Swarm Approach to Chemical Source Localization

Dimitri Zarzhitsky and Diana F. Spears
Computer Science Department
University of Wyoming
Laramie, WY, 82071
dimzar@uwyo.edu

Abstract – *We present a physics-based approach to the localization of chemical sources with autonomous swarms. Robotic vehicles with short-range sensors use local pairwise interactions to self-assemble into structured lattice formations, serving as distributed sensor and computational meshes. The robots use fluid flow information to navigate toward the chemical emitter. We develop a new search algorithm from first principles of fluid mechanics that outperforms the leading biomimetic competitors for the chemical source localization task. A validation of the scalability of our solution via simulation of plume and vehicle dynamics is given.*

Keywords: emergent behavior, robotic swarms, chemical plume tracing, adaptive sensor networks.

1 Introduction

The reduction of manufacturing costs and increased availability of mass-produced robotic platforms makes it possible to deploy large teams of inexpensive vehicles suitable for a variety of important assignments, such as search and rescue missions, battlefield surveillance, and facilitation of emergency communications. However, a system consisting of many independent agents requires a practical, scalable control framework to achieve the desired behavior, with a robust mechanism for dealing with system failures, and a capability to operate autonomously without global control or supervision. Ability to formally assure the system's long-term behavior is also needed. This paper examines a physics-based approach to managing swarm complexity in a theoretically sound and practically efficient manner. We test our control method on the chemical plume tracing (CPT) task, which consists of a robotic swarm tracing a toxic chemical cloud to the source emitter. Our experiments show that vehicle coordination is beneficial in terms of shortening the search time and increasing the overall success of source localization. The improvement in performance follows from the self-organizing capabilities of our framework. Dynamic computational meshes are constructed for real-time fluid analysis, using only local interactions between nearby vehicles. A brief overview of the control architecture is presented first, followed by a description of the plume tracing application.

2 Motivation

For the past several years, many military and civilian agencies have signaled an increased concern in protecting chemical manufacturing and storage facilities from acts of terror and sabotage [1, 2]. Conducting search and rescue operations in a contaminated area is both risky and costly; thus the use of robotic, autonomous units is seen as the logical approach to addressing this issue. Because chemical toxins usually propagate through a fluid medium, such as air and water, techniques are needed to both track the extent of the poison, and discover the location of the source. Mapping of the chemical plume is usually called plume *tracking*, while the localization of the emitter is known as plume *tracing*. This paper presents a new plume tracing algorithm for robotic swarms based on the dynamics of fluid flow.

Because the swarm will operate under hazardous conditions, to be successful in locating and possibly neutralizing the toxic source, the vehicles must cooperate in an autonomous, self-organizing manner. System design must minimize human involvement in the operation, and the required hardware should remain simple and inexpensive. The advantage of a multi-unit system is the resiliency and robustness due to the ability of the functioning vehicles to take over the responsibilities of their damaged neighbors; however, achieving this level of distributed control from simple robots with limited sensory and computational abilities is difficult.

Distributed vehicle control methods can be categorized into several schools of thought. Some techniques consist of modeling animal behavior, such as bird flocking [3]. Others design sophisticated rule structures to associate sensor readouts with vehicle actions [4, 5]. Control theoretic work has allowed for both formalization and optimization of certain swarm systems [6]. Besides these solutions, an approach based on physics modeling, e.g., [7], has emerged as an efficient way to coordinate very large teams of mobile robotic agents. Transcription of real physical properties and laws into virtual equivalents for artificial systems is the operating principle of the *physicomimetics*, or *artificial physics* (AP) framework [8]. The AP framework operates in the middle of the control hierarchy, between the low-level actuator control and the high-level task planning. Our design of a decentralized CPT system is one practical application of this approach.

Physical principles underlying AP make it possible to derive a comprehensive analysis of swarm components and their interaction by using standard physics analysis tools [9]. In a similar manner, we apply mathematical analyses to common chemical plume configurations to design and validate our CPT method [10]. This allows us to construct a theoretically sound approach to the problem of finding a source of a chemical cloud, one that outperforms heuristic methods in current use [11].

