

Distributed Robotics Approach to Chemical Plume Tracing

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

Department of Computer Science

University of Wyoming

Laramie, WY 82071

dimzar@uwyo.edu

Abstract—This paper presents an application of a physics-based framework for distributed control of autonomous vehicles. The autonomous swarm uses local information to self-organize into dynamic sensing and computation grids during localization of the source of a toxic plume. Using physics of fluid flow we develop a new plume-tracing algorithm, and then use computational fluid dynamics simulations to show that the new approach outperforms the leading biomimetic competitors for this task.

Index Terms—plume tracing, sensor networks, swarms.

I. INTRODUCTION

We are interested in scalable, efficient, and robust swarms of autonomous vehicles that combine sensing, computation, and networking with mobility, thus enabling deployment, self-organization, and reconfiguration of the swarm. One important application for such systems is distributed environment monitoring, where tens to hundreds of mobile units collect and fuse sensor data into a comprehensive assessment, thus improving situational awareness and reducing response time. The chemical plume tracing (CPT) task concerns detection of a toxic plume and subsequent localization of the source emitter. Distributed vehicle control is done via the artificial physics (AP) framework, which is efficient, scalable, robust, and subject to rigorous theoretical analysis and predictions [19]. We use the self-organizing capabilities of AP to construct a mobile sensor grid that finds the toxic source by processing important plume observations. The paper focuses on an evaluation of three plume-tracing algorithms on a large suite of realistic flow simulations, including a challenging set of obstacle courses, and shows that a theoretically-guided use of fluid mechanics improves swarm performance on the CPT task.

II. MOTIVATION

The objective of the CPT task is rapid localization of the chemical emitter. Plume-tracing conditions are difficult to predict – it may be an urban setting after a deliberate terror act, or an industrial facility after a chemical accident. Our multi-agent CPT approach is founded upon the AP control framework – because this framework is physics-based and is therefore easily amenable to physics-based analyses [19]. The design and application of our CPT solution is likewise analyzable with fluid physics [21]. In particular, this paper demonstrates that by using AP for formation control, and fluid dynamics for perceptual guidance, our comprehensive approach to the CPT problem is theoretically sound.

III. PRIOR APPROACHES

Common methods for managing 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 [3], [6], [8]. 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., [14]. This is the approach embodied by the *physicomimetics*, or *artificial physics* (AP) framework [18]. 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. AP is a distributed control framework for applying virtual physics forces to construct multi-robot formations. Potential energy is minimized, but without the expense of global potential fields calculations [12]. Unlike explicit kinetic energy formulations [1], AP uses inter-agent forces directly. Even though AP does not attempt to match the optimality offered by control-theoretic approaches [5], [6], its performance is excellent [19].

CPT solutions typically derive from the biological olfactory systems of lobsters and moths [4], [7], [13]. Since a CPT system must measure 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 [2], [9], [16]. Another popular biomimetic approach is *anemotaxis*. An anemotaxis agent observes the direction of fluid flow, and navigates upstream inside the plume [8], [11].

Wandel, et al. [20] 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 [15] 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. [10].

IV. CHEMICAL MASS FLUX

Our new CPT method relies heavily on ideas borrowed from the field of computational fluid dynamics, so this section provides 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. The law of mass conservation is most relevant to our current discussion; it can be 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. Our CPT algorithm uses the lattice formations, built with AP, as a distributed sensor and computation network, that acts as a parallel computer for performing fluid flow analyses to assist in the making of navigational decisions.

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 called the *divergence of mass flux*, and in 2D is given by:

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

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

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

where W is the control volume and S is the bounding surface of the volume. This theorem formally defines the intuitive notion that a control volume containing the emitter (i.e., the source) will have a positive mass flux divergence, while a control volume containing a sink (potentially a CPT “trap”) will have a negative mass flux divergence. In particular, (3) shows that if the robots surround a suspected emitter, and the total mass flux measured by the the sensor grid consistently exceeds some small, empirically-determined threshold, then the robots have localized the emitter. This result serves as the important criterion for *theoretically identifying a chemical emitter*. To the best of our knowledge, previous criteria for CPT emitter identification are mostly heuristic [3], [4], [8]. Our algorithm, which we appropriately named *fluxotaxis*, is unique in its application of fluid mechanics in the context of a mobile swarm driven by local plume observations.

