{T_p} - \frac{s}{L} \right) + \phi \right] \quad (2.1)$$

Here, L is the length of the swimmer, $T_p = 1$ is the tail-beat frequency, ϕ is the phase-difference. The curvature amplitude $A(s)$ varies linearly from the head to the tail thus reducing head motion and amplifying tail-beat amplitude. We estimate the swimming-efficiency using a modified form of the Froude efficiency proposed by Tytell & Lauder (2004):

$$\eta = \frac{P_{thrust}}{P_{thrust} + \max(P_{def}, 0)} = \frac{Tu}{Tu + \max(-\iint \mathbf{u}_{def} \cdot d\mathbf{F}, 0)} \quad (2.2)$$

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 $T = \iint (\mathbf{u} \cdot d\mathbf{F} + |\mathbf{u} \cdot d\mathbf{F}|)/2$, 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:

$$d\mathbf{F}_V = 2\mu \mathbf{D} \cdot \mathbf{n} dS \quad \text{and} \quad d\mathbf{F}_P = -P\mathbf{n} dS \quad (2.3)$$

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 surface-pressure is obtained by solving a Poisson's equation ($\nabla^2 P = -\rho (\nabla \mathbf{u}^T : \nabla \mathbf{u}) + \rho \lambda \nabla \cdot (\chi(\mathbf{u}_s - \mathbf{u}))$) (Verma *et al.* submitted).

2.2. Reinforcement learning

The two swimmers either follow a-priori defined swimming patterns, or the follower adapts its body-deformation to synchronise its motion with that of the leader. This adaptation is achieved using a Reinforcement Learning (RL) (Sutton & Barto 1998), a potent machine learning algorithm for model-free flow control (Gautier *et al.* 2015; Gazzola *et al.* 2016). In RL, the agents receive information about their *State* and chose *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 $r(s, a, s')$, 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 $\pi(s')$:

$$Q_\pi(s, a) = \mathbb{E} [r(s, a, s') + \gamma Q(s', \pi(s'))] \quad (2.4)$$

The discount factor $\gamma \in [0, 1]$ determines the trade-off between immediate and future rewards and was set to $\gamma = 0.8$ for all present results. The learning process terminates upon convergence of $Q(s, a)$, and the swimmer can make optimal decisions by following a 'greedy' policy ($\pi(s) = \arg \max_a Q_\pi(s, a)$). The result of this process is usually a tabular approximation of state-action-reward values. An important note for the present study is




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 Eq. 2.5, and the blue dash-dot line indicates the unmodified curvature (i.e., the sinusoidal part of Eq. 2.1). (b) The impact of varying the control-amplitude on the modified shape. The unmodified shape corresponds to $b = 0$.




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

the realisation that the motion of the swimmer implies a continuous state space. Hence, in this work the swimmer learns a parametrized approximation of $Q(s, a)$ by training a Neural Network (NN) with experience replay (Lin 1992). The algorithm involves storing all observed transitions $\{s, a, s', r\}$, which are sampled to update 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 Eq. 2.1:

$$k_{\text{learner}}(s, t) = k(s, t) + k'(s, t) = A(s) \left(\sin \left[2\pi \left(\frac{t}{T_p} - \frac{s}{L} \right) + \phi \right] + M \left(\frac{t}{T_p} - \frac{s}{L} \right) \right) \quad (2.5)$$

where $M \left(\frac{t}{T_p} - \frac{s}{L} \right)$ defines a travelling natural cubic spline, computed using three evenly-spaced nodes separated by a distance of $L/4$ (figure 1a). The spline shape remains constant as it proceeds from the head to the tail, and is determined by a fixed control-amplitude ‘ b ’ associated with a particular action. The swimmers are permitted to execute an action every half tail-beat period ($T_p/2$), which allows them to either increase or reduce the undulation amplitude (figures 1a and 1b).

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 also considers whether it is taking an action in the first or second half of its tail-beat period, as part of its state (modulo(time, 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 1a). 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 laterally strays too far 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

We distinguish externally imposed motions on the swimmers to those that are achieved by the deformation of the body of self-propelled swimmers. 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, are discussed in this section.

