SIAM J. SCI. COMPUT.
Vol. 36, No. 3, pp. B622–B639
© 2014 Society for Industrial and Applied Mathematics

Downloaded 06/28/14 to 128.103.224.4. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

REINFORCEMENT LEARNING AND WAVELET ADAPTED VORTEX METHODS FOR SIMULATIONS OF SELF-PROPELLED SWIMMERS*

MATTIA GAZZOLA†,‡, BABAK HEJAZIALHOSSEINI‡, AND PETROS KOUMOUTSAKOS‡

Abstract. We present a numerical method for the simulation of collective hydrodynamics in self-propelled swimmers. Swimmers in a viscous incompressible flow are simulated with a remeshed vortex method coupled with Brinkman penalization and projection approach. The remeshed vortex methods are enhanced via wavelet based adaptivity in space and time. The method is validated on benchmark swimming problems. Furthermore the flow solver is integrated with a reinforcement learning algorithm, such that swimmers can learn to adapt their motion so as to optimally achieve a specified goal, such as fish schooling. The computational efficiency of the wavelet adapted remeshed vortex method is a key aspect for the effective coupling with learning algorithms. The suitability of this approach for the identification of swimming behaviors is assessed on a set of learning tasks.

Key words. vortex methods, reinforcement learning, fish schooling, wavelet adapted grids

AMS subject classifications. 76Z10, 74F10, 68T05, 76M23

DOI. 10.1137/130943078

1. Introduction.

The fascinating complexity of schooling dynamics exhibited by a wide variety of swimmers across taxa has been the subject of intense research efforts in recent years. Experimental and theoretical studies have suggested a number of potential social benefits that may arise from such coordinated behaviors, ranging from enhanced protection [1, 2, 3, 4, 5, 6] to exploration capabilities for foraging [7, 8, 9].

Furthermore, it has been argued that collective hydrodynamics phenomena may correspond to energy saving mechanisms, increasing (in average) the swimming efficiency of school members [10, 11, 12, 13, 14]. For example, vortical structures generated by the unsteady motions of leading swimmers and shed into their wakes, may be exploited by trailing fish to increase their locomotion efficiency [15]. The role of vorticity in fish schooling, however, remains largely unexplored due to difficulties in obtaining quantitative information from experiments and in devising simulations that can handle effectively multiple self-propelled swimmers.

The work of Weihs [13] suggested the existence of optimally efficient schooling arrangements based on analytical considerations derived from a two-dimensional inviscid model. In turn the identification of such optimal lattice formations for swimmers may be exploited for the design and arrangement of energy harnessing devices [16, 17]. More recently, field studies of human-led migratory flights of birds have revealed that individuals flying in a V-flock position themselves in an energetically favorable arrangement [18].

A number of inviscid flow models have led the way in studying hydrodynamic interactions between model swimmers [19, 20, 21]. Inviscid flow solvers are attractive as they can provide enhanced understanding and they are computationally inexpensive

*Submitted to the journal's Computational Methods in Science and Engineering section October 28, 2013; accepted for publication (in revised form) February 10, 2014; published electronically June 26, 2014.

http://www.siam.org/journals/sisc/36-3/94307.html

†School of Engineering and Applied Sciences, Harvard University (mgazzola@seas.harvard.edu).

‡Chair of Computational Science, ETH Zurich, CH-8092 Zurich, Switzerland (hbabak@mavt.ethz.ch, petros@ethz.ch).

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 06/28/14 to 128.103.224.4. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

so they can be used to investigate different schooling scenarios. However, inviscid models do not capture processes such as flow separation and dynamic interactions between vortices in fish wakes. Viscous effects have been shown [22] to dramatically affect not only the quantitative but also the qualitative response of a system of bodies interacting in a flow field.

Several works have simulated in recent years individual swimmers. However only a limited number of studies have focused on simulating multiple self-propelled swimmers. Individual tethered swimmers have been simulated in three dimensions through immersed boundary methods [23, 24, 25] and finite differences [26]. The more challenging case of individual self-propelled swimmers has been tackled in two dimensions via finite differences [27, 28] and vortex methods [29, 30], while three-dimensional simulations have been performed using immersed boundary methods [31], Lagrangian multipliers [32], finite volumes [33], remeshed smoothed particle hydrodynamics [34], and remeshed vortex methods [30, 35, 36]. The case of multiple self-propelled bodies was simulated in two dimensions by Bergmann and Iollo via finite differences for a pair of fish, while remeshed vortex methods [37] have been used to simulate groups of fish including up to five members [30].

In addition to the computational challenges inherent in simulating collective hydrodynamics, a critical issue in examining schooling is the implementation of a behavioral model which allows individual swimmers to react to the unsteady flows generated by themselves and their neighbors. Theoretical studies have shown that collective behaviors emerge from simple, local interaction rules, such as alignment or repulsion among neighboring individuals [38, 39, 40]. These studies have focused on handcrafting a priori rules, based on experimental observations [41, 42]. Such interaction rules when applied to agent swimmers without accounting for hydrodynamics reproduce patterns observed in nature [43, 44, 45]. However, it is unclear how these interaction rules can accommodate the feedback mechanisms of the flow fields inherent to multiple interacting swimmers. In the context of group swimming in a viscous flow environment, behavioral policies cannot be a priori defined due to the highly nonlinear nature and unsteadiness of flow mediated interactions.

In this work, in order to develop simulations of fish schooling, the behavioral policies of multiple self-propelled two-dimensional swimmers are identified by coupling a reinforcement learning (RL) technique to a multiresolution remeshed vortex method with Brinkman penalization and projection approach [30, 22]. To the best of our current knowledge, this coupling of RL algorithms and simulations of the Navier–Stokes equations has never been attempted before. Here we demonstrate its suitability for the characterization of schooling dynamics by presenting the single and multiple swimmers with several learning tasks.

The paper is organized as follows. The geometrical model, the governing equations, the flow solver, and the RL approach are presented in section 2. The proposed numerical scheme is validated and its convergence properties are assessed in section 3. Learning applications are reported in section 4, and our findings are summarized and discussed in section 5.

2. Methodology.

2.1. Geometrical model of swimmers. We consider geometrically identical self-propelled anguilliform swimmers of length L. The swimmers' geometry is captured by a characteristic function $\chi_i(t)$ $(0 \leq \chi_i \leq 1)$ that assumes the value 1 in the interior of the body and zero outside it. The shape of this function is defined by the half width $w(s)$ of the body along its arc-length s [30]. Swimming motion follows

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 06/28/14 to 128.103.224.4. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

the parameterization presented in [35] based on the instantaneous curvature κ_s of the midline of the body,

$$\kappa_s(s,t) = B(s) + K(s) \cdot \sin\left[2\pi(t/T - \tau(s))\right]. \tag{2.1}$$

where $B(s)$ and $K(s)$ are natural cubic splines through the baseline $\mathbf{B} = \{B_1, \ldots, B_6\}$ and undulatory $\mathbf{K} = \{K_1, \ldots, K_6\}$ curvature values at the six interpolation points $S_1 = 0$, $S_2 = 0.05L$, $S_3 = 0.33L$, $S_4 = 0.67L$, $S_5 = 0.95L$, $S_6 = L$ along the body of the fish, t is the time, and T is the swimming period ($T = 1$ in this work). The phase shift $\tau(s)$ is linearly proportional to the arc-length $\tau(s) = (s/L)\tau_{tail}$ and is responsible for a traveling wave along the body of the fish. The midline motion gives rise to the deformation velocity field $\mathbf{u}_i^{def}(t)$ of the whole body, as detailed in [30].

We emphasize that the time evolution of $\chi_i(t)$ and $\mathbf{u}_i^{def}(t)$ depends uniquely on the swimming pattern (determined by $\mathbf{B}$, $\mathbf{K}$, τ_{tail}) adopted by a fish. The swimmer's rigid components of motion, i.e., translational ($\mathbf{u}_i^t$, t for translation) and rotational ($\mathbf{u}_i^r$, r for rotation) velocities, and consequently the swimming trajectories, are instead the result of the flow-structure interaction.

2.2. Governing equations and numerical scheme. We consider a two-dimensional system in which a collection of N deforming bodies $\Omega_{i=1,\ldots,N}$ is immersed in a viscous, incompressible flow in the infinite domain Σ. Furthermore, bodies are assumed to be of the same density of the fluid ($\rho = 1$), a rather reasonable approximation for a fish mass composition. The governing Navier–Stokes equations are expressed in their velocity-vorticity ($\mathbf{u}$–ω, where $\omega = \nabla \times \mathbf{u}$) form, and no-slip boundary conditions at the surface of the bodies $\partial\Omega_i$ are approximated via Brinkman penalization [30]

$$\frac{\partial \omega}{\partial t} + \nabla \cdot (\mathbf{u}\omega) = \nu \nabla^2 \omega + \sum_{i=1}^{N} \lambda \nabla \times \chi_i (\mathbf{u}_i - \mathbf{u}), \qquad \mathbf{x} \in \Sigma, \tag{2.2}$$

$$\nabla \cdot \mathbf{u} = 0, \qquad \mathbf{x} \in \Sigma \backslash \Omega_i, \tag{2.3}$$

$$\nabla \cdot \mathbf{u}_i^{def} = 0, \qquad \mathbf{x} \in \Omega_i, \tag{2.4}$$

where ν is the kinematic viscosity, $\lambda \gg 1$ is the penalization factor, and $\mathbf{u}_i = \mathbf{u}_i^t + \mathbf{u}_i^r + \mathbf{u}_i^{def}$ are the velocity fields characterizing each body.