V. 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 search for the plume.

A. Finding the Chemical

The most common method for finding the chemical plume is called *casting*; it typically consists of a zigzag or spiraling motion to increase exploration [4], [8]. The use of several (rather than two) collaborating robots improves traditional casting, and in our experiments, all three CPT algorithms use such casting. If the swarm is large, the emergent behavior of the formation is to fragment while moving past obstacles or passing through narrow openings or disjoint regions. Because our framework is completely decentralized, each subgroup of the original robot collective continues to operate independently of its former neighbors.

B. Tracing the Plume

Three CPT algorithms have been implemented: chemotaxis, anemotaxis, and fluxotaxis. Each one uses AP to arrange the vehicles into a hexagonal formation, so that robots can share flow-field variables with their neighbors, and use these values to calculate the next waypoint. Each robot maps the shared sensor data to its own coordinate axes, independently calculates the best direction to search next, and then translates it into a virtual goal force. The final driving force for each vehicle consists of a vector sum of the plume-force and lattice-preserving AP-force vectors. The emergent effect is that the entire collective moves in formation toward the goal. Conceptually, each robot continually executes the following:

```

while emitter-not-found() do
  foreach neighbor in scan() do
    range = compute-range(neighbor);
    bearing = compute-bearing(neighbor);
    data = get-plume-data(neighbor);
    force += (AP-force(range, bearing) +
              plume-force(data));
  end foreach
  velocity += force * deltaT; // F = ma
end while

```

The three CPT algorithms are distinguished by the implementation of the plume-following force, *plume-force*, as follows:

CHEMOTAXIS. The chemotaxis approach simply follows the chemical gradient, or $\nabla \rho$, so that for chemotaxis, the direction of the largest chemical concentration is the goal direction. In our implementation, this direction is estimated by seeking out a neighbor with the largest detected chemical concentration, and computing the bearing toward this neighbor. If a vehicle has no neighbors, it executes the casting procedure.

ANEMOTAXIS. The anemotaxis strategy always moves upstream while inside the chemical plume. “Upstream” (also called “upwind”) is the goal direction for the robots. Even though anemotaxis agents do not share flow data, they still coordinate their actions implicitly via the formation force. Vehicles performing anemotaxis will follow the casting procedure whenever the output of the chemical or wind sensors indicates the loss of the plume or a region of stagnant air flow.

FLUXOTAXIS. The concept of mass flux, explained above, forms the basis of our novel fluxotaxis algorithm. With fluxotaxis, the robotic lattice computes the amount of chemical flux passing through virtual surfaces formed by neighboring swarm vehicles. Mathematically, the lattice acts like a distributed sensor grid, and each individual robot independently calculates the amount of local chemical flux, $\rho \vec{V}$, relative to the current position of its neighbors. For flow regimes with zero mass flux, fluxotaxis employs additional strategies, as described in [17], designed to maximize the use of available sensor data. Because fluxotaxis combines information about both fluid velocity and chemical density, and since it can also estimate the total mass flux from (3), we can show that it will find the true chemical source in a number of common CPT configurations [17], [21].