3 Related Work

The impressive olfactory performance of certain biological organisms, such as insects and crustaceans, motivated the design of most of the proposed CPT systems [12–14]. Biomimetic CPT methods frequently use key elements of animal behavior in an attempt to reproduce the plume tracing success of these natural plume tracers. The dominant plume tracing technique is called *chemotaxis*, and it consists of first measuring the local chemical gradient and then moving in the direction of increasing chemical concentration [15–17]. Chemotaxis is a simple, straightforward method with only modest hardware requirements. However, since the local chemical gradient becomes a poor indicator of the emitter's location with increasing distance away from the source, robots using chemotaxis are often misled by local concentration maxima [6, 18].

A biology-inspired CPT approach that relies more on locally sensed fluid flow is called *anemotaxis*. To execute this plume-tracing strategy the vehicle needs to be equipped with an anemometer capable of reporting the direction of airflow in the immediate vicinity of the vehicle, so that the robot can move upstream within the plume. This strategy frequently overcomes the problem of local chemical concentration maxima; however, regions with unstable or turbulent airflow create difficulties for navigation [14, 19, 20].

Simulations of fluid mechanics for robot navigation are reported by [21–23]. Findings by [24] indicate that full dissipation of indoor chemical plumes is a slow process, that occurs on the order of several hours. An odor recognition array capable of using wind sensors to discern the location of the source is described in [25]. The odor compass designed by [26] uses a small fan to improve sensor response to low chemical concentrations, and also can estimate the direction of the local chemical concentration peak.

4 Physicomimetics Framework

One crucial aspect of our work is the design of an automated, coordinated, multi-vehicle sensor network, which does not depend on vehicle size, e.g., vehicles could be nanobots, micro-air vehicles (MAVs), or micro-satellites. In our model, each robot observes the environment through its sensors, and then alters both itself and the environment through its effectors. Vehicle sensors can only detect and affect nearby vehicles, and a hallmark of our design is the use of only "local" control rules. A combination of local self-organization and fault-tolerance of the vehicle collective means that the global swarm behavior degrades slowly as individual units are disabled. Many commonplace physical systems already possess these abilities, and therefore physics-based approaches hold a great deal of promise.

The Artificial Physics (AP) control framework uses virtual forces to minimize the system's virtual potential energy. The system faithfully implements $F = ma$ dynamics, including forces, and therefore preserves all of the properties of Newtonian physics. Structured grid vehicle formations emerge without explicit global control. Although AP minimizes the potential energy of the entire system, it never needs the expensive calculation of global potential fields, and unlike explicit kinetic energy formulations [27], it uses inter-agent forces directly to control robot velocities. Also, AP does not aim to match the optimality offered by control-theoretic approaches [6,28]; however, this has not shown any detrimental effects on AP's suitability for the CPT task.

5 Physics of Chemical Plume Flow

The following brief summary of the physics behind our new CPT algorithm will provide the necessary background for understanding the new approach. The three familiar physical laws: conservation of mass, conservation of momentum, and conservation of energy govern the flow of fluids, and are consequently known as the *governing equations*. In a simple flow, the conservation of mass is the most relevant equation:

$$-\frac{\partial \rho}{\partial t} = \nabla \cdot (\rho \vec{V}) \tag{1}$$

where t denotes time, ρ is the mass density of the chemical, and $\vec{V}$ is its velocity (ρ and $\vec{V}$ are called the *flow-field variables*). The product $\rho\vec{V}$ measures the *mass flux*, or the rate of change of mass flow per unit area. The RHS of (1) is known as the *divergence of mass flux*, and assuming a 2D flow, it is:

$$\nabla \cdot (\rho \vec{V}) = u\frac{\partial \rho}{\partial x} + \rho\frac{\partial u}{\partial x} + v\frac{\partial \rho}{\partial y} + \rho\frac{\partial v}{\partial y} \tag{2}$$

Vector calculus indicates that if $\nabla \cdot (\rho\vec{V}) > 0$ at some point P, then the space at P is a *source* of mass flux. Likewise, $\nabla \cdot (\rho\vec{V}) < 0$ indicates a *sink* (i.e., a build-up) of mass flux. The role of mass flux in the CPT task becomes clear as the result of applying the Divergence theorem:

$$\int_W \nabla \cdot (\rho\vec{V}) \mathrm{d}W = \oint_S (\rho\vec{V}) \cdot \mathrm{d}\vec{S} \tag{3}$$

assuming that W is some control volume and S is the corresponding control surface. What this equation means in