The system of equations (2.2)–(2.4) is discretized using remeshed vortex methods [37, 46, 47], in which vortex particles are regularized on an underlying cartesian grid through high order interpolation (in this work M_4' was used [48]) to prevent excessive gaps between particles [46]. The underlying regular grid allows for the fast evaluation of differential operators using finite difference schemes and enables the use of wavelet based adaptivity [49, 50]. In turn remeshing imposes constraints on the types of geometries that can be implemented [46]. However, when combined with penalization techniques such constraints are readily bypassed as the cartesian grid is distributed throughout the computational domain. The wavelet based, multiresolution analysis of vorticity and velocity fields captures the emergence (*grid refinement*) or dissipation (*grid compression*) of flow scales, providing an efficient allocation of computational elements [49]. In this work, we aim at simulating multiple self-propelled bodies swimming within large domains (relative to the fish dimensions); therefore the ability of focusing computational resources only where necessary (i.e., in the proximity of the fish and their wakes) is key.

The full algorithm, from time t^n to t^{n+1}, is detailed in (2.5)–(2.20), where all quantities are assumed to be known up to time t^n. We denote with $\mathcal{R}(a)|_{t_r}$ and with

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 06/28/14 to 128.103.224.4. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

$\mathcal{C}(a)|_{t_c}$, respectively, the refinement and compression stages based on the wavelet multiresolution analysis of the field a, using the refinement and compression thresholds t_r and t_c [49]. While evolving the solution of the system in time, refinement (2.5), computing (2.6)–(2.19), and compression (2.20) stages regularly alternate according to the adaptation frequency AF (i.e., $AF = 5$ corresponds to one readaptation stage every five computing stages).

At the beginning of each time step, (2.5)–(2.20), the fields χ_i and $\mathbf{u}_i^{def}$ are computed independently and subsequently interpolated at the appropriate location on the computational grid [30].

$$\textbf{refinement:} \quad \mathcal{R}(\omega^n \text{ and } \mathbf{u}^n)|_{t_r}, \tag{2.5}$$

$$\sigma^n = \sum_{i=1}^{N} \chi_i^n \left(\nabla \cdot \mathbf{u}_i^{def,n}\right), \tag{2.6}$$

$$\nabla^2 \psi^n = -\omega^n, \tag{2.7}$$

$$\nabla^2 \phi^n = \sigma^n, \tag{2.8}$$

$$\mathbf{u}^n = \nabla \times \psi^n + \nabla \phi^n, \tag{2.9}$$

$$\mathbf{u}_i^{t,n} = \frac{1}{m_i} \int_\Sigma \rho \chi_i^n \mathbf{u}^n d\mathbf{x}, \tag{2.10}$$

$$\dot{\theta}_i^n = \frac{1}{I_i} \int_\Sigma \rho \chi_i^n \left(\mathbf{x} - \mathbf{x}_i^n\right) \times \mathbf{u}^n d\mathbf{x}, \tag{2.11}$$

$$\mathbf{u}_i^{r,n} = \dot{\theta}_i^n \times \left(\mathbf{x} - \mathbf{x}_i^n\right), \tag{2.12}$$

$$\mathbf{u}_\lambda^n = \frac{\mathbf{u}^n + \lambda \Delta t \sum_{i=1}^N \chi_i^n \left(\mathbf{u}_i^{t,n} + \mathbf{u}_i^{r,n} + \mathbf{u}_i^{def,n}\right)}{1 + \lambda \Delta t \sum_{i=1}^N \chi_i^n}, \tag{2.13}$$

$$\omega_\lambda^n = \nabla \times \mathbf{u}_\lambda^n, \tag{2.14}$$

$$\frac{\partial \omega_\lambda^n}{\partial t} = \nu \nabla^2 \omega_\lambda^n, \tag{2.15}$$

$$\frac{\partial \omega_\lambda^n}{\partial t} + \nabla \cdot \left(\mathbf{u}_\lambda^n \, \omega_\lambda^n\right) = 0, \tag{2.16}$$

$$\omega^{n+1} = \omega_\lambda^{n+1}, \tag{2.17}$$

$$\mathbf{x}_i^{n+1} = \mathbf{x}_i^n + \mathbf{u}_i^{t,n} \Delta t^n, \tag{2.18}$$

$$\theta_i^{n+1} = \theta^n + \dot{\theta}_i^n \Delta t^n, \tag{2.19}$$

$$\textbf{compression:} \quad \mathcal{C}(\omega^{n+1} \text{ and } \mathbf{u}^{n+1})|_{t_c}. \tag{2.20}$$

The velocity field $\mathbf{u}$ is computed via Helmholtz decomposition (2.9) from vorticity ω and potential velocity σ (2.6) by solving the unbounded Poisson equations (2.7)–(2.8). The solution of Poisson equations on multiresolution regular grids is achieved through the Barnes–Hut multipole method [51], thus naturally accommodating the far field boundary condition.

Rigid components of motion ($\mathbf{u}_i^t$, $\mathbf{u}_i^r$) are captured through the projection approach [52, 53, 30] by (2.10)–(2.12), where $\mathbf{x}_i$, m_i, I_i, and $\dot{\theta}_i$ are, respectively, the position of the center of mass, the mass, the moment of inertia, and the angular velocity of the ith body.

Time integration is carried out via Godunov splitting. Velocity penalization (2.13)–(2.14) is performed by the implicit Euler scheme. Diffusion (2.15) and particle advection (2.16) are performed through a second order Runge–Kutta scheme coupled

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 06/28/14 to 128.103.224.4. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

with local time stepping [49, 54] and body advection (2.18)–(2.19) with an explicit Euler scheme. The use of a first order implicit Euler operator for velocity penalization and use of a Godunov splitting approach, which is first order in the present form, does not justify the use of a higher order time integration scheme for the advection of the body. All spatial operators are discretized with fourth order accurate finite difference schemes, except for the fifth order interpolating wavelets. The present algorithm is implemented using the multiresolution adapted grids library [49, 50, 54].

Throughout this work, if not specified otherwise, we consider the domain $\Sigma = [0,1] \times [0,1]$, set $\lambda = 10^4$, $t_r = 10^{-4}$, $t_c = 10^{-6}$, and the Lagrangian CFL condition is [47] $LCFL = 0.1$. Characteristic functions χ_i are computed as given in [30] with smoothing length $\epsilon = 2\sqrt{2}h^e$, where h^e denotes the effective grid spacing, i.e., the minimum grid spacing in the multiresolution representation. Moreover, hereafter we will refer to $ER = 1/h^e$ as the effective resolution.

2.3. Reinforcement learning. In this work the swimmer's behavior is identified through the use of an RL technique based on the one-step Q-learning algorithm [55, 56]. The choice of this learning algorithm is dictated by its simplicity as well as by the fact that it has been shown to converge to an optimal behavioral policy for finite Markov decision processes [55].

The swimming agent progressively learns how to act though a trial and error interaction with the environment. The experience gained by the agent i is represented by the tuple $\langle s_i, a_i, r_i, s_i' \rangle$, where s_i' is the next state reached by the agent i given that the action a_i was taken from state s_i, and r_i is the corresponding reward. The term s_i is a discrete set, encoding the sensorial perception of the agent. The term a_i identifies an action among the discrete set of actions attainable by the agent. The term r_i is a scalar value that encodes the goal of the swimmer. A large value of r_i indicates that the action taken was successful relatively to the agent's goal. The actual definition of state, actions, and reward is problem specific, and in the present work it will be introduced case by case.

The one-step Q-learning algorithm aims at approximating the action-value function $Q_i(s_i, a_i)$ which represents the expected long term reward for taking the action a_i given the state s_i. In essence, Q_i captures the intelligence or behavior of the swimming agent. In fact, the optimal behavior is determined from Q_i by choosing, given the current state s_i, the action $a_i \to \max_{a_i} Q_i(s_i, a_i)$ with probability $1 - \epsilon$. In turn, a suboptimal (given the current agent's knowledge) action is randomly selected with the probability ϵ. This procedure is referred to as the ϵ-greedy selection scheme [56], where the parameter $0 \le \epsilon \le 1$ determines the trade-off between exploitation of the gained knowledge and the exploration of new solutions. The flow of information in RL is sketched in Figure 1.

Throughout this work, we employ a variation of the one-step Q-learning algorithm in which $Q_i = Q$, i.e., only one global Q is used and all swimmers update Q based on their personal experience. This strategy, denoted as shared policy, has been computationally shown to accelerate the learning process almost linearly with the number of agents [57]. Therefore, at uniform time intervals δt (here $\delta t = T$), swimmers update Q according to

$$Q(s_i, a_i) = Q(s_i, a_i) + \varphi(\Delta Q)_i, \qquad i = 1, \dots, N, \tag{2.21}$$

$$(\Delta Q)_i = r_i + \gamma \max_{a_i} Q(s_i', a_i) - Q(s_i, a_i), \qquad i = 1, \dots, N, \tag{2.22}$$

where $0 \le \varphi \le 1$ is the learning rate and $0 \le \gamma \le 1$ is the discount parameter that values past experiences.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

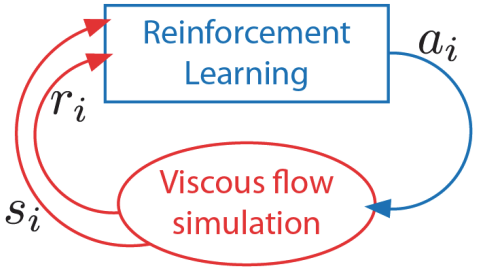

FIG. 1. *Schematic of RL coupled with viscous flow simulations of self-propelled swimmers. The goal of the i-th agent is encoded into the numerical reward r_i and the agent learns, through trial and error, how to map states s_i into actions a_i to maximize the expected long-term reward $Q_i(s_i, a_i)$ [56]. The system is simulated via direct numerical simulations that account for swimmer-swimmer dynamics mediated by the viscous flow.*

TABLE 1

Attainable swimmer's motion patterns (curvatures normalized by L). Note that curvature values at the extrema (B_1, B_6, K_1, K_6) of the midline are always set to zero.