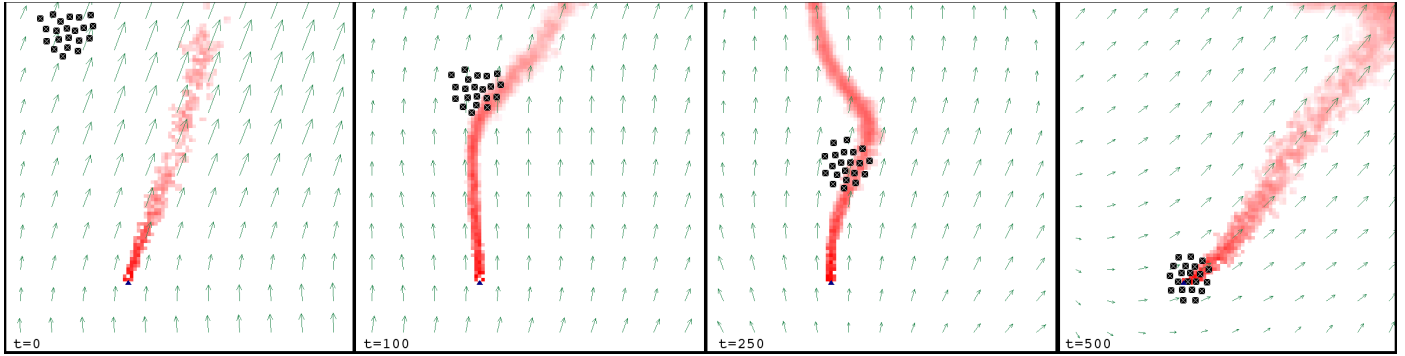


Fig. 1. A sample, simulated fluxotaxis plume tracing sequence: heavy chemical concentration is shown with darker colors; the collection of black dots is the swarm, which starts out in the top left corner, and successfully finds the emitter, denoted by a triangle on the bottom. Arrows depict wind velocity.

C. Source Emitter Identification

The Divergence theorem 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 ring of robots surrounding the conjectured source emitter, and exactly this mechanism provides a straightforward implementation of the `emitter-not-found()` procedure for swarms performing fluxotaxis. However, standard implementations of anemo- and chemotaxis lack a method for identifying the source. In the empirical studies reported in this paper, in order to standardize the testing conditions, we effectively ignored the emitter identification ability of fluxotaxis.

VI. SIMULATION EXPERIMENTS

Empirical studies detailed in this paper fall into two major categories. First, we evaluate a set of simulated scenarios designed to accurately model the hardware in our laboratory, and then we increase the number of agents in simulation to investigate the scalability and robustness of our approach. In the first set of experiments, a hexagonal ring formation of seven robots is used, coupled with a success metric that is well-suited to such a ring formation, namely, the *frequency of enclosure* of the source emitter by the ring. In the second set of experiments, a swarm of robots of increasing size is used. Because we are interested in studying the effects of additional agents on CPT performance, a simple enclosure metric is no longer adequate. Instead, we keep a track of the number of sensor contacts with the emitter made by all of the vehicles, called the *localization frequency*. The second set of experiments also uses an additional performance metric, namely, the amount of time it takes to locate the emitter, called the *time of first detection* by any robot. This second metric illuminates the relationship between swarm size and CPT speed.

A. Experimental Model of the Plume

All of the reported experimental results rely on a fluid solver¹ by Farrell et al. [4], chosen 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 using a continuous, time-averaged model, our chemical plume is discrete, consisting of a collection of localized “puffs” of the trace element (see Fig. 1).

B. Experimental Models of the Robots

In the first set of experiments, all of the plume tracing algorithms employ a hexagonal lattice of seven robots, with one in the center of the formation. To compensate for the small number of vehicles, the CPT algorithms can expand and contract the radius of the formation to improve the spatial resolution of plume sensors. This swarm behavior is accomplished within the AP framework via a local change of the neighbor separation for each vehicle. Given the small size of the swarm, the AP parameters are configured to produce a strong formation cohesion force, so that a solid-like lattice formation is preserved during navigation around obstacles.

The second experimental set is oriented toward the investigation of the scalability characteristics of our physics-based approach. We still use the same vehicle model as before, but allow the swarm size to increase. This change also requires us to modify the CPT algorithms slightly, as the larger swarm size dictates the need for more agile formations in order to pass through narrow passages and satisfy other environmental constraints. The navigation planning routines execute independently on each vehicle, resulting in a completely distributed control, fluid-like swarm motion, and increased effectiveness in negotiating around complex multi-obstacle geometries. These experiments demonstrate that the self-organization and autonomous recovery capabilities of the AP framework translate directly into a strong and consistent improvement in the swarm performance on the CPT task.