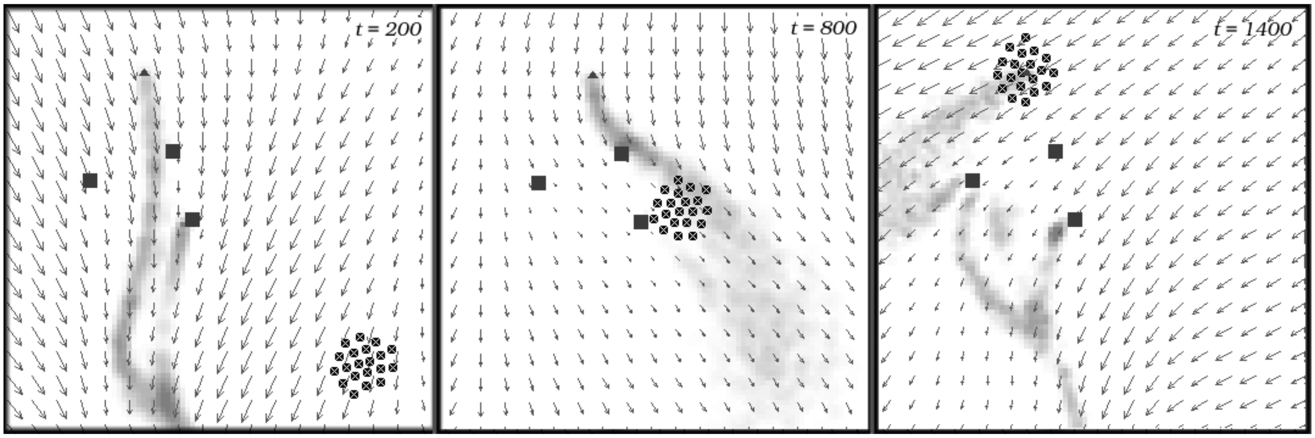

Figure 1. Fluxotaxis CPT simulation with a 20-vehicle lattice, three rectangular obstacles, and a meandering plume originating in the top left corner. The lattice starts out in the lower right corner and successfully traces the plume to its source (the triangle).

the CPT task, is that a sensor lattice enclosing the chemical source will measure a positive mass flux divergence. Hence (3) provides a method for *theoretically identifying a chemical emitter*, and it forms the foundation of our algorithm, which we call *fluxotaxis*. In order to be of practical use, we replace the continuous partial derivatives with corresponding discrete approximations, and use the AP-constructed swarm formation to measure the flow-field variables at several locations simultaneously. Therefore, the swarm can be viewed as a distributed computer for analyzing fluid flow and locating the source of the chemical.

6 CPT Algorithms

The CPT task has three subtasks: find the chemical, trace it to the source, and identify the source. These subtasks are elaborated next, starting with the search for the plume.

6.1 Finding the Chemical

The search for the chemical plume is accomplished via *casting*, which is implemented by treating a boundary of the tracing area as the "goal" location. Upon reaching the designated boundary, the vehicle selects another boundary, and continues this process until a plume is located. Each unit starts out in the casting mode, and may return to casting many times during a single CPT scenario as the result of losing the plume. A periodic exchange of simple binary state information between neighboring vehicles serves as the swarm coordination mechanism when locating the plume [4].

6.2 Tracing the Plume

Three emitter localization algorithms are implemented:

Chemotaxis follows the chemical gradient, or $\nabla\rho$. Chemotaxis is performed by a vehicle when it has at least one neighbor, and the accuracy of the gradient approximation improves as the number of neighbors increases. The swarm is attracted to high chemical concentrations within the plume.

Anemotaxis moves in the direction opposite the local airflow. A vehicle can obtain wind information even when it has no neighbors, but it must be inside the plume, as determined by the output of its chemical sensors. Swarms performing anemotaxis look for wind sources within the plume.

Fluxotaxis advances the swarm by seeking out points of large mass flux $\rho\vec{V}$. Additional strategies are available for regions where the flux is zero [11]. Conceptually, every group of three or more vehicles forms a virtual control volume, and the estimate for the divergence of mass flux from (2) is obtained through application of (3) to the mass flux measured on the virtual surfaces formed by the lattice sensor nodes. As can be seen, our fluxotaxis approach combines spatial information about both fluid velocity and chemical density into one plume-tracing method.