	B_2	B_3	B_4	B_5	K_2	K_3	K_4	K_5	τ_{tail}
Forward	0.00	0.00	0.00	0.00	0.50	0.16	1.91	0.91	1.44
Right	0.50	0.50	0.50	0.50	0.50	0.16	1.91	0.91	1.44
Left	−0.50	−0.50	−0.50	−0.50	0.50	0.16	1.91	0.91	1.44

2.3.1. Actions. Maneuverings are determined by the available set of actions that swimming agents are able to perform. Different actions are modeled through combinations of baseline curvatures **B**, undulatory curvatures **K**, and phase shift at the tail τ_{tail}. We chose the undulatory parameters **K** and τ_{tail} based on Kern and Koumoutsakos [33], therefore obtaining anguilliform swimming. Baseline parameters **B** were instead determined to allow forward swimming as well as left and right turning maneuvers. The parameters corresponding to these actions are reported in Table 1.

Each transition from one action to the other is carried out through a cubic interpolation between starting and final curvatures (setting the first derivative at the extrema to zero to ensure smoothness) within the time interval T. The parameter τ_{tail} is initially ($t = 0$) ramped up via cubic interpolation from zero to its designated value at T.

We note that our motion parameterization allows us to model a wide range of maneuverings, as well as accelerations or decelerations, by modifying the swimming period T. Nevertheless, in this work we limited the number of actions to three ($N_a = 3$) in order to reduce as much as possible the state-action space dimension and therefore favor the learning process.

3. Validation of the flow solver.

3.1. Single self-propelled anguilliform swimmer. We present computations of a two-dimensional self-propelled swimmer and perform an error analysis in space and time of the present methodology. Furthermore, we compare our findings to reference simulations by Kern and Koumoutsakos [33], obtained using a finite volume method and body fitted grids.

All simulation in this section are performed considering, within the computational domain $\Sigma = [0, 1] \times [0, 1]$, an anguilliform swimmer characterized by the Carling motion pattern [27, 30]. Its shape and Reynolds number $\mathrm{Re}_{\mathrm{fish}} = L^2/(T\nu) = 7143$ are the same as reported by Kern and Koumoutsakos [33]. All simulations are started

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 06/28/14 to 128.103.224.4. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

with the body at rest and the motion is ramped up through a sinusoidal function within the first swimming cycle.

Simulations are carried out by setting $L = 0.3$, $AF = 1$, $t_r = 10^{-6}$, and $t_c = 10^{-8}$. The systems are simulated up to time $t = T$ and convergence orders are determined by computing the L^1, L^2, and L^∞ norm of the error $e(t) = \| v_{cm\ best\ resolved}(t) - v_{cm}(t) \|$, where v_{cm} is the speed of the fish center of mass.

Concerning space convergence, two different studies are performed by varying the effective resolution ER between 512×512 and 4096×4096 with 8192×8192 as the best resolved case and setting $LCFL = 0.001$. In the first study, we fix the model by setting the mollification length proportional to the coarsest effective grid spacing ($\epsilon = 2\sqrt{2}h^e_{512}$), while in the second study, the ratio $\epsilon/h^e = 2\sqrt{2}$ is chosen to be constant to investigate the convergence to the actual geometry. As can be seen in Figures 2(a) and 2(b), the method shows, respectively, second order convergence ($L^1 = 1.96$, $L^2 = 2.2$, $L^\infty = 1.96$) fixing ϵ and between first and second order ($L^1 = 1.4$, $L^2 = 1.4$, $L^\infty = 1.4$) fixing ϵ/h^e.

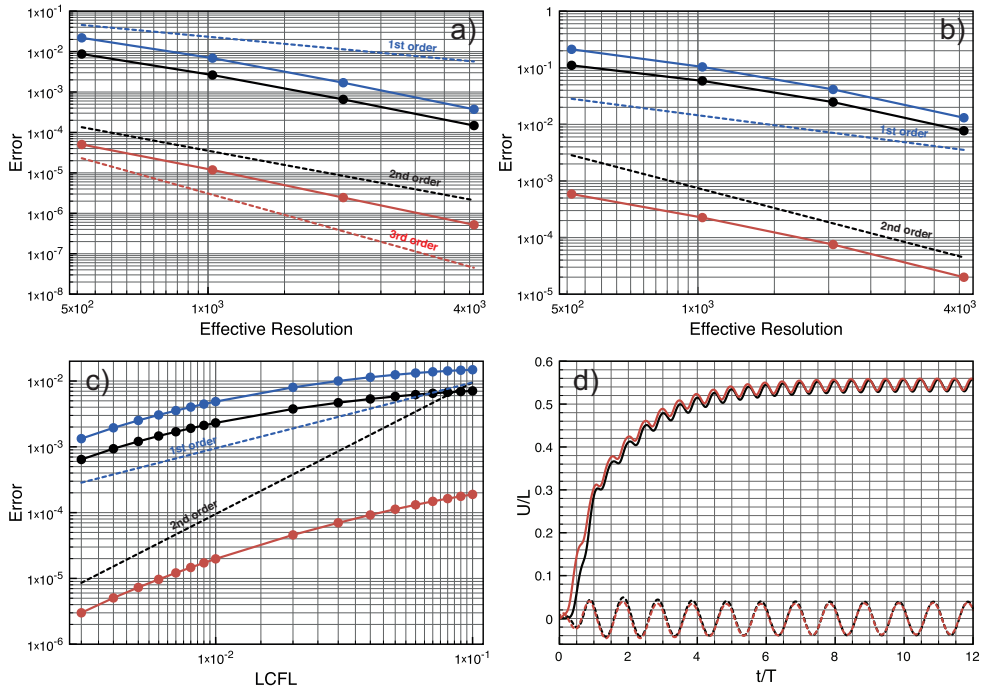

FIG. 2. *Convergence study for a swimmer characterized by the Carling motion pattern* [27, 30] *and a* $Re_{fish} = L^2/(T\nu) = 7143$. *Simulations were carried out up to physical time* $t = T$. (a) *Space convergence (mollification length fixed based on ER,* $\epsilon = 2\sqrt{2}h^e_{512}$*):* $L^\infty(e)$ *(blue),* $L^1(e)$ *(black), and* $L^2(e)$ *(red) are plotted against domain resolution. LCFL was set to* 0.001. (b) *Space convergence (ratio* ϵ/h^e *fixed to* $2\sqrt{2}$*):* $L^\infty(e)$ *(blue),* $L^1(e)$ *(black), and* $L^2(e)$ *(red) are plotted against domain resolution. LCFL was set to* 0.001. (c) *Time convergence:* $L^\infty(e)$ *(blue),* $L^1(e)$ *(black), and* $L^2(e)$ *(red) are plotted against LCFL. ER was set to* 4096×4096 *and* $\epsilon = 2\sqrt{2}h^e$. *For all studies we used* $\lambda = 10^4$. *Dashed blue, black, and red lines represent, respectively, first, second, and third order slopes.* (d) *Comparison with reference finite volume solution* [33]*: solid and dashed red lines represent, respectively, forward* $U_\parallel$ *and lateral velocity* $U_\perp$ *obtained by the present method, and solid and dashed black lines indicate, respectively, reference forward* $U_\parallel$ *and lateral velocity* $U_\perp$. *Settings:* $L = 0.1$, $LCFL = 0.01$, $AF = 5$, $t_r = 10^{-4}$, $t_c = 10^{-6}$, $ER = 2^{14} \times 2^{14}$, *and* $\epsilon = 2\sqrt{2}h^e$.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 06/28/14 to 128.103.224.4. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

Downloaded 06/28/14 to 128.103.224.4. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