3.1. Rigid objects with pre-specified motion

A swimming pattern often observed in schooling involves the exchange of the leader and follower positions. We examine this scenario by first studying two rigid, airfoil-shaped bodies (shape identical to swimmers) dragged along prescribed sinusoidal paths to exchange their position as leader and follower (see Movie 1):

$$a_x(t) = \begin{cases} -\frac{u_{max}-u_{min}}{T_s} & \text{for } 2nT_s \leq t < (2nT_s + 1) \\ \frac{u_{max}-u_{min}}{T_s} & \text{for } (2nT_s + 1) \leq t < 2n(T_s + 1) \end{cases} \quad (3.1)$$

Here, n is an integer, $u_{max} = 4.5L/T_s$, $u_{min} = 1.5L/T_s$, 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 (Eq 3.1), whereas the other object starts with u_{min} on a trough and is subjected to constant acceleration. This arrangement of positions and velocities alternates between the two airfoils every time-period T_s .

A snapshot of the resulting 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. We observe that the follower in figure 3a experiences a dramatic reduction in drag, as indicated by the dash-dotted line in figure 3b ($t \approx 3.9$). This can be attributed to a decrease in relative velocity, due to the presence of the positive vortex highlighted in figure 3a. 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 3b indicates that both the leader and the follower may experience a reduction in drag as a result of mutual interaction. 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.

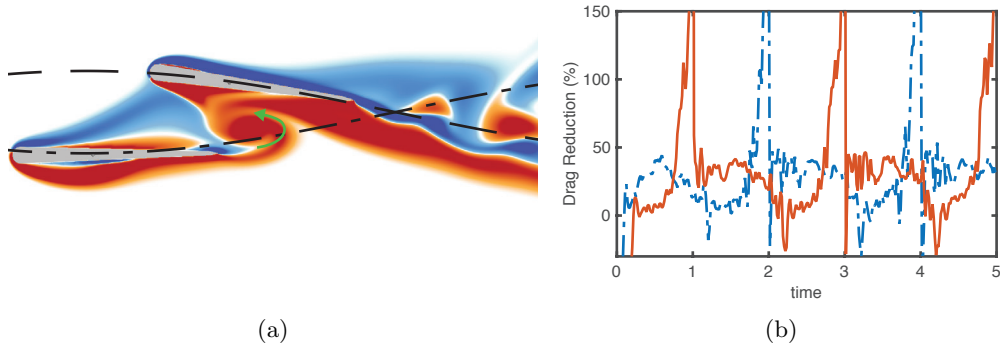


Figure 3: (a) Vorticity field generated by the two fish-shaped airfoils dragged with the prescribed sinusoidal pattern ($Re = u_{max}L/\nu = 2250$). The snapshot shown corresponds to $t = 3.88$. (b) Percentage of drag reduction for the two airfoils, with respect to a single airfoil executing the same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

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. Both swimmers have a-priori defined body-deformations. Their motions are determined by their deformations and, in particular for the follower, by their interaction with the flow field. The kinematics imposed for body-undulations are identical for both the leader and the follower (Eq. 2.1), and correspond to a Reynolds number of $Re = L^2/T_P\nu = 5000$. The crucial difference from the configuration studied in Section 3.1 is that both the swimmers are now self-propelled (Section 2.1), and their trajectories are not imposed, but emerge from their interaction with the flow.

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 4a), and at $\delta_0 = 2.15L$ with $\phi = \pi/2$ (figure 4b). The vorticity fields shown in both figures 4a and 4b correspond to instances when the follower first encounters the leader's wake. We observe an increase of 9.5% in the follower's maximum velocity in the first case (figure 4a, $t \approx 4.4$), whereas the follower in the second case experiences a velocity reduction of up to -9.1% (figure 4b, $t \approx 4.9$).

These results suggest that hydrodynamic interaction with 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). This suggests the need for active modulation of the trailing swimmer's actions when navigating the wake of a leader, 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%. The policy $\pi_{\Delta y}$ is determined using approximately 100,000 state-actions-reward sets.

The follower actively seeks to maintain its position in the center of the leader's wake, as can be observed from the time-evolution of Δy in figure 5a (and supplementary Movie 3).

