

SWARMS FOR CHEMICAL PLUME TRACING

Dimitri Zarzhitsky and Diana F. Spears and William M. Spears

Department of Computer Science
University of Wyoming
Laramie, WY 82071
wspears@cs.uwo.edu

ABSTRACT

This paper presents a physics-based framework for managing distributed sensor networks of autonomous vehicles, e.g., robots, which self-organize into structured lattice arrangements using only local information. The vehicles remain in formation during obstacle avoidance and search for a chemical emitter that is actively ejecting a toxic chemical into the air. We discuss a new plume tracing algorithm, based on the principles of fluid physics, that outperforms the leading biomimetic competitors for this task.

1. INTRODUCTION

The focus of our research is to design and build rapidly deployable, scalable, adaptive, efficient, and robust swarms of autonomous distributed vehicles. This combines sensing, computation and networking with mobility, thus enabling deployment, self-organization, and reconfiguration of the swarm. Our objective is to provide a scientific, yet practical, approach to the design and analysis of swarm systems.

The general purpose for deploying tens to hundreds of such vehicles can be summarized as “volumetric control.” Volumetric control means monitoring, tracking, reporting, and responding to environmental conditions within a specified physical region. This is done in a distributed manner by deploying numerous vehicles, each carrying one or more sensors, to collect, aggregate, and fuse distributed data into a tactical assessment, yielding an enhanced situational awareness and the potential for a more timely response.

The chemical plume tracing (CPT) task involves detection of a toxic chemical and localization of the source emitter. This paper addresses both aspects of the problem, by first presenting an efficient and robust framework for distributed control of simple autonomous vehicles, and then demonstrating how the framework is exercised to construct a mobile sensor network for measuring important fluid properties. An evaluation of the overall system performance for common plume tracing scenarios in simulation is given.

2. MOTIVATION

The objective of the chemical plume tracing task is rapid localization of the emitter that is ejecting the chemical. Plume *tracing* is different from plume *tracking* – the objective of the latter is to acquire a map of the plume. The underlying motivation for CPT is typically to minimize human exposure to the toxin; therefore, the use of autonomous robots seems especially appropriate. It is difficult to predict ahead of time the conditions that the plume tracing agents will encounter at the site of the contamination – it may be an urban setting after a deliberate terror act, or an industrial facility after a chemical accident. Preparedness for such events is a concern for government and commercial sectors [3, 5, 14].

Self-organization, autonomous operation and error recovery, along with consistently reliable performance are the chief design objectives for an autonomous plume tracing system. The use of multiple robots to accomplish the task may reduce the time it takes to localize the emitter; however, massively parallel systems require efficient, scalable control mechanisms in order to realize the benefits of multiplicity.

Common approaches to managing such distributed autonomous systems, and robotic *swarms* in particular, include those that use ad-hoc and loosely connected structures, model specific behaviors, follow a set of rules, or employ control-theoretic methods based on optimization of system equations [7, 10, 12]. Over the course of the last decade, a physics-inspired methodology has emerged as a successful alternative for achieving complex aggregate behavior from a group of simple robots, e.g., [21]. This is the approach embodied by the *physicomimetics*, or *artificial physics* (AP) framework [25]. The term “artificial” emphasizes the fact that although we are motivated by real physical forces, we are not restricted to emulation of commonplace physics, but have the freedom to develop our own force laws to achieve our objective. The AP framework is the foundation for our new CPT approach.

Since the AP control framework is based on physics, it is amenable to physics-based analyses [27]. The chemical plume tracing approach is likewise analyzable in both

application and design using fluid physics [30]. In particular, this paper demonstrates that by basing our system design upon physicomimetics for formation control and upon fluid dynamics for perceptual guidance, our approach to the CPT problem can be characterized as “theoretically sound.” Most importantly, this principled foundation has resulted in substantial performance improvements over alternative, more heuristic approaches [24].

3. PRIOR APPROACHES

Many current CPT solutions are inspired by the biological olfactory systems of lobsters and moths [9, 11, 20]. Since any system that traces a chemical plume must have the ability to sense the concentration of the trace chemical, the most widely utilized approach is that of *chemotaxis*, which follows a local gradient of the chemical concentration within a plume [6, 15, 23]. Chemotaxis is simple to understand and implement, but it may lead to locations in the plume that are far away from the true source [7, 10].