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

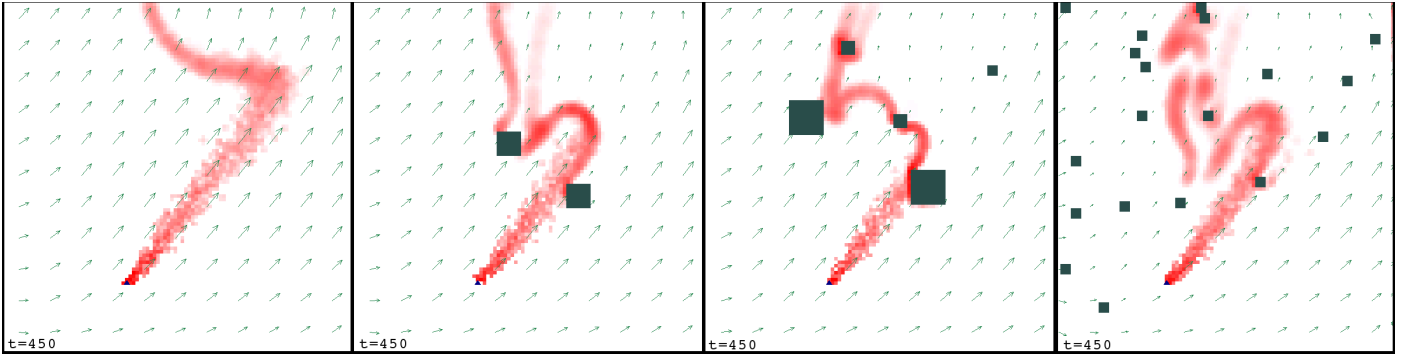


Fig. 2. Changes in the chemical plume due to obstacles: identical flow parameters are visualized in each image (solid black squares are the obstacles).

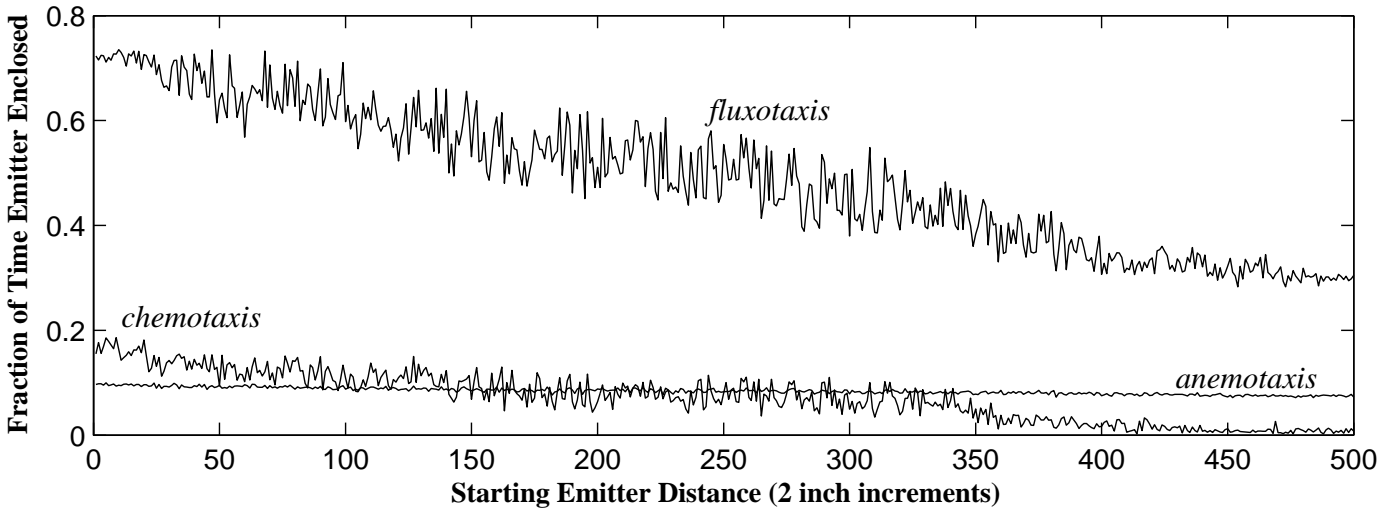


Fig. 3. Emitter enclosure results for seven robots, averaged over 150 plume configurations. The x axis is the distance between the initial starting location of the lattice and the emitter, and the y axis is the fraction of total CPT time (7200 simulation steps) spent by the lattice surrounding the emitter.