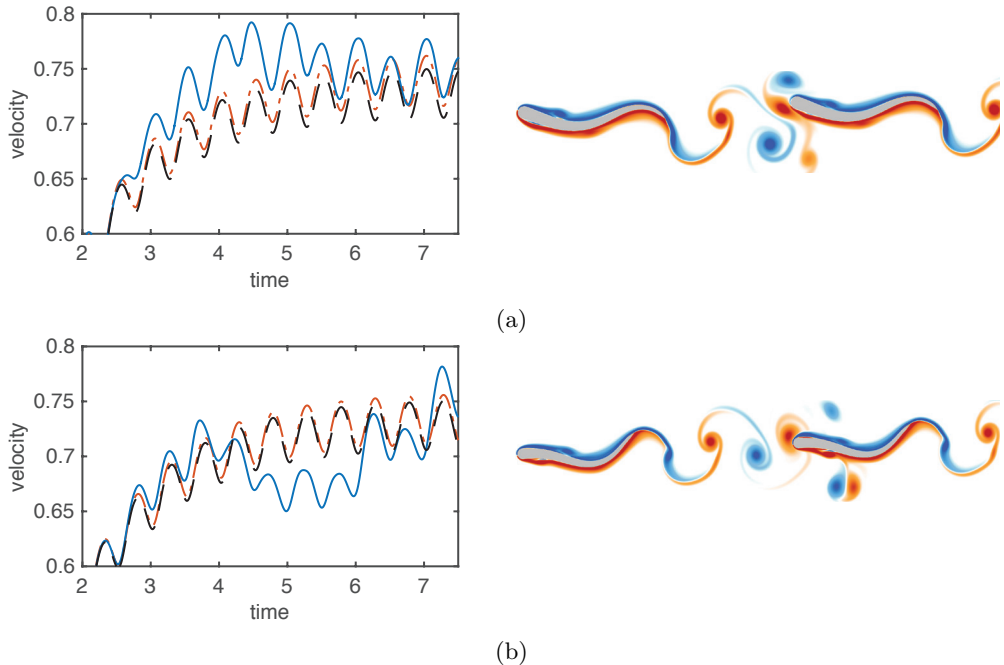


Figure 4: The velocity of the leader (dash-dot orange line), the follower (solid blue line), and a solitary swimmer (dashed black line) for (a) $\delta_0 = 1.75L$ (with the vorticity field shown at $t = 2.6$) and (b) $\delta_0 = 2.15L$ (vorticity field shown at $t = 3.6$).




Figure 5: (a) Lateral displacement, and (b) CoT (equation 3.2) of the follower, normalized with the CoT of a solitary swimmer.

Without active control, the follower may get deflected away from the wake, as discussed in section 3.2. 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:

$$\text{CoT}(t) = \frac{\int_{t-T_p}^t \max(P_{def}, 0) dt}{\int_{t-T_p}^t \|\mathbf{u}\| dt} \quad (3.2)$$

The CoT indicates the energy spent per unit distance travelled. We remark that the max operator in Eq. 3.2 precludes negative values of P_{Def} , and accounts for the fact that the trailing swimmer may not ‘store’ energy from the flow. Figure 5b 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