Another popular biomimetic approach is *anemotaxis*. An anemotaxis-driven agent measures the direction of the fluid’s velocity and navigates “upstream” within the plume [9, 12, 17]. Such a strategy is successful for problems where the flow has no large-scale turbulence and remains mostly laminar for the duration of the plume tracing activity, but in general, the fluid can have large turbulent eddies that curl and circulate, creating a region where upstream movement causes anemotaxis to fail.

Other approaches derive from a wide variety of disciplines. Application of statistical analysis to a CPT problem is reported by [2], and concerns the development of an algorithm that estimates the probability of the source’s location. A prototype radar network capable of operating under real-world conditions is presented in [13] as a method for tracking and localizing a target. Unlike our approach, theirs assumes fixed sensors and some degree of centralization. Early use of the fluid dynamic metaphor within a robotic system [8, 18, 19] adopted a simplified model of fluid flow in simulation to navigate a single robot in a dynamic environment. Our approach uses similar concepts, but provides the theory applicable to a large collection of robots operating within a real plume [24, 30].

Wandel, *et al.* [29] performed laboratory plume mapping experiments with a single mobile robot, and found that the chemical density profiles are spatially stable over significant lengths of time. Russell and Purnamadajaja [22] developed an odor recognition system consisting of an array of chemical and airflow sensors capable of identifying the type of a trace chemical, and estimating the direction of the source. Both of these systems used a small suction fan to move the air past the stationary chemical sensors, similar in principle to the odor compass developed by Ishida, *et al.* [16].

4. PHYSICOMIMETICS FRAMEWORK

Recall that one of our objectives is the design of fully automated, coordinated, multi-vehicle sensor systems. The team vehicles could vary widely in type, as well as size, e.g., from nanobots or micro-electromechanical systems (MEMS) to micro-air vehicles (MAVs) and micro-satellites. To date, our implementation has been on ground-based robotic vehicles. Nevertheless, the approach is easily extendable to other types of vehicles. A vehicle’s sensors perceive the environment, including other vehicles, and a vehicle’s effectors affect the vehicle and its environment. It is assumed that vehicles can only sense and affect nearby vehicles; thus, a key challenge has been to design “local” control rules. Not only do we want the desired global behavior to emerge from the local interaction between vehicles (self-organization), but we also require fault-tolerance; that is, the global behavior degrades very gradually if individual vehicles are damaged. Self-repair is also desirable, in the event of damage. Self-organization, fault-tolerance, and self-repair are precisely those principles exhibited by natural physical systems. Thus, answers to the problems of distributed control can be found in physics.

AP is designed on physics principles. In AP, virtual physics forces drive a multi-robot system to a desired configuration or state. The desired configuration is one that minimizes virtual system potential energy. Unlike behavior-based approaches [1], AP uses an $F = ma$ dynamics simulation, hence its forces are such that the system obeys Newtonian physics and is amenable to Newtonian analyses. AP has been used for self-assembly and self-repair of robotic lattices, where grid formations are accomplished via virtual physics forces. Although the inter-vehicle forces are virtual, the robots *act* as if they were real. Thus, the vehicle’s sensors must see enough to allow it to compute the force to which it is reacting, and the vehicle’s effectors must allow it to respond to this perceived force.

To construct mobile sensing grids, we define a force law $F = G/r^p$, where F is the magnitude of the force between two robots, r is the range between the two robots, and p is some user-defined power. The “gravitational constant” G is set at initialization. The desired separation between robots is R , and the force is repulsive if $r \leq R$ and attractive if $r > R$. By varying R , the lattice can expand and contract. The only effector is vehicle movement with velocity \vec{V} . To ensure that the force laws are local in nature, robots have a restricted sensor and communication range. Also, due to the discrete-time nature of the model, we define a maximum system force F_{\max} . The force law works well for constructing several types of vehicle formations, including hexagonal lattices [25, 27, 28], which were shown [4] to have superior computational characteristics for important types of fluid flow boundary conditions.



Figure 1. Seven AP robots self-organize and move in a hexagonal formation toward a light source. Each robot is autonomous.

Depending on the value of G , it is possible for robots to cluster at individual nodes in the lattice, which may not be desired. To prevent clustering of the robots, a force balance analysis indicates that G must be less than $F_{\max}R^p/(2\sqrt{3})$.