C. Experiment Set 1: Simulation of Lab Vehicles

To evaluate the CPT performance of a model swarm based on our existing hardware prototype, we selected different flow conditions containing a dynamic, airborne plume spread over a 10,000 sq. ft. region. The suite of plumes was divided equally between laminar, transitional, and turbulent flows. In order to evaluate the CPT algorithms under realistic conditions, fluid dynamic equations were resolved over randomly generated environments with: no obstacles, one $10' \times 10'$ obstacle, two $10' \times 10'$ obstacles, four $5' \times 5'$ obstacles, and eight $5' \times 5'$ obstacles. Flow parameters were identical for each environment, but the plumes differed because of the obstacles (see Fig. 2). Typically, environment contamination precedes search and containment operations; hence, in our simulation the chemical was ejected for 3600 simulation steps (about an hour of real plume time) prior to lattice deployment. We then advanced the plume for 7200 steps (corresponding to a realistic two hour CPT run [20]), and computed the percentage of time that the chemical source was surrounded by the seven-agent lattice, i.e., the enclosure frequency. This metric allows us to capture the stability and robustness characteristics of each CPT

algorithm. Initial distance of the lattice from the emitter, which varied from 0 to 80 feet, served as the independent variable for these evaluations, with a total of 500 randomly selected starting locations tested in each flow and obstacle configuration.

In earlier work [17], we demonstrated that the AP framework is robust for a large range of lattice radii. In this case, the maximum formation radius for fluxotaxis was set at ten feet, to match the scale of plume features. Based on prior results indicating that chemo- and anemotaxis performance is more consistent with a fixed-radius formation, the AP force law for these two algorithms executed with the desired inter-vehicle spacing set at ten feet. Each robot navigated via a set of common waypoints spaced two feet apart, and was equipped with a set of chemical and wind sensors. The simulation of the chemical sensor accounted for the recovery period, which for a typical metal oxide sensor, such as Figaro's TGS 2620, is less than a minute [3].

1) *Enclosure Results:* The outcome of enclosure tests is plotted in Fig. 3. The amount of time the lattice spent encircling the emitter is an important evaluation criterion because it measures how precise and consistent a CPT algorithm is

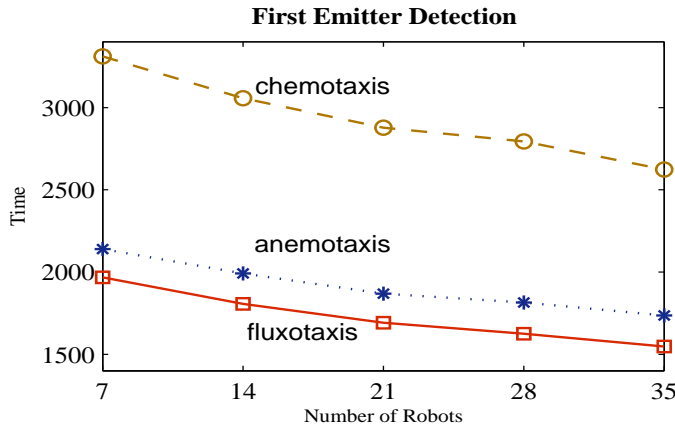


Fig. 4. Arrival time results for each CPT algorithm averaged over 81 plumes and 50 starting locations. Lower values indicate better performance.

at determining the location of the plume’s stationary source. To measure this criterion, we used a global observer (strictly during the evaluation stage), and computed the time fraction that the robotic lattice spent surrounding the chemical emitter.

From the plot, note that fluxotaxis (mean=0.491) significantly² outperformed chemotaxis (mean=0.072) and anemotaxis (mean=0.084). Since anemotaxis always moves upwind in the plume, it often overshoots the emitter, and spends 92% of its CPT time in a cycle of first moving upstream in the plume and then casting after missing the emitter. Anemotaxis performance is low even when it begins near the emitter, and the performance curve slowly decreases with increasing distance. Chemotaxis does somewhat better when it starts near the emitter, but its performance falls off rapidly, caused by the increase in the number of obstacle-induced local density maxima. However, even when chemotaxis begins its tracing within ten feet of the emitter (mean=0.143), the lattice still fails to contain the source 86% of the time. This is due to the turbulence in the flow, and the resulting periodic chemical shedding effect.