The CPT algorithm creates one of the virtual AP forces, which pulls the swarm toward the hypothesized location of the plume emitter. Each vehicle shares the local flow measurements with its neighbors, and then uses the CPT algorithm to independently calculate the next waypoint, which translates into a virtual "goal" force. Note that the presence of obstacles can sometimes alter the goal force via an incremental correction that will gently deflect the vehicle away from a collision. In order to calculate the next move, each robot combines its goal and formation-preserving forces, to compute the net driving force acting on the vehicle. The swarm collective thereby moves toward the goal while maintaining the lattice formation.

6.3 Source Emitter Identification

As shown by (3), locating a region with positive mass flux divergence (e.g., where the emitter is located) can be done by measuring the total mass flux exiting the region. This can be performed by a lattice of robots surrounding the conjectured location of the source emitter [10]. However, because anemo- and chemotaxis typically do not identify the

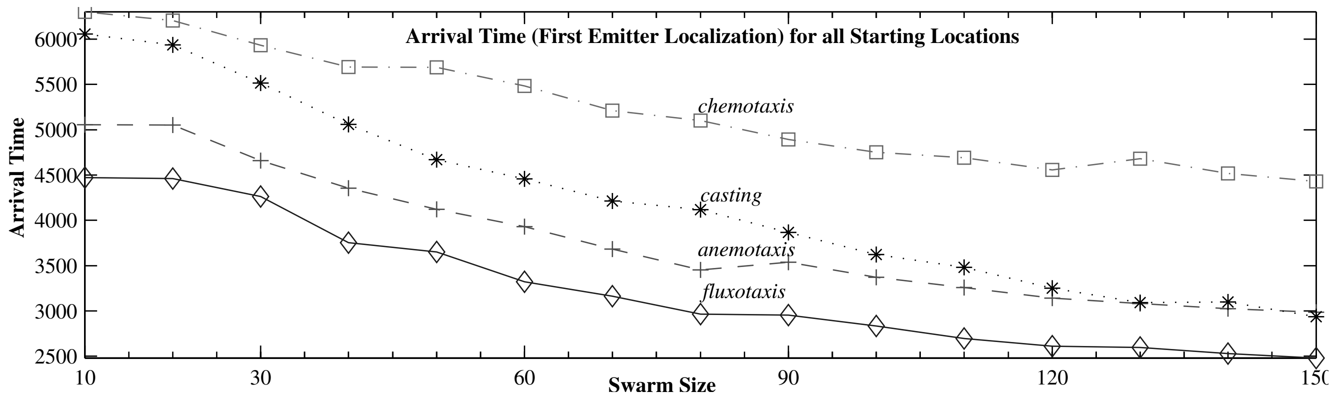

Figure 2. Arrival time metric for each CPT algorithm over a range of swarm sizes (smaller values indicate faster localization). Arrival event occurs when any swarm vehicle makes the first sensor contact with the emitter (determined by a global observer).

emitter, for the sake of fairness we effectively ignored this feature of fluxotaxis in the experiments presented here. This ablation from fluxotaxis also helped us focus on the purpose of the experiments – to analyze the gain in CPT performance obtained by methodically increasing the swarm size.

7 Simulation Experiments

Experiments detailed in this paper use a fluid solver of Farrell, et al. [14]. The simulator creates meandering sinuous plumes composed of localized chemical filaments or "puffs". Each puff has a template shape that grows in size and disperses over a large area with time, thus causing the entire plume to dissipate at locations far away from the emitter (see Fig. 1). Periodic boundary conditions with a variable gain produce many physically distinct flow configurations. We extended the original model through additional equations to accommodate rectangular obstacles. The resulting simulation environment is computationally efficient and incorporates realistic multi-scale properties – including intra-plume filament mixing and advective transport of the chemical.