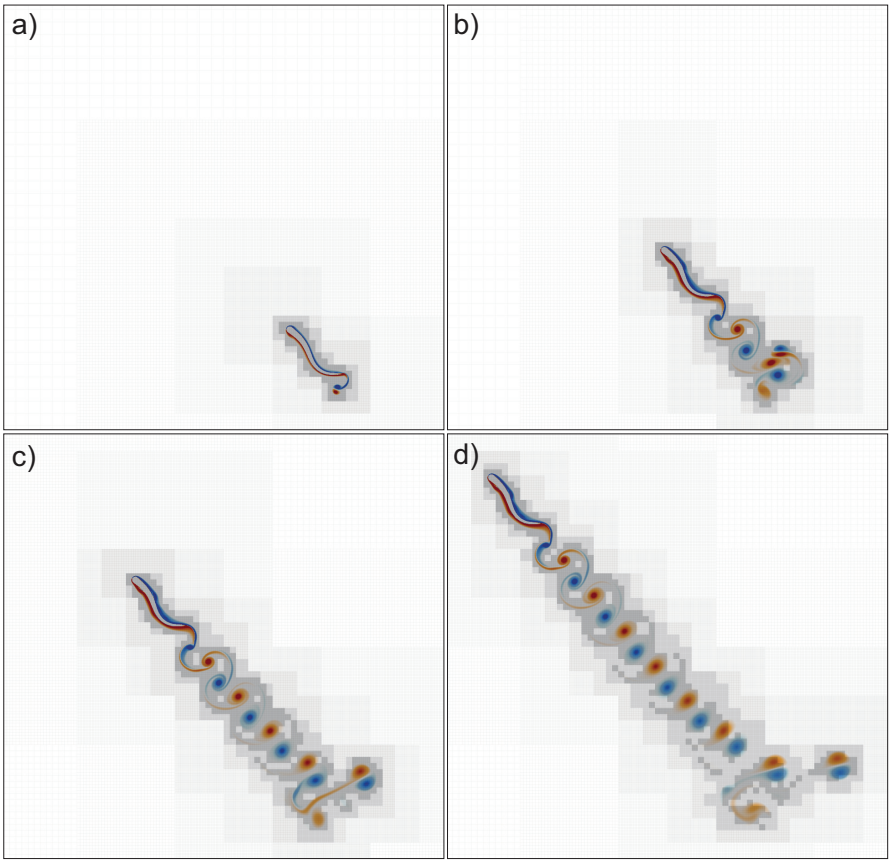

FIG. 3. *Self-propelled anguilliform swimmer. Wavelet adapted vorticity fields:* (a)–(d) $t = 1T$, $t = 4T$, $t = 7T$, $t = 10T$.

Time convergence study is performed setting $ER = 4096 \times 4096$ and $\epsilon = 2\sqrt{2}h^e$ and varying LCFL between 0.1 and 0.003 with 0.001 as the best resolved case. The order of convergence was found to be between first and second ($L^1 = 1.2$, $L^2 = 1.7$, $L^\infty = 1.2$) (Figure 2(c)). The convergence rates are therefore in agreement with our spatial and time discretization schemes.

Moreover, as can be seen in Figure 2(d), the forward and lateral swimming velocities obtained with the proposed method agree well with benchmark simulations carried out by Kern and Koumoutsakos [33]. Results reported in Figure 2(d) were obtained by setting $L = 0.1$, $LCFL = 0.01$, $AF = 5$, $t_r = 10^{-4}$, $t_c = 10^{-6}$, $ER = 2^{14} \times 2^{14}$, and $\epsilon = 2\sqrt{2}h^e$. The corresponding wavelet adapted vorticity fields are illustrated in Figure 3. As can be noticed the computational elements are efficiently concentrated where flow features appear.

3.2. Multiple self-propelled anguilliform swimmers. In this section, we illustrate the capability of the present flow solver to account for multiple self-propelled geometries doubly connected through the flow. In this context, we consider different swimming arrangements, namely, a pair of antiphase fish (Figure 4), a parallel formation (Figure 5), and a V-shaped (Figure 6) five-member school. The swimmers considered here are geometrically identical to that presented in the previous section

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

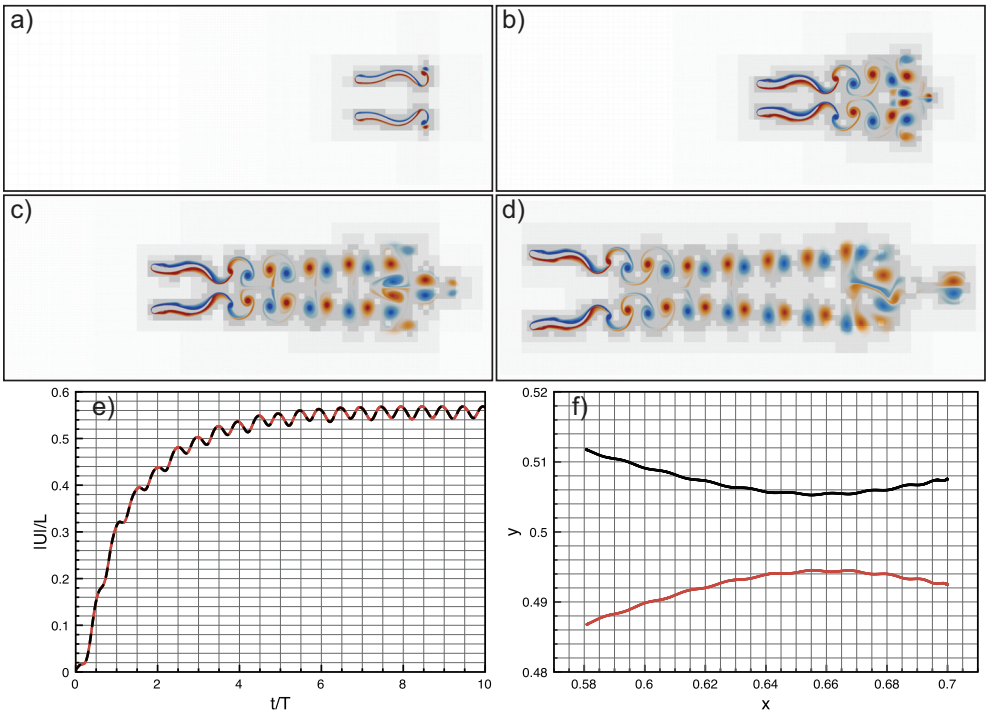

FIG. 4. *Antiphase self-propelled anguilliform swimmers.* (a)–(d) *Wavelet adapted vorticity fields at, respectively,* $t = T$, $t = 4T$, $t = 7T$, $t = 10T$. (e) *Absolute normalized velocities* $|\mathbf{U}|/L$. (f) *Swimmers' center of mass trajectories.*

and are characterized by the Carling motion. We refer to arrangements of such fish as passive schools, because swimmers cannot adjust their motion pattern to react to the flow. The reported simulations were obtained by settings $\mathrm{Re}_{\mathrm{fish}} = 7143$, $L = 0.025$, $LCFL = 0.05$, $AF = 5$, $t_r = 10^{-4}$, $t_c = 10^{-6}$, $ER = 2^{15} \times 2^{15}$, and $\epsilon = 2\sqrt{2}h^e$.

As can be notice in Figure 4, the vortical structures generated by one of the two swimmers influence its dynamics as well as the dynamics of the partner. The swimmers initially converge toward each other, while in a second phase the stronger velocity field (velocity decays inversely proportional to the distance from the vorticity source) generated by the boundary layers and the shed dipoles, leads to divergent trajectories (Figure 4(f)), consistently with [30]. Figures 4(a)–4(d) illustrate the corresponding wavelet adapted vorticity fields. As can be seen, the computational elements are efficiently placed to capture the generated vortical structures. Furthermore, the computed vorticity fields (and swimming speeds and trajectories) are remarkably symmetric, as expected given the system configuration, thus demonstrating the numerical robustness of the proposed methodology. Only after 10 swimming cycles (Figure 4(d)) a slight asymmetry in the far end of the wake is observed, due to numerical noise. Such effect can be amended for by decreasing refinement and compression tolerances.

We simulate two schooling formations, parallel (Figure 5) and V-shaped (Figure 6), to illustrate how different schooling arrangements, due to nonlinear flow mediated interactions, affect swimmers' speed and trajectories. As can be noticed in Figure 6(f), V-shaped schools favor rectilinear trajectories, contrary to parallel

Downloaded 06/28/14 to 128.103.224.4. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 06/28/14 to 128.103.224.4. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

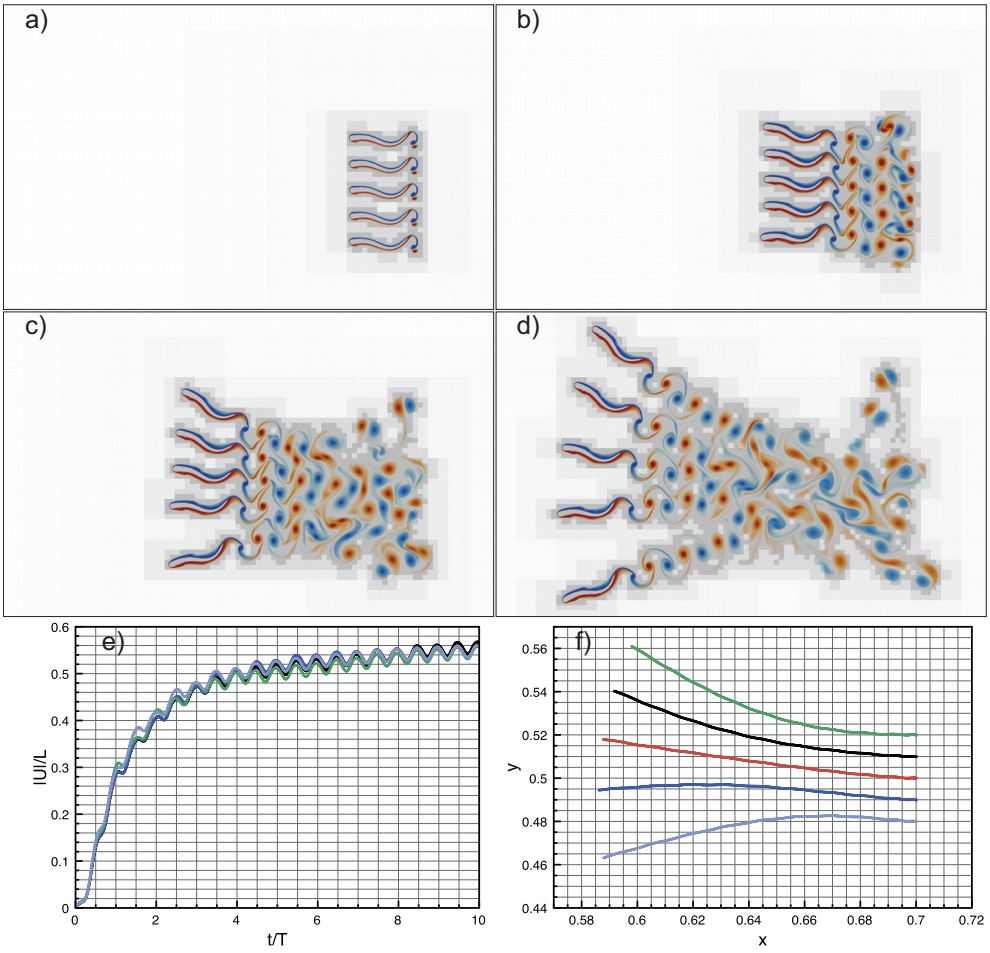

FIG. 5. *Parallel schooling formation.* (a)–(d) *Wavelet adapted vorticity fields at, respectively,* $t = T$, $t = 4T$, $t = 7T$, $t = 10T$. (e) *Absolute normalized velocities* $|\mathbf{U}|/L$. (f) *Swimmers' center of mass trajectories.*