D. Experiment Set 2: Swarm Simulations

The set of experiments presented in this section focuses on the impact that increasing the swarm size has on CPT performance. The experimental methodology in this second investigation is consistent with that of the first study: we compare chemotaxis, anemotaxis, and fluxotaxis on a suite of 81 simulated, realistic plume scenarios. Nine distinct flow configurations were selected, each containing a chemical-gas plume evolving over a larger 90,000 sq. ft. region, with a representative combination of laminar, transitional, and turbulent flows. Robotic swarms of varying sizes, controlled by the AP framework, navigated through a test suite of randomly created environments with no obstacles, with nine, eighteen, twenty-seven, thirty-six $5' \times 5'$ obstacles, and with two, four, seven,

²A Wilcoxon rank sum test applied to each experiment showed that the CPT performance improvement obtained with fluxotaxis is statistically significant in 94% of the experiments at $p = 0.01$ and in 96% at $p = 0.05$ levels of significance.

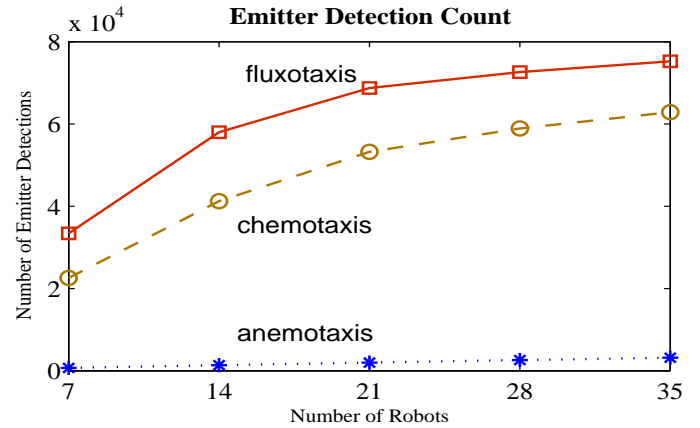


Fig. 5. Emitter detection by swarms of different size, averaged over 81 plumes and 50 starting locations. Larger values indicate higher performance.

and nine $10' \times 10'$ obstacles. The CPT performance scores are reported as averages over the combination of each flow configuration and obstacle course, using swarm size as the independent variable.

Each simulated robot is the analogue of our laboratory platform, which is approximately one foot in diameter, with the maximum speed of 3 in/sec. The sensor and communication range of simulated agents is limited to 40 inches, matching the corresponding limitations of the hardware. Consistent with our previous simulations, we allow for 3600 steps of initial chemical contamination before deploying the robots. The starting location of the swarm 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 35 robots, with 5 different swarm sizes per plume, obstacle, and initial starting location (a total of 20,250 CPT evaluation runs were performed for this study). The evolution of each plume spanned 7200 time steps, and for each CPT algorithm we recorded the time of the first sensor contact with the emitter by any robot in the swarm (the *time of first detection* metric), and the total number of emitter detections by the whole swarm during the entire CPT scenario (the *localization frequency* metric).

1) *Scalability Results:* The amount of time each CPT algorithm spent searching for the emitter is plotted in Fig. 4. The count of successful source localizations for each swarm size is shown in Fig. 5. The first observation from the two plots is that only fluxotaxis performs well on *both* metrics; it combines the speed performance demonstrated by anemotaxis, with the good emitter detection frequency of chemotaxis. Fluxotaxis is a robust, efficient, theory-based solution for a complex problem, that uses the available information about the underlying physical process to consistently improve the swarm’s performance.

The intelligent utilization of data from the distributed sensor nodes is an emergent property of the swarm, arising from the lattice arrangement of the agents, which in turn is 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 application of fluid mechanics, the set of vehicle waypoints selected by fluxotaxis is shown to be more effective in finding the emitter than the paths chosen by the more heuristic algorithms.

Figure 5 highlights another important finding: CPT is inherently a swarm application, because of the gain in performance