Let us assume that each robot senses a goal direction. We cannot assume that such sensing will always be accurate. Hence, on occasion robots may attempt to move in different directions toward their incorrectly sensed goal. Furthermore, if one or more robots are temporarily halted in their movement (due to environmental or hardware problems), we would like the formation to maintain its cohesion. We can again use a “force balance” analysis to show that the goal force F_{goal} must be less than $\sqrt{3}G/R^p$ for lattice cohesion to be maintained.

The current laboratory implementation of the physicomimetics framework on a team of robots consists of seven vehicles. Our first experimental objective was to form a stable hexagon that moves toward a light source. Each robot ran the same piece of software, and could detect only the range and bearing to neighboring robots. The desired distance, R , between robots was 20 inches, $p = 2$, and $F_{\max} = 2$. Using our theory [27], $G < 308$ will prevent clustering. Also, F_{goal} needs to be less than 1.33, and we used 1.0 for our experiments. The results are shown in Fig. 1, and were consistent over ten runs, maintaining the formation and never showing clustering.

Vehicle sensors exhibit noise and their readings can be occluded by the other robots. If the cohesion of the formation were weak, individual robots could wander in different directions. However, under strong cohesion the formation acts as a solid geometrical object. Hence, in the aggregate, sensor noise is automatically minimized and the formation moves in the direction that the majority agrees on. It is important to note that this “majority vote” is implicit – it is an emergent property of the collective. In summary, due to the virtual cohesive forces holding the lattice together, the entire lattice will move in the goal direction determined by the majority (without explicit voting) [26–28, 30], regardless of the nature of the goal. In our first experiment, the search

goal was a source of light. In the following section, we consider a different goal – that of finding the source of a toxic plume.

5. PHYSICS OF CHEMICAL PLUME FLOW

Our CPT algorithm relies on methods and concepts from computational fluid dynamics, so we begin with a short review of the relevant material. Fluid flow is governed by three fundamental laws: conservation of mass, conservation of momentum, and conservation of energy. These laws are in turn expressed by the Governing Equations; for example, the law of conservation of mass is written in a differential form as:

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

where ρ is the mass density of the chemical, \vec{V} is the fluid’s velocity (ρ and \vec{V} are called the *flow-field variables*), and t denotes time. In practice, an analytical solution of the Governing Equations is impossible to obtain due to the inherent non-linearity. However, we can replace the continuous partial derivatives with discretized finite-difference approximations, and solve for the flow-field variables using a sensor grid. Our CPT algorithm uses the lattice formations, built with physicomimetics, as distributed sensor networks. These networks act as a parallel computer for performing fluid flow analyses to assist in the making of navigational decisions.

Note that the product $\rho \vec{V}$ is called the *mass flux*, or the rate of change of mass flow per unit area. The RHS of (1) is the *divergence of mass flux*, and in 2D it is given by:

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

where $\vec{V} = u\hat{i} + v\hat{j}$ and \hat{i} and \hat{j} are unit vectors in the x and y coordinate directions respectively. If at some spatial point location P , $\nabla \cdot (\rho \vec{V}) > 0$, then point P is a *source* of mass flux, while $\nabla \cdot (\rho \vec{V}) < 0$ indicates a *sink* of mass flux. The




Figure 2. Fluxotaxis plume tracing sequence: heavy chemical concentration is shown with darker colors; the black dots are the robots, which start out in the bottom left corner, and successfully find the emitter in the top right corner. Arrows depict wind velocity.

importance of mass flux in the chemical plume tracing task comes from the Divergence theorem of vector calculus:

$$\int_W \nabla \cdot (\rho \vec{V}) dW = \oint_S (\rho \vec{V}) \cdot d\vec{S} \quad (3)$$

This equation, where W is the control volume and S is the bounding surface of the volume, allows us to formally define the intuitive notion that a control volume containing the emitting source will have a positive mass flux divergence, while a control volume containing a sink (a type of CPT “trap”) will have a negative mass flux divergence. In particular, (3) shows that if the robot lattice surrounds a suspected emitter, and the total mass flux exiting the circle of robots consistently exceeds some small, empirically-determined threshold, then the robots have localized the true source. This result serves as the important criterion for *theoretically identifying a chemical emitter*. To the best of our knowledge, previous criteria for emitter identification are purely heuristic [7, 9, 12]. Our algorithm, which we appropriately named *fluxotaxis*, is the first to apply principles of fluid mechanics within a CPT solution.