arrangements, which lead to the immediate divergence of their school members (5(f)). These results are consistent with the simulations reported in [30]. Wavelet adapted vorticity fields associated with parallel and V-shaped schools are illustrated, respectively, in Figures 5(a)–(d) and Figures 6(a)–(d).

To the best of our current knowledge, these are the first multiresolution simulations of individual and multiple self-propelled swimmers.

4. Learning applications. We present two studies in which the one-step Q-learning algorithm (section 2.3) is coupled to the present flow solver. These studies involve individual and multiple self-propelled swimming agents and demonstrate the potential of our approach for the investigation of schooling behavior.

4.1. Swimming in circles. In this section we consider a task in which the swimming agent learns how to swim in order to follow a predefined circular trajectory Γ_t (Figure 7). The goal of following a specified path is mathematically cast into a

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 06/28/14 to 128.103.224.4. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

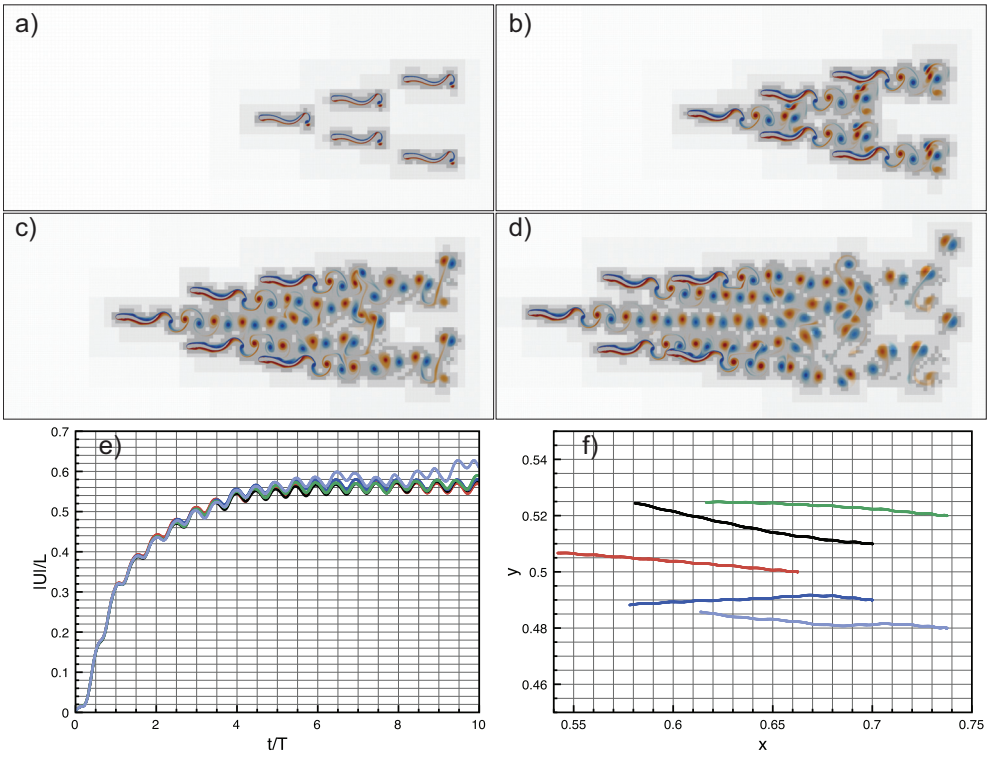

FIG. 6. *V-shape schooling formation.* (a)–(d) *Wavelet adapted vorticity fields at, respectively,* $t = T$, $t = 4T$, $t = 7T$, $t = 10T$. (e) *Absolute normalized velocities* $|\mathbf{U}|/L$. (f) *Swimmers' center of mass trajectories.*

function of the radial distance d from the fish center of mass to Γ_t. The reward r is therefore defined as

$$r = 1 - \frac{d}{R_{\text{out}} - R_{\text{in}}}, \tag{4.1}$$

where R_{out} and R_{in} are, respectively, the outer and inner radii of a two-dimensional toroidal region located at the center of the domain Σ (Figure 7). The radius of Γ_t is defined as $R_{\Gamma_t} = (R_{\text{out}} + R_{\text{in}})/2$. We note that the reward r reaches the maximum value of one only if the fish center of mass is exactly on top of Γ_t ($d = 0$); otherwise it linearly decreases for $d > 0$.

The swimming agent defines its state by sensing the radial distance d and orientation β relative to Γ_t (Figure 7(d)). The quantities d and β are mapped into a set of, respectively, $L_d = 5$ and $L_\beta = 8$ states. Such states correspond to the discretization of the ranges $\Delta_d = [0, L]$ and $\Delta_\beta = [0, 2\pi]$ into, respectively, L_d and L_β intervals according to $s = \{\max(L_d, \min(0, \lfloor dL_d/\Delta d \rfloor)), \max(L_\beta, \min(0, \lfloor \beta L_\beta/\Delta\beta \rfloor))\}$. In total, the swimming agent is characterized by the state-action space of dimension $L_d \times L_\beta \times N_a = 120$.

In this section, simulations have been performed considering swimmers of length $L = 0.025$, $AF = 5$, $ER = 2^{14} \times 2^{14}$ and a flow regime characterized by $\text{Re}_{\text{fish}} = 550$. This value has been chosen because of biological relevance to larval anguilliform

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 06/28/14 to 128.103.224.4. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