We compared chemotaxis, anemotaxis and fluxotaxis on a suite of simulated plume scenarios. Results for swarms that only executed the casting strategy are included as the baseline. Ten different flow conditions were chosen, each containing a dynamic chemical-gas plume evolving over a 90,000 sq. ft. region, with a mix of laminar, transitional, and turbulent flows. Each plume tracing area contained ten $5' \times 5'$ randomly placed obstacles. Because toxin contamination typically precedes the robot activity, a chemical was ejected for 3600 simulations steps (an hour of real plume time) before the swarm was first deployed. We then advanced each plume for 7200 steps (corresponding to a realistic two-hour time frame) and collected CPT performance statistics for each CPT algorithm. To evaluate CPT performance we identified two metrics as being most relevant: the amount of elapsed time before the first emitter detection (the "arrival time" metric), and the length of time the swarm maintained sensor contact with the emitter (the "localization frequency" metric). Since only fluxotaxis offers a fluid-theoretic method for emitter identification, in order to compare all three CPT algorithms, we assumed that another emitter detection method is available [20, 24].

Simulated robots have a dimension of 1 sq. ft. and a maximum speed of 3 in/sec, closely modeling the equipment in our laboratory. To match the available hardware, vehicle communication and sensor ranges are limited to 3 feet, with the desired inter-robot lattice spacing of 20 inches. All robots are equipped with an identical suite of chemical detectors and anemometers. Each CPT algorithm is evaluated on a matching set of simulation parameters consisting of 15 swarm sizes (ranging from 10 to 150 vehicles), and 30 random starting locations per swarm, uniformly distributed in the emitter-centered polar coordinate system, with distances ranging from 0 to 200 feet away from the emitter.

7.1 Results and Analysis

Results of our experiments for different swarm sizes are shown in Figures 2 and 3, with the averages computed in Table 1. We note that larger swarms find the emitter faster, and localize the source more frequently, supporting our conjecture that CPT is a swarm application that benefits from the multiplicity and mobility of the sensing units.

Another observation is that fluxotaxis, which we derived from the physics of fluid flow, consistently outperformed the biomimetic approaches. The general shape of the fluxotaxis performance curve mirrors those of chemotaxis

Table 1. Average performance of the three CPT algorithms over 30 starting locations, 15 swarm sizes, and 10 plumes.

Algorithm	*Arrival Time* (lower is better)	*Localization Frequency* (higher is better)
Casting	4224.4	224.1
Chemotaxis	5208.5	1496.5
Anemotaxis	3780.5	1983.6
Fluxotaxis	3249.9	2384.8

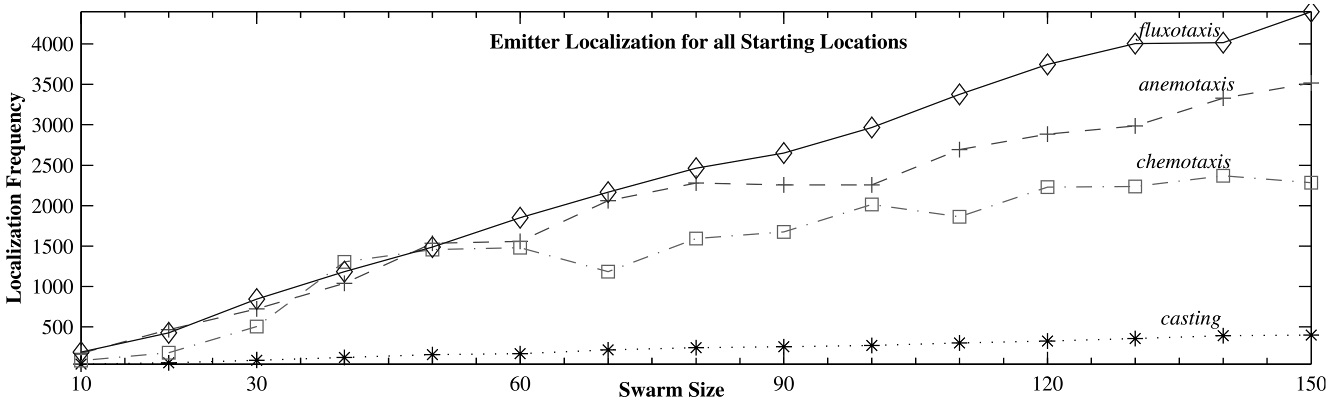

Figure 3. Count of cumulative chemical emitter localizations by each CPT algorithm. The emitter is considered to have been localized when a swarm vehicle makes sensor contact with the emitter, as determined by a global evaluation algorithm.

and anemotaxis, which is expected, since fluxotaxis incorporates both the ρ and $\vec{V}$ elements of these two techniques. Mathematically, fluxotaxis is a second order method, while the other approaches are of the first order [10], allowing our method to outperform the more heuristic CPT algorithms.