6. CHEMICAL PLUME TRACING ALGORITHMS

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

6.1. Tracing the Plume

Three tracing algorithms have been implemented: chemotaxis, anemotaxis, and fluxotaxis. We implemented all three with AP as the foundation for vehicle control. For all algorithms, robots share flow-field variables with their neighbors and use these values to calculate the next waypoint. These flow-field measurements are mapped to the robots’ own local coordinate axes. Based on these values, each

robot independently decides the best direction to move, which translates into a virtual force. Finally, each robot takes the resultant vector of the plume-following force and the lattice-preserving force to compute the driving force experienced by each robot. The emergent effect is to move the collective toward the goal while preserving the lattice. The plume-following force acts as a force in the goal direction. The nature of this force distinguishes the three algorithms.

Chemotaxis. The chemotaxis approach simply follows the chemical gradient, or $\nabla \rho$. For chemotaxis, the direction of the largest chemical concentration is the goal direction.

Anemotaxis. The anemotaxis strategy always moves upstream while detecting an above-threshold level of chemical concentration. “Upstream” (also called “upwind”) is the goal direction for the robots.

Fluxotaxis. Divergence of mass flux (2) forms the basis of our novel fluxotaxis algorithm. With fluxotaxis, the robotic lattice computes the local divergence of mass flux, and then follows its gradient (the direction of steepest increase). For flow regimes with zero flux divergence, fluxotaxis has additional strategies, as described in [24]. Mathematically, the lattice acts like a distributed sensor grid, and each individual robot independently calculates the gradient:

$$\nabla(\nabla \cdot \rho \vec{V}) = \nabla \left(u \frac{\partial \rho}{\partial x} + \rho \frac{\partial u}{\partial x} + v \frac{\partial \rho}{\partial y} + \rho \frac{\partial v}{\partial y} \right) \quad (4)$$

From the equation above it is clear that fluxotaxis combines information about both fluid velocity and chemical density, and the fact that it also measures mass flux from (3) provides assurance that it will find the source of the chemical [24, 30].

Fluxotaxis addresses all subtasks of the CPT problem; it is the only algorithm that performs source identification.

6.2. Finding the Chemical

The most common method for finding the chemical is called *casting*; it typically consists of a sweeping, zigzag, or spiraling motion to increase exploration [9, 12]. The use of




Figure 3. Chemical plume changes due to obstacles (dark rectangles): images show flows with identical initial conditions.

several (rather than two) collaborating robots improves traditional casting, and in our experiments, all three CPT algorithms use such casting.

6.3. Source Emitter Identification

The Divergence theorem (3) can be used by fluxotaxis to unambiguously identify a chemical emitter from local flux measurements. One can logically infer from (3) that to identify a region with positive mass flux divergence, it is sufficient to measure the total flux exiting the region. This can be done with a hexagonal lattice of robots surrounding the conjectured source emitter. Because anemo- and chemotaxis typically do not have a way to identify an 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 evaluate the scalability and robustness of the physicomimetics framework under a range of environmental constraints.

7. SIMULATION EXPERIMENTS

The experimental results reported in this paper use a fluid solver of Farrell *et al.* [9].¹ It is especially well-suited to our needs because of its efficiency, realism (i.e., its instantaneous and time-averaged results match measurements of actual plumes), and multi-scale properties – including chemical diffusion and advective transportation. Rather than a continuous, time-averaged model, Farrell models the chemical plume as a filament-based collection of discrete and localized “puffs.”

We compared random casting, chemotaxis, anemotaxis, and fluxotaxis on a suite of 81 simulated, realistic plume scenarios. Nine physically distinct flow configurations were selected, each containing a dynamic chemical-gas plume evolving over a 90,000 sq. ft. region, with an equal mix

of laminar, transitional, and turbulent flows (see Fig. 2). The effect of obstacles on the CPT performance of the AP-driven robotic swarms of varying sizes was evaluated by solving fluid dynamic equations over randomly created environments with: no obstacles, with nine, eighteen, twenty-seven, thirty-six $5' \times 5'$ obstacles, and with two, four, seven, and nine $10' \times 10'$ obstacles. The entire experimental test suite is therefore a combination of each flow configuration and obstacle course; Fig. 3 illustrates the impact obstacles may have on the development of a plume.