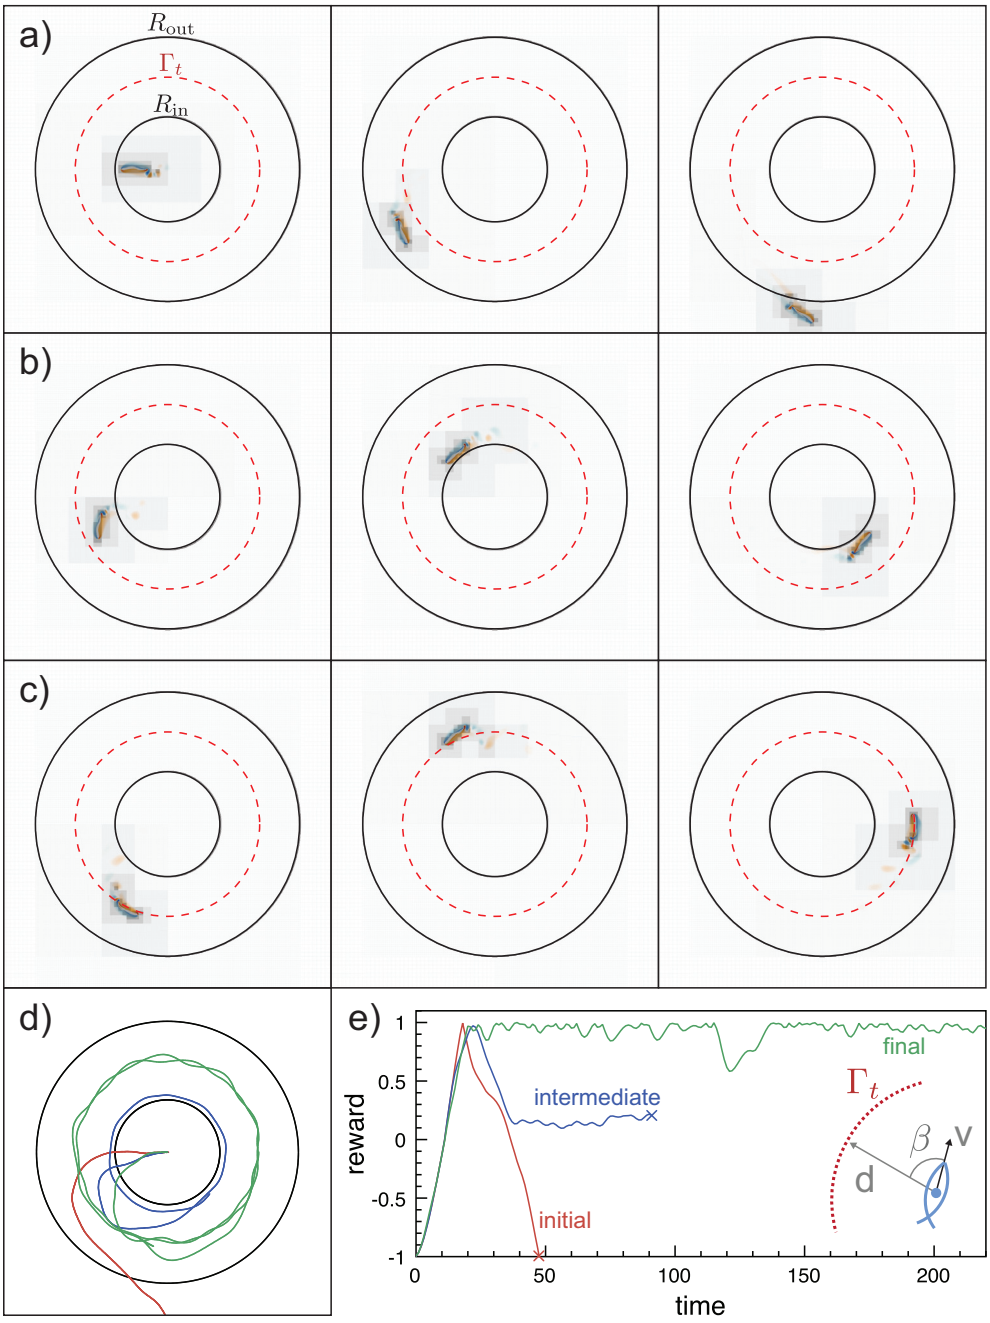

FIG. 7. *Single swimmer's vorticity fields corresponding to* (a) *initial,* (b) *intermediate, and* (c) *final learning stages and* (d) *associated trajectories and* (e) *rewards.*

swimmers [58, 35, 36]. Learning parameters have been set to $\varphi = 0.05$, $\gamma = 0.99$, and $\epsilon = 0.05$.

As can be noticed in Figure 7(a), the swimmer initially does not know how to appropriately act to achieve its goal and crosses the toroidal region of interest, quickly leaving the domain. Nevertheless, after a trial and error exploration period

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 06/28/14 to 128.103.224.4. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

the swimming agent learns that is preferable to swim inside the two-dimensional torus (Figure 7(b)). Eventually, within a learning period of approximately ~ 2500 swimming cycles, the agent learns to accurately follow its assigned trajectory Γ_t (Figures 7(c), (d)) in the attempt of maximizing the reward r (Figure 7(e)). In this scenario, swimmers are restarted while retaining the learned action-value matrix Q.

The same experiment is performed considering four swimmers. In the beginning, they are located tangentially to the circular path in a north-south and west-east configuration. Similar to the single swimmer case, after an exploration phase (Figures 8(a), (b)) they learn how to alternate actions in order to follow the prescribed path (Figures 8(c), (d)) and maximize the reward (Figure 8(e)). We note that in this case, fish also have to learn how to deal with the disturbances related to the wake of the leading swimmer. Despite this additional difficulty, they manage to appropriately execute their task (Figures 8(d), (e)). It must be noted that all four swimmers share the same policy and therefore interact with the same global action-value matrix Q.

4.2. Swimming in a parallel arrangement. In this section, we consider a task in which five swimmers learn to swim in a structured arrangement. In particular they have to maintain the initially equally spaced (interdistance set to $0.8L$) parallel configuration of Figure 9. As highlighted in previous studies [30] and in section 3.2, this is a nontrivial task. In fact, flow mediated interactions affect the fish dynamics, causing either the separation of school members or collisions among themselves. Therefore, maintaining a given group configuration requires constant and accurate decision making by the fish.

The goal of maintaining a given position within a predefined schooling arrangement is achieved by assigning to each swimmer a target point $\mathbf{x}_i^T$ and by letting the agents learn how to follow their targets. The points $\mathbf{x}_i^T$ that define the school structure are advected in a specified schooling direction $\mathbf{D}_S$ (Figure 9) according to $\mathbf{u}_i^T = \mathbf{u}_i^t \cdot \mathbf{D}_S$, where $\mathbf{u}_i^t$ is the swimmer translational velocity and $\mathbf{u}_i^T$ is the target point velocity. Target points are initialized one body length L in front of the corresponding swimmer according to $\mathbf{x}_i^T = \mathbf{x}_i^t + L\mathbf{D}_S$. Therefore, the reward is chosen to reflect how well the swimmer can follow its assign point,

$$r_i = -\frac{d_i^2}{L^2} + \frac{2d_i}{L}, \tag{4.2}$$

where $d_i = |\mathbf{x}_i^T - \mathbf{x}_i^t|$ is the distance between the agent center of mass and its target point. We note that the reward reaches the maximum value of one only if the relative position between fish center of mass and target point matches exactly the initial configuration.

The swimming agent defines its state by sensing the distance d_i and the orientation $\beta_i = \arg(\mathbf{x}_i^T - \mathbf{x}_i) - \theta_i$ relative to its target point (Figure 9). The quantities d_i and β_i are mapped into a set of, respectively, $L_d = 20$ and $L_\beta = 16$ states. Such states correspond to the discretization of the ranges $\Delta_d = [0, L]$ and $\Delta_\beta = [0, 2\pi]$ into, respectively, L_d and L_β intervals according to $s_i = \{\max(L_d, \min(0, \lfloor d_i L_d/\Delta d \rfloor)), \max(L_\beta, \min(0, \lfloor \beta_i L_\beta/\Delta\beta \rfloor))\}$. In total, the swimming agent is characterized by the state-action space of dimension $L_d \times L_\beta \times N_a = 960$.

In this section, simulations have been performed considering swimmers of length $L = 0.025$, $ER = 2^{14} \times 2^{14}$ and a flow regime characterized by $\mathrm{Re}_{\mathrm{fish}} = 550$. Learning parameters have been set to $\varphi = 0.05$, $\gamma = 0.99$, and $\epsilon = 0.05$. Overall the learning process entailed approximately $\sim 10^4$ swimming cycles.

As can be noticed in Figure 10, the swimmers learn a behavioral policy that allow them to maintain the prescribed schooling formation throughout the entire length of

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 06/28/14 to 128.103.224.4. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

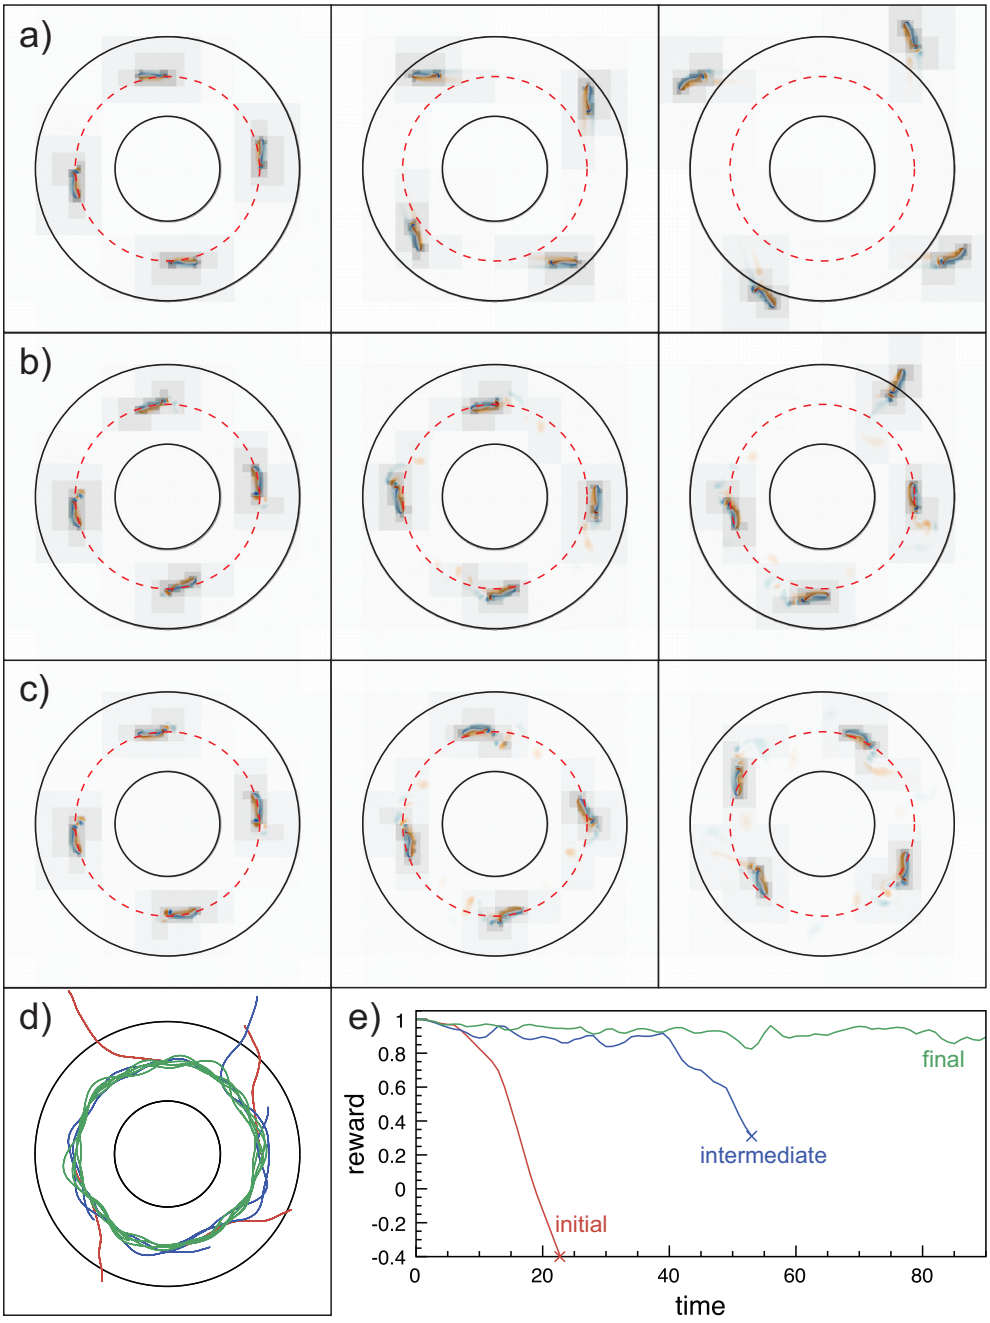

FIG. 8. *Multiple swimmers' vorticity fields corresponding to* (a) *initial,* (b) *intermediate, and* (c) *final learning stages and* (d) *associated trajectories and* (e) *normalized total reward* $R = \frac{1}{4}\sum_{i=1}^{4} r_i$.