Of the CPT techniques, note that chemotaxis is most easily led astray. Even the naïve, uninformed casting strategy yields better arrival time results than chemotaxis. The chemotaxis-driven swarm is frequently misled by the local concentration maxima around obstacles, lowering its performance. However, in the vicinity of the emitter chemotaxis performs better, since the observed peaks in the concentration landscape are in fact close to the real source.

With small swarms, anemotaxis outperforms simple casting in terms of the arrival time, but this advantage vanishes with a larger swarm size. This is due to the fact that instantaneous wind velocity is generally a poor indicator of the true direction of the source. As fluid flow periodically changes its direction, the anemotaxis strategy is misled by the isolated pockets of a fragmented plume, resulting in a time-inefficient zigzag plume traversal pattern often observed in moths, lobsters, and other organisms [12, 14, 19]. This is an important observation, since it shows that just adding more sensors is not enough: *the sensory input must be managed in an informed and theory-supported fashion*. Given the decentralized swarm architecture, each team member must provide the means to facilitate this emergent sensor processing ability. Our fluxotaxis algorithm provides a practical example of how such functionality can be engineered into a complex system, such that the implementation provides a mathematically formal assurance of the long-term performance, and also guarantees efficient scalability and low operational cost.

8 Summary and Future Work

In this paper we showed that the use of the artificial physics control framework, in combination with a physics-based chemical plume tracing algorithm resulted in a theoretically derived CPT solution. Our method outperformed the popular biomimetic alternatives on a large suite of realistic plume simulations on two metrics that measured both the speed and the accuracy of the CPT swarm. Using the AP framework, a swarm of simple autonomous vehicles constructs a large adaptive sensor grid in a decentralized fashion, using limited-range virtual forces between the neighboring vehicles and the environment. The chief contribution of this work is the demonstration of how the emergent behavior of the AP-controlled swarm can be used to solve a challenging problem. We illustrated how a physics approach to the CPT task produced a novel plume tracing algorithm, and yielded a new theoretic measure for chemical emitter identification.

Continuation of this work will improve fluid simulations and sensor models (e.g., recovery time, noise). Our current goal is to port the CPT simulation to a laboratory setting. Our UW Distributed Robotics Laboratory created seven prototypes that demonstrated the ability to form hexagonal arrangements, avoid obstacles, and locate a source of light [9]. We are adding chemical sensors and anemometers to the hardware platform, which will allow us to test the fluxotaxis algorithm with volatile organic compound gas plumes.

References

[1] Board on Atmospheric Sciences and Climate. *Tracking and Predicting the Atmospheric Dispersion of Hazardous Material Releases: Implications for Homeland Security*. National Academies Press, 2003.

[2] S. L. Caldwell et al. Combating terrorism: Federal agencies' efforts to implement national policy and strategy, 1997. Congressional report GAO/NSIAD-97-254.

[3] T. Balch and R. Arkin. Behavior-based formation control for multi-robot teams. In *IEEE Transactions on Robotics and Automation*, volume 14, 1998.

[4] S. T. Kazadi, E. Kondo, and A. Cheng. A robust centralized linear spatial search flock. In *Proceedings of*

the IASTED International Conference on Robotics and Applications, pages 52–59, 2004.

[5] A. Howard, M. J. Matarić, and G. S. Sukhatme. An incremental self-deployment algorithm for mobile sensor networks. *Autonomous Robots*, 13(2):113–126, 2002.

[6] J. T. Feddema, R. D. Robinett, and R. H. Byrne. An optimization approach to distributed controls of multiple robot vehicles. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'03)*, 2003.

[7] M. Polycarpou, Y. Yang, and K. Passino. Cooperative control of distributed multi-agent systems. *IEEE Control Systems Magazine*, 2001.

[8] W. Spears and D. Gordon. Using artificial physics to control agents. In *Proceedings of the IEEE Conference on Information, Intelligence, and Systems (ICIIS'99)*, 1999.