Because contamination typically precedes CPT efforts, the trace chemical was ejected for 3600 simulation steps (about an hour of real plume time) before a swarm of CPT robots was first deployed. The initial swarm starting location varied from precisely over the emitter to 200 ft. away from the emitter in 4 ft. increments. The number of vehicles in the swarm varied from 7 to 70 robots, with a total of 10 different swarm sizes per plume, obstacle, and initial starting location (a total of 40,500 CPT evaluation runs were performed). We advanced each plume for 7200 steps (corresponding to a realistic two hour time frame [29]), and recorded the time of the first sensor contact with the emitter by any robot in the swarm (the *arrival time* metric), in addition to keeping track of the total number of emitter detections by the entire swarm during a CPT scenario (the *localization frequency* metric).

In simulation, each robot is approximately one foot in diameter and has a maximum speed of 3 in/sec, closely modeling the equipment in our laboratory. To match the hardware communication range of our lab platform, the sensor and communication range is limited to 40 inches. Obstacle avoidance is implemented using a hierarchical architecture, with the CPT algorithms usually navigating around obstacles before collision avoidance becomes necessary. The obstacle avoidance problem is simpler within a plume, since the lattice can always select a new bearing when blocked by an obstacle. We confirm the observations of [8, 18, 19], that using direction of fluid flow is an efficient and successful strategy for navigating out of maze-like areas. This is an

¹Periodic boundary conditions with variable gain produce different sinusoidal flows; we extended the original equations to incorporate obstacles.

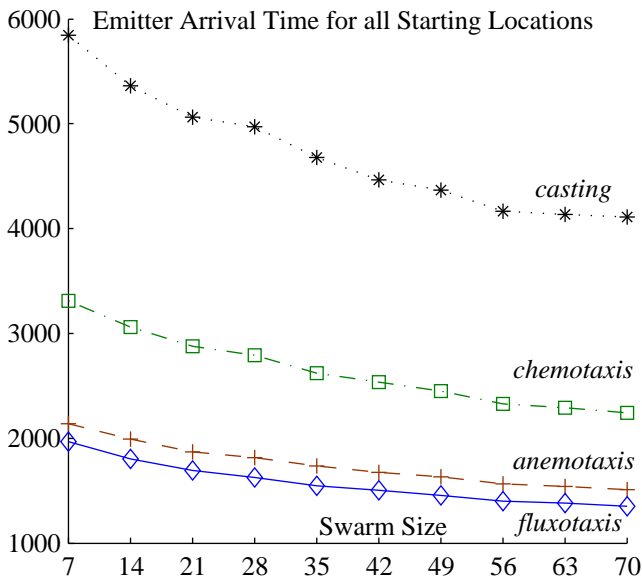


Figure 4. Arrival time results for each CPT algorithm averaged over 81 plumes and 50 starting locations. Lower values indicate better performance.

important finding, since CPT system deployment may occur in a hazardous area, littered with debris or industrial machinery.

7.1. Empirical Results

The results in this paper focus on two CPT performance metrics: the elapsed trace time before the first emitter detection (see Fig. 4 and Table 1), and the number of times the emitter is successfully localized by the members of the swarm (see Fig. 5 and Table 2). The first observation from the two charts and tables is that only fluxotaxis performs well on *both* metrics; it combines the speed performance demonstrated by anemotaxis, with the localization frequency advantage shown by chemotaxis, illustrating a robust, theory-guided approach to solving a complex problem by intelligently exploiting the available knowledge of the underlying

Table 1. Arrival time metric of each CPT algorithm versus the world obstacle coverage, averaged over 10 swarm sizes, 50 starting locations, and 81 plumes.