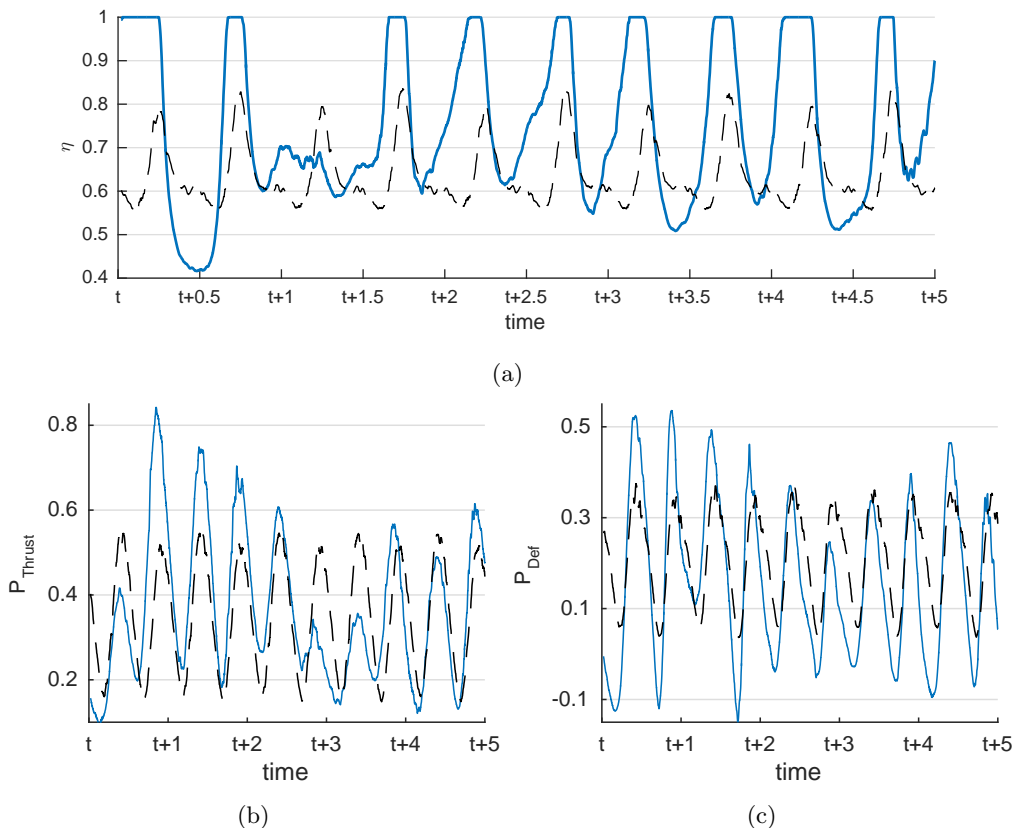


Figure 6: (a) Efficiency, (b) P_{thrust} , and (c) P_{def} measured over 5 swimming periods for a solitary passive swimmer (dashed black line), and for the follower acting according to $\pi_{\Delta y}$ (solid blue line). Power measurements are non-dimensionalized by mL^2/T_p^3 (m represents the mass of the swimmer).

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. Over a course of 40 time-periods, actively minimizing the lateral distance from the leader results in substantial energy savings per unit distance travelled; the 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 that can be obtained by a follower when exploiting a leader's wake.

In addition to the CoT, we examine the efficiency η (Eq. 2.2), 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 6). The active follower experiences an increase in swimming-efficiency (figure 6a). 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 over a duration of 40 time-periods is approximately 19.4%. These gains do not arise due to an increase in P_{Thrust} (figure 6b), but rather due to a reduction in P_{Def} (figure 6c). The time-averaged P_{Thrust} over 40 time-periods varies only moderately for the active swimmer compared to the solitary swimmer (-1.2%), but the reduction in P_{Def} is substantial (-36.6%). This may be

related to the fact that the follower tends to favour a decrease, rather than an increase, in the undulation-amplitude (section 2.2).

The thrust- and deformation-power for the follower show a noticeable variation in amplitude, compared to those for a solitary swimmer (figure 6b). The fluctuating power-output is related to the fact that the distance-based reward $\mathcal{R}_{\Delta y}$ solely aims to minimize Δy . This can affect efficiency adversely for relatively short durations, as observed at times $(t + 0.5)$, $(t + 3.5)$ and $(t + 4.5)$ in figure 6a. Nonetheless, the ‘smart’ follower is more efficient on average than a solitary swimmer (19.4% higher 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.

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. 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, for the first time to the best of our knowledge, 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.

Current work extends these simulations to three-dimensional flows by deploying wavelet adapted vortex to massively parallel computing architectures. We envision that this work will be important for the design of energetically efficient robotic devices that need to account for strong hydrodynamic interactions.

Acknowledgement

This work utilized computational resources granted by the Swiss National Supercomputing Centre (CSCS) under project ID ‘s436’. We gratefully acknowledge support by the European Research Council Advanced Investigator Award.

REFERENCES

- ABRAHAMS, M. V. & COLGAN, P. W. 1985 Risk of predation, hydrodynamic efficiency and their influence on school structure. *Environ. Biol. Fish.* **13** (3), 195–202.
- BREDER, C. M. 1965 Vortices and fish schools. *Zoologica-N.Y.* **50** (2), 97–114.
- BROCK, V. E. & RIFFENBURGH, R. H. 1960 Fish schooling: A possible factor in reducing predation. *ICES J. Mar. Sci.* **25** (3), 307–317.