[9] W. Spears, D. Spears, J. Hamann, and R. Heil. Distributed, physics-based control of swarms of vehicles. In *Autonomous Robots*, volume 17(2-3), August 2004.

[10] D. Zarzhitsky, D. Spears, D. Thayer, and W. Spears. Agent-based chemical plume tracing using fluid dynamics. In *Lecture Notes in Computer Science*, volume 3228. Springer-Verlag, 2004.

[11] D. Spears, D. Zarzhitsky, and D. Thayer. Multi-robot chemical plume tracing. In *Proceedings of the Third International Workshop on Multi-Robot Systems*, 2005.

[12] F. W. Grasso. Invertebrate-inspired sensory-motor systems and autonomous, olfactory-guided exploration. *Biol. Bull.*, 200:160–168, 2001.

[13] M. Koehl et al. Lobster sniffing: Antennule design and hydrodynamic filtering of information in an odor plume. *Science*, 294:1948–1951, 2001.

[14] J. Farrell, J. Murlis, W. Li, and R. Carde. Filament-based atmospheric dispersion model to achieve short time-scale structure of odor plumes. In *Environmental Fluid Mechanics*, volume 2. Kluwer, 2002.

[15] G. Sandini, G. Lucarini, and M. Varoli. Gradient driven self-organizing systems. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'93)*, 1993.

[16] J. Crimaldi, M. Koehl, and J. Koseff. Effects of the resolution and kinematics of olfactory appendages on the interception of chemical signals in a turbulent odor plume. In E. Cowen, editor, *Environmental Fluid Mechanics*, volume 2, pages 35–63. Kluwer, 2002.

[17] J. E. Hurtado, R. D. Robinett, C. R. Dohrmann, and S. Y. Goldsmith. Distributed sensing and cooperating control for swarms of robotic vehicles. In *Proceedings of the 1998 International Conference on Control and Applications (IASTED'98)*, 1998.

[18] X. Cui, C. Hardin, R. Ragade, and A. Elmaghraby. A swarm-based fuzzy logic control mobile sensor network for hazardous contaminants localization. In *Proceedings of the IEEE International Conference on Mobile Ad-hoc and Sensor Systems (MASS'04)*, 2004.

[19] A. Hayes, A. Martinoli, and R. Goodman. Swarm robotic odor localization. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'01)*, 2001.

[20] S. T. Kazadi. Extension of plume tracking behavior to robot swarms. In *Proceedings of the SCI2003 Conference*, 2003.

[21] J. Decuyper and D. Keymeulen. A reactive robot navigation system based on a fluid dynamics metaphor. In H.-P. Schwefel and R. Männer, editors, *Parallel Problem Solving from Nature (PPSN I)*, volume 496 of *Lecture Notes in Computer Science*. Springer-Verlag, 1991.

[22] D. Keymeulen and J. Decuyper. The fluid dynamics applied to mobile robot motion: the stream field method. In *Proceedings of the 1994 International Conference on Robotics and Automation (ICRA'94)*, 1994.

[23] D. Keymeulen and J. Decuyper. The stream field method applied to mobile robot navigation: a topological perspective. In A. G. Cohn, editor, *Proceedings of the Eleventh European Conference on Artificial Intelligence (ECAI'94)*, pages 699–703, 1994.

[24] M. Wandel, A. Lilienthal, T. Duckett, U. Weimar, and A. Zell. Gas distribution in unventilated indoor environments inspected by a mobile robot. In *Proceedings of the IEEE International Conference on Advanced Robotics (ICAR'03)*, 2003.

[25] R. A. Russell and A. H. Purnamadjaja. Odor and airflow: Complementary senses for a humanoid robot. In *Proceedings of the 2002 IEEE International Conference on Robotics and Automation*, 2002.

[26] H. Ishida, T. Nakamoto, T. Moriizumi, T. Kikas, and J. Janata. Plume-tracking robots: A new application of chemical sensors. *Biol. Bull.*, 200:222–226, April 2001.

[27] C. Belta and R. V. Kumar. Trajectory design for formations of robots by kinetic energy shaping. In *Proceedings of the 2002 IEEE International Conference on Robotics and Automation (ICRA'02)*, 2002.

[28] J. A. Fax and R. M. Murray. Information flow and cooperative control of vehicle formations. In *Proceedings of the 2002 IFAC World Congress*.