Algorithm	Obstacle Coverage				
	0.0%	0.25%	0.50%	0.75%	1.0%
Fluxotaxis	1508	1523	1615	1572	1610
Anemotaxis	1885	1760	1725	1700	1731
Chemotaxis	2451	2603	2887	2437	2778
Casting	5292	4848	4739	4498	4490

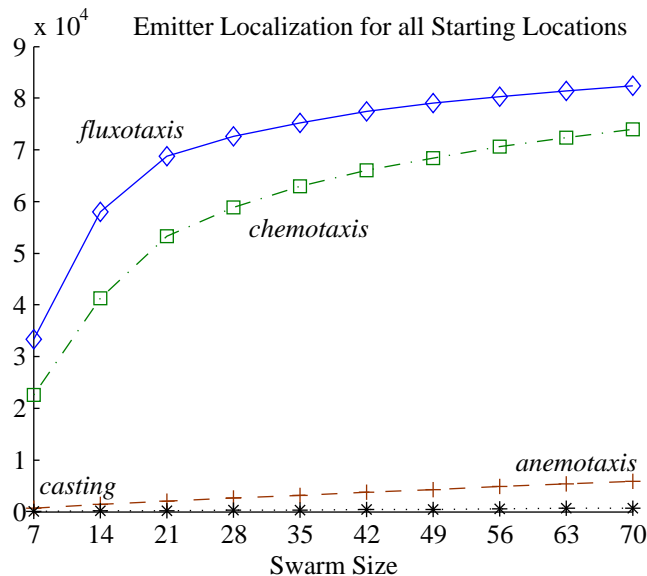


Figure 5. Frequency of emitter localization by the swarm, averaged over 81 plumes and 50 starting locations. Larger values indicate higher performance.

physical process.

This clever utilization of the local sensor data is an emergent property of the swarm which arises due to the lattice formation of the agents, facilitated by the AP framework. By sharing the information about local flow conditions between neighboring vehicles in the group, each team member is able to construct a more accurate view of the surrounding plume. Furthermore, through the use of insight from fluid mechanics, the set of vehicle waypoints selected by fluxotaxis is shown to be consistently more effective in achieving the CPT goals than the tracing paths taken by the more heuristic algorithms.

Figure 5 highlights another important finding: CPT is inherently a swarm application, because of the large gain in performance achieved with even a modest increase in the swarm size. However, careful management of the swarm is necessary in order to realize the performance benefit, as

Table 2. Cumulative count of emitter detections by the swarm for each obstacle coverage category, averaged over 10 swarm sizes, 50 starting locations, and 81 plumes.

Algorithm	Obstacle Coverage				
	0.0%	0.25%	0.50%	0.75%	1.0%
Fluxotaxis	73253	72189	70085	70657	69369
Chemotaxis	62925	59398	55802	62649	56368
Anemotaxis	3604	3667	3479	3429	2939
Casting	368	308	330	593	271

can be seen from the extremely low localization success rate achieved with the simple random casting algorithm, even with large swarm sizes. A clue to the difference in localization performance between chemotaxis and anemotaxis lies in the fact that anemotaxis makes little use of sensor information available from its neighbors, and each vehicle simply navigates against the local current. Chemotaxis however, computes the local *gradient* of the chemical concentration, thus gaining more information about the behavior of the plume in the vicinity of each vehicle, maximizing the benefits of the collaboration within the robotic lattice.

Data in Tables 1 and 2 is the evidence for the robustness and scalability of the physicomimetics framework, and its ability to adapt to a changing environment. The force law parameters were the same in each experiment, regardless of the number of agents, or size and the number of obstacles; however, note that the resulting performance of each algorithm is largely unaffected by the obstacles, and variations in the actual performance of each algorithm are due to the natural characteristics of each CPT approach. For instance, the arrival time of the random casting strategy shortens with increasing obstacle coverage because greater obstacle coverage implies less open area that the random searcher has to explore. Likewise, the ability of the chemotaxis strategy to maintain sensor contact with the emitter is somewhat inhibited in the environments with many obstacles, as each obstacle induces a local chemical sink in its vicinity, thus misleading this naïve gradient follower.

8. SUMMARY AND FUTURE WORK

This paper presented a physics-based approach to the chemical emitter localization problem using a distributed system of simple autonomous robots that self-organize into a mobile, adaptive sensor network. The control framework, called “artificial physics” (AP), is efficient, scalable, robust, and amenable to theoretical analysis and predictions. The chemical plume tracing task is a complex swarm application, which we address using a collection of simple robots, with emergent behavior under the AP control.

Our contribution to this area is the development of a new plume-tracing algorithm, called fluxotaxis, that is based on the theoretical principles of fluid flow and utilizes AP-built lattices. We showed that our theoretically-founded fluxotaxis algorithm is able to improve *both* emitter localization time and detection frequency over the two most popular alternatives, for a large set of different flows, plumes, and tracing environments.