the domain (approximately ∼500 swimming cycles). Their ability in parallel schooling is reflected by the overall normalized reward (Figure 10(c)) which oscillates between 96% and 100% of the maximum achievable value. School member trajectories are reported in Figure 10(b), together with the trajectories of the corresponding passive parallel school. As can be noticed, the parallel arrangement is responsible for relevant

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

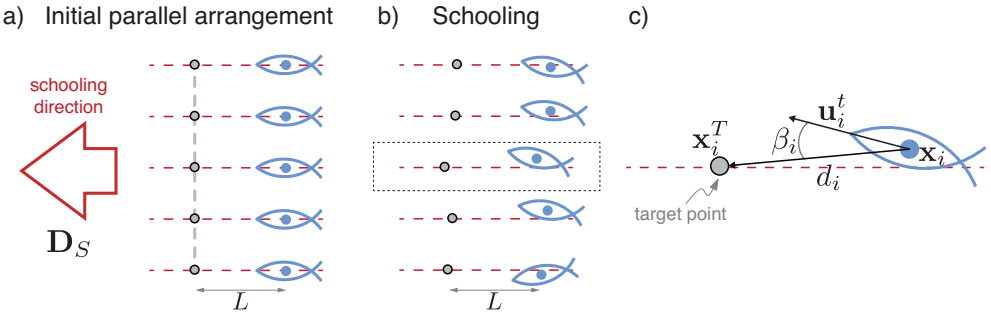

FIG. 9. *Swimmers state representation.* (a) *Initial conditions for a parallel arrangement.* (b) *Swimmers and relative target points configuration during schooling.* (c) *State representation of an individual swimmer attempting to follow its assigned target point.*

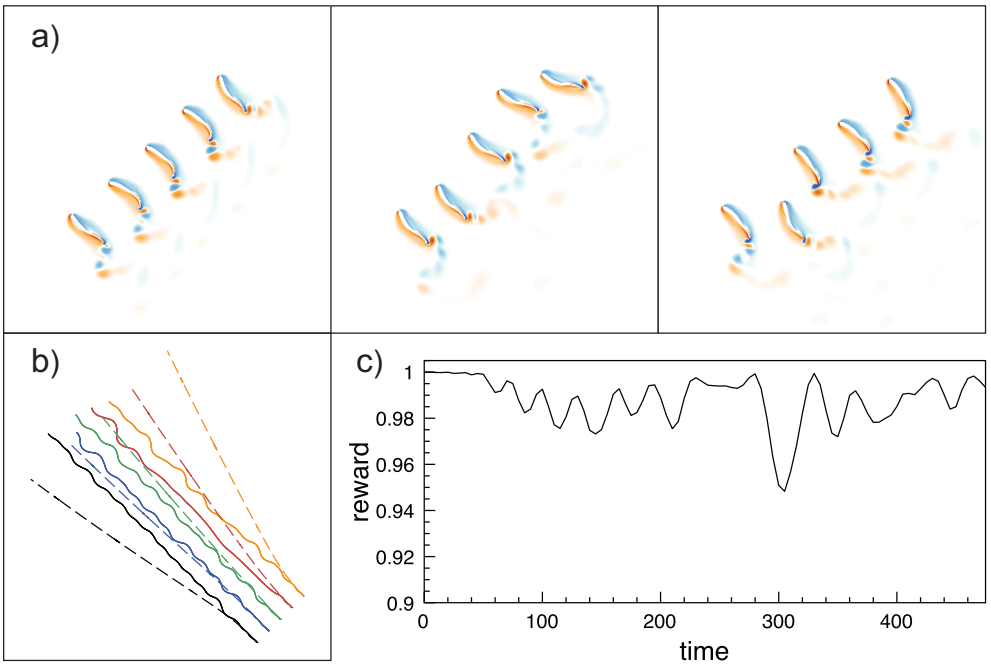

FIG. 10. *Vorticity fields generated by a five-member parallel school* (a) *at, from left to right,* $t \approx 10T$, $t \approx 200T$, *and* $t \approx 400T$. *Relative trajectories (continuous lines) are illustrated in* (b) *together with those of the corresponding passive school (dashed lines) as reference.* (c) *Normalized total reward* $R = \frac{1}{5}\sum_{i=1}^{5} r_i$.

hydrodynamic side forces on the lateral school members, leading to substantial trajectory deflections. This observation suggests that the use of a shared policy, although effective, may not be optimal. In fact, depending on the relative locations within the group, swimmers may experience substantially different hydrodynamic forces. In this context, an approach in which the main behavioral traits are learned through a shared policy (computationally cheaper) and are subsequently refined via individual policy specializations may be beneficial. A thorough exploration of different learning strategies goes beyond the scope of this paper and is the subject of future work.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 06/28/14 to 128.103.224.4. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

Downloaded 06/28/14 to 128.103.224.4. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

5. Conclusion. We have developed a computational method for the study of multiple self-propelled swimmers coordinating their motion to exhibit collective behavior. The flow solver relies on a remeshed vortex method, coupled with a Brinkman penalization and projection approach and enhanced with wavelet based adaptivity for the efficient allocation of vortex particles. The behavioral policy of the swimmers is determined by an RL algorithm that maps the state of the swimmers to suitable actions so as to achieve a priori defined tasks. The applicability of this approach to the study of schooling hydrodynamics and behavior is demonstrated in several learning tasks involving single and multiple swimmers. We consider the present approach as the first of its kind that enables the study of viscous flow mediated interactions in collective hydrodynamics.

The RL algorithm is shown to be essential in achieving schooling behavior. In turn the feasibility of these simulations is due to the use of multiresolution techniques. In addition to the economy of representation afforded by a velocity-vorticity formulation and the Lagrangian character of the vortex methods, wavelet adaptivity enables the allocation of computational elements to critical areas of the flow field. These techniques allow us to simulate large domains in which several swimmers can perform their maneuverings, an intractable scenario in a uniform resolution context.

Ongoing work focuses on the definition of a new set of actions to enable the swimmers to perform different maneuvers and on the use of sensory information available from optical flow and the fish lateral line. Furthermore, we pursue the investigation of different schooling formations and their characterization from an energetic standpoint.

Acknowledgments. We thank the Swiss National Science Foundation for financial support.

REFERENCES

[1] B. L. Partridge, *The structure and function of fish schools*, Scientific American, 246 (1982), pp. 114–123.
[2] W. D. Hamilton, *Geometry for selfish herd*, J. Theoret. Biol., 31 (1971), pp. 295–311.
[3] S. R. J. Neill and J. M. Cullen, *Experiments on whether schooling by their prey affects the hunting behaviour of cephalopods and fish predators*, J. Zoology, 172 (1974), pp. 549–569.
[4] P. F. Major, *Predator-prey interactions in two schooling fishes, Caranx ignobilis and Stolephorus purpureus*, Animal Behaviour, 26 (1978), pp. 760–777.
[5] J. Lazarus, *The early warning function of flocking in birds: An experimental study with captive quelea*, Animal Behaviour, 27 (1979), pp. 855–865.
[6] L. Landeau and J. Terborgh, *Oddity and the "confusion effect" in predation*, Animal Behaviour, 34 (1986), pp. 1372–1380.
[7] J. R. Krebs, M. H. MacRoberts, and J. M. Cullen, *Flocking and feeding in the great tit parus major—an experimental study*, Ibis, 114 (1972), pp. 507–530.
[8] T. J. Pitcher, A. E. Magurran, and I. J. Winfield, *Fish in larger shoals find food faster*, Behavioral Ecology Sociobiology, 10 (1982), pp. 149–151.
[9] K. S. Norris and C. R. Schilt, *Cooperative societies in three-dimensional space: On the origins of aggregations, flocks, and schools, with special reference to dolphins and fish*, Ethology Sociobiology, 9 (1988), pp. 149–179.
[10] J. Herskin and J. F. Steffensen, *Energy savings in sea bass swimming in a school: Measurements of tail beat frequency and oxygen consumption at different swimming speeds*, J. Fish Biology, 53 (1998), pp. 366–376.
[11] C. M. Breder, *Vortices and fish schools*, Zoologica New York, 50 (1965), pp. 97–114.
[12] V. V. Belyayev and G. V. Zuyev, *Hydrodynamic hypothesis of school formation in fishes*, Problems of Ichthyology, 9 (1969), pp. 578–584.
[13] D. Weihs, *Hydromechanics of fish schooling*, Nature, 241 (1973), pp. 290–291.
[14] P. B. Lissaman and C. A. Shollenberger, *Formation flight of birds*, Science, 168 (1970), pp. 1003–1005.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