- CUSHING, D. H. & HARDEN JONES, F. R. 1968 Why do fish school? *Nature* **218**, 918–920.
- DAGHOOGHI, M. & BORAZJANI, I. 2015 The hydrodynamic advantages of synchronized swimming in a rectangular pattern. *Bioinspir. Biomim.* **10** (5), 056018.
- GAUTIER, N., AIDER, J.-L., DURIEZ, T., NOACK, B. R., SEGOND, M. & ABEL, M. 2015 Closed-loop separation control using machine learning. *J. Fluid Mech.* **770**, 442–457.
- GAZZOLA, M., CHATELAIN, P., VAN REES, W. M. & KOUMOUTSAKOS, P. 2011 Simulations of single and multiple swimmers with non-divergence free deforming geometries. *J. Comput. Phys.* **230**, 7093–7114.
- GAZZOLA, M., TCHIEU, A. A., ALEXEEV, D., DE BRAUER, A. & KOUMOUTSAKOS, P. 2016 Learning to school in the presence of hydrodynamic interactions. *J. Fluid Mech.* **789**, 726–749.
- HEMELRIJK, C. K., REID, D. A. P., HILDENBRANDT, H. & PADDING, J. T. 2015 The increased efficiency of fish swimming in a school. *Fish Fish.* **16** (3), 511–521.
- HERSKIN, J. & 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** (2), 366–376.
- KILLEN, S. S., MARRAS, S., STEFFENSEN, J. F. & MCKENZIE, D. J. 2011 Aerobic capacity influences the spatial position of individuals within fish schools. *P. Roy. Soc. Lond. B Bio.* **279** (1727), 357–364.
- LIAO, J. C., BEAL, D. N., LAUDER, G. V. & TRIANTAFYLLOU, M. S. 2003 Fish exploiting vortices decrease muscle activity. *Science* **302**, 1566–1569.
- LIN, L. 1992 Self-improving reactive agents based on reinforcement learning, planning and teaching. *Mach. Learn.* **8** (3-4), 293–321.
- PARKER, F. R. 1973 Reduced metabolic rates in fishes as a result of induced schooling. *T. Am. Fish. Soc.* **102** (1), 125–131.
- PARTRIDGE, B. L., JOHANSSON, J. & KALISH, J. 1983 The structure of schools of giant bluefin tuna in Cape Cod Bay. *Environ. Biol. Fish.* **9** (3), 253–262.
- PARTRIDGE, B. L. & PITCHER, T. J. 1979 Evidence against a hydrodynamic function for fish schools. *Nature* **279**, 418–419.
- PAVLOV, D. S. & KASUMYAN, A. O. 2000 Patterns and mechanisms of schooling behavior in fish: A review. *J. Ichthyol.* **40** (2), 163–231.
- PITCHER, T. J., MAGURRAN, A. E. & WINFIELD, I. J. 1982 Fish in larger shoals find food faster. *Behav. Ecol. Sociobiol.* **10** (2), 149–151.
- ROSSINELLI, D., HEJAZIALHOSSEINI, B., VAN REES, W. M., GAZZOLA, M., BERGDORF, M. & KOUMOUTSAKOS, P. 2015 MRAG-I2D: Multi-resolution adapted grids for remeshed vortex methods on multicore architectures. *J. Comput. Phys.* **288**, 1–18.
- SHAW, E. 1978 Schooling Fishes: The school, a truly egalitarian form of organization in which all members of the group are alike in influence, offers substantial benefits to its participants. *Am. Sci.* **66** (2), 166–175.
- SUTTON, R. S. & BARTO, A. G. 1998 *Reinforcement learning: An introduction*. MIT press.
- SVENDSEN, J. C., SKOV, J., BILDSE, M. & 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** (4), 834–846.
- TAGUCHI, M. & LIAO, J. C. 2011 Rainbow trout consume less oxygen in turbulence: the energetics of swimming behaviors at different speeds. *J. Exp. Biol.* **214** (9), 1428–1436.
- TSANG, A. C. H. & KANSO, E. 2013 Dipole interactions in doubly periodic domains. *J. Nonlinear Sci.* **23** (6), 971–991.
- TYTELL, E. D. & LAUDER, G. V. 2004 The hydrodynamics of eel swimming. *J. Exp. Biol.* **207** (11), 1825–1841.
- VERMA, S., ABBATI, G., NOVATI, G. & KOUMOUTSAKOS, P. submitted Computing the force distribution on the surface of complex, deforming geometries using vortex methods and brinkman penalization. *Submitted*.
- WEIHS, D. 1973 Hydromechanics of fish schooling. *Nature* **241**, 290–291.
- WEIHS, D. 1975 *Swimming and Flying in Nature: Volume 2*, chap. Some Hydrodynamical Aspects of Fish Schooling, pp. 703–718. Boston, MA: Springer US.