Further algorithm development in simulation will include advanced turbulence models, online learning of thresholds, an increased number of obstacles, and modeling of sensor characteristics (e.g., size, number, noise). An important near-term focus will be on porting the simulation to

actual robots in a laboratory plume emission setting. Our UW Distributed Robotics Laboratory built a team of several small robots that have successfully demonstrated self-organization into lattices, robust obstacle avoidance, and goal seeking in formation [27]. The next step is to integrate chemical sensors with the robot processors and test the fluxotaxis algorithm with emissions of volatile organic compound gases. In the long term, we plan to transition to outdoor robots, including heterogeneous teams of ground-based and micro-air vehicle platforms.

9. REFERENCES

- [1] T. Balch and R. Arkin. Behavior-based formation control for multi-robot teams. In *IEEE Transactions on Robotics and Automation*, volume 14, pages 1–15, 1998.
- [2] E. Balkovsky and B. Shraiman. Olfactory search at high reynolds number. *PNAS*, 99(20):12589–12593, October 2002.
- [3] S. L. Caldwell et al. Combating terrorism: Federal agencies’ efforts to implement national policy and strategy, September 1997. Congressional report GAO/NSIAD-97-254.
- [4] E. S. Carlson, H. Sun, D. H. Smith, and J. Zhang. *Third Order Accuracy of the 4-Point Hexagonal Net Grid. Finite Difference Scheme for Solving the 2D Helmholtz Equation*, 2003. Technical Report No. 379-03, Department of Computer Science, University of Kentucky.
- [5] A. H. Cordesman. Defending America: Asymmetric and terrorist attacks with chemical weapons, September 2001.
- [6] 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.
- [7] X. Cui, C. T. Hardin, R. K. Ragade, and A. S. 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.
- [8] 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*, pages 356–362. Springer-Verlag, 1991.
- [9] J. A. Farrell, J. Murlis, W. Li, and R. T. Carde. Filament-based atmospheric dispersion model to achieve short time-scale structure of odor plumes. In E. Cowen, editor, *Environmental Fluid Mechanics*, volume 2. Kluwer, 2002.
- [10] 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.
- [11] F. W. Grasso. Invertebrate-inspired sensory-motor systems and autonomous, olfactory-guided exploration. *Biol. Bull.*, 200:160–168, April 2001.
- [12] 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.
- [13] B. Horling et al. Distributed sensor network for real time tracking. In *Proceedings of the International Conference on Autonomous Agents (AGENTS'01)*, pages 417–424, 2001.
- [14] S. S. Hsu. Sensors may track terror's fallout. *Washington Post*, page A01, June 2003.
- [15] 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.
- [16] 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.
- [17] S. T. Kazadi. Extension of plume tracking behavior to robot swarms. In *Proceedings of the SCI2003 Conference*, 2003.
- [18] 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)*, volume 4, pages 378–385. IEEE Computer Society Press, 1994.
- [19] 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. John Wiley & Sons, 1994.
- [20] M. Koehl et al. Lobster sniffing: Antennule design and hydrodynamic filtering of information in an odor plume. *Science*, 294:1948–1951, 2001.
- [21] M. Polycarpou, Y. Yang, and K. Passino. Cooperative control of distributed multi-agent systems. *IEEE Control Systems Magazine*, June 2001.
- [22] R. A. Russell and A. H. Purnamadajaja. Odor and airflow: Complementary senses for a humanoid robot. In *Proceedings of the 2002 IEEE International Conference on Robotics and Automation*, pages 1842–1847, 2002.
- [23] 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.
- [24] D. Spears, D. Zarzhitsky, and D. Thayer. Multi-robot chemical plume tracing. In *Proceedings of the Third International Workshop on Multi-Robot Systems*, 2005. (in press).
- [25] 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.
- [26] W. Spears, R. Heil, D. Spears, and D. Zarzhitsky. Physicomimetics for mobile robot formations. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multi Agent Systems (AA-MAS'04)*, volume 3, pages 1528–1529, 2004.
- [27] 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.
- [28] W. Spears, D. Spears, R. Heil, W. Kerr, and S. Hettiarachichi. An overview of physicomimetics. In *Lecture Notes in Computer Science - State of the Art Series*, volume 3342. Springer-Verlag, 2005.
- [29] 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.
- [30] 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.