[15] J. C. Liao, D. N. Beal, G. V. Lauder, and M. S. Triantafyllou, *Fish exploiting vortices decrease muscle activity*, Science, 302 (2003), pp. 1566–1569.
[16] R. W. Whittlesey, S. Liska, and J. O. Dabiri, *Fish schooling as a basis for vertical axis wind turbine farm design*, Bioinspiration Biomimetics, 5 (2010).
[17] J. O. Dabiri, *Potential order-of-magnitude enhancement of wind farm power density via counter-rotating vertical-axis wind turbine arrays*, J. Renewable Sustainable Energy, 3 (2011), 043104.
[18] S. J. Portugal, T. Y. Hubel, J. Fritz, S. Heese, D. Trobe, B. Voelkl, S. Hailes, A. M. Wilson, and J. R. Usherwood, *Upwash exploitation and downwash avoidance by flap phasing in ibis formation flight*, Nature, 505 (2014), pp. 399–402.
[19] S. Nair and E. Kanso, *Hydrodynamically coupled rigid bodies*, J. Fluid Mech., 592 (2007), pp. 393–411.
[20] A. A. Tchieu, D. Crowdy, and A. Leonard, *Fluid-structure interaction of two bodies in an inviscid fluid*, Physics of Fluids, 22 (2010), 107101.
[21] A. A. Tchieu, E. Kanso, and P. K. Newton, *The finite-dipole dynamical system*, Proc. Roy. Soc. A, 468 (2012), pp. 3006–3026.
[22] M. Gazzola, C. Mimeau, A. A. Tchieu, and P. Koumoutsakos, *Flow mediated interactions between two cylinders at finite re numbers*, Physics of Fluids, 24 (2012), 043103.
[23] A. Gilmanov and F. Sotiropoulos, *A hybrid cartesian/immersed boundary method for simulating flows with* 3*d, geometrically complex, moving bodies*, J. Comput. Phys., 207 (2005), pp. 457–492.
[24] I. Borazjani and F. Sotiropoulos, *Numerical investigation of the hydrodynamics of carangiform swimming in the transitional and inertial flow regimes*, J. Experimental Biology, 211 (2008), pp. 1541–1558.
[25] H. Dong, M. Bozkurttas, R. Mittal, P. Madden, and G. V. Lauder, *Computational modelling and analysis of the hydrodynamics of a highly deformable fish pectoral fin*, J. Fluid Mech., 645 (2010), pp. 345–373.
[26] C. L. Yu, S. C. Ting, M. K. Yeh, and J. T. Yang, *Three-dimensional numerical simulation of hydrodynamic interactions between pectoral-fin vortices and body undulation in a swimming fish*, Physics of Fluids, 23 (2011), 091901.
[27] J. Carling, T. L. Williams, and G. Bowtell, *Self-propelled anguilliform swimming: Simultaneous solution of the two-dimensional navier-stokes equations and Newton's laws of motion*, J. Experimental Biology, 201 (1998), pp. 3143–3166.
[28] M. Bergmann and A. Iollo, *Modeling and simulation of fish-like swimming*, J. Comput. Phys., 230 (2011), pp. 329–348.
[29] J. D. Eldredge, *Numerical simulations of undulatory swimming at moderate reynolds number*, Bioinspiration Biomimetics, 1 (2006), pp. S19–S24.
[30] M. Gazzola, P. Chatelain, W. M. van Rees, and P. Koumoutsakos, *Simulations of single and multiple swimmers with non-divergence free deforming geometries*, J. Comput. Phys., 230 (2011), pp. 7093–7114.
[31] I. Borazjani, F. Sotiropoulos, E. D. Tytell, and G. V. Lauder, *Hydrodynamics of the bluegill sunfish c-start escape response: Three-dimensional simulations and comparison with experimental data*, J. Experimental Biology, 215 (2012), pp. 671–684.
[32] A. A. Shirgaonkar, M. A. MacIver, and N. A. Patankar, *A new mathematical formulation and fast algorithm for fully resolved simulation of self-propulsion*, J. Comput. Phys., 228 (2009), pp. 2366–2390.
[33] S. Kern and P. Koumoutsakos, *Simulations of optimized anguilliform swimming*, J. Experimental Biology, 209 (2006), pp. 4841–4857.
[34] S. E. Hieber and P. Koumoutsakos, *An immersed boundary method for smoothed particle hydrodynamics of self-propelled swimmers*, J. Comput. Phys., 227 (2008), pp. 8636–8654.
[35] M. Gazzola, W. M. Van Rees, and P. Koumoutsakos, *C-start: Optimal start of larval fish*, J. Fluid Mech., 698 (2012), pp. 5–18.
[36] W. M. van Rees, M. Gazzola, and P. Koumoutsakos, *Optimal shapes for anguilliform swimmers at intermediate reynolds numbers*, J. Fluid Mech., 722 (2013).
[37] P. Koumoutsakos, *Inviscid axisymmetrization of an elliptical vortex*, J. Comput. Phys., 138 (1997), pp. 821–857.
[38] C. W. Reynolds, *Flocks, herds and schools: A distributed behavioural model*, in SIGGRAPH '87: Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques, 1987, pp. 71–87.
[39] I. D. Couzin and J. Krause, *Self-organization and collective behavior in vertebrates*, Advances in the Study of Behavior, 32 (2003), pp. 1–75.

Downloaded 06/28/14 to 128.103.224.4. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

[40] T. Vicsek, A. Czirok, E. Benjacob, I. Cohen, and O. Shochet, *Novel type of phase-transition in a system of self-driven particles*, Phys. Rev. Lett., 75 (1995), pp. 1226–1229.
[41] M. Ballerini, N. Cabibbo, R. Candelier, A. Cavagna, E. Cisbani, I. Giardina, V. Lecomte, A. Orlandi, G. Parisi, A. Procaccini, M. Viale, and V. Zdravkovic, *Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study*, Proc. Natl. Acad. Sci. USA, 105 (2008), pp. 1232–1237.
[42] I. D. Couzin, C. C. Ioannou, G. Demirel, T. Gross, C. J. Torney, A. Hartnett, L. Conradt, S. A. Levin, and N. E. Leonard, *Uninformed individuals promote democratic consensus in animal groups*, Science, 334 (2011), pp. 1578–1580.
[43] I. D. Couzin, J. Krause, R. James, G. D. Ruxton, and N. R. Franks, *Collective memory and spatial sorting in animal groups*, J. Theoret. Biol., 218 (2002), pp. 1–11.
[44] S. V. Viscido, J. K. Parrish, and D. Grunbaum, *The effect of population size and number of influential neighbors on the emergent properties of fish schools*, Ecological Modelling, 183 (2005), pp. 347–363.
[45] T. Vicsek and A. Zafeiris, *Collective motion*, Phys. Rep., 517 (2012), pp. 71–140.
[46] G. H. Cottet and P. Koumoutsakos, Vortex Methods, Theory and Practice, Cambridge University Press, New York, 2000.
[47] P. Koumoutsakos, *Multiscale flow simulations using particles*, Ann. Rev. Fluid Mechanics, 37 (2005), pp. 457–487.
[48] J. J. Monaghan, *Extrapolating b splines for interpolation*, J. Comput. Phys., 60 (1985), pp. 253–262.
[49] B. Hejazialhosseini, D. Rossinelli, M. Bergdorf, and P. Koumoutsakos, *High order finite volume methods on wavelet-adapted grids with local time-stepping on multicore architectures for the simulation of shock-bubble interactions*, J. Comput. Phys., 229 (2010), pp. 8364–8383.
[50] D. Rossinelli, B. Hejazialhosseini, M. Bergdorf, and P. Koumoutsakos, *Wavelet-adaptive solvers on multi-core architectures for the simulation of complex systems*, Concurrency and Computation: Practice and Experience, 23 (2011), pp. 172–186.
[51] J. Barnes and P. Hut, *A hierarchical o(n-log-n) force-calculation algorithm*, Nature, 324 (1986), pp. 446–449.
[52] N. A. Patankar and N. Sharma, *A fast projection scheme for the direct numerical simulation of rigid particulate flows*, Commun. Numer. Methods Eng., 21 (2005), pp. 419–432.
[53] M. Coquerelle and G. H. Cottet, *A vortex level set method for the two-way coupling of an incompressible fluid with colliding rigid bodies*, J. Comput. Phys., 227 (2008), pp. 9121–9137.
[54] D. Rossinelli, *Multiresolution Flow Simulations on Multi/Many-Core Architectures*, Ph.D. thesis, ETH Zurich, 2011.
[55] C. Watkins and P. Dayan, *Q-learning*, Machine Learning, 8 (1992), pp. 279–292.
[56] S. R. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, MIT Press Cambridge, MA, 1998.
[57] M. Gazzola, *Simulation, Optimization and Learning of Artificial Swimmers*, Ph.D. thesis, ETH Zurich, 2012.
[58] U. K. Muller, J. G. M. van den Boogaart, and J. L. van Leeuwen, *Flow patterns of larval fish: Undulatory swimming in the intermediate flow regime*, J. Experimental Biology, 211 (2008), pp. 196–205.

Downloaded 06/28/14 to 128.103.224.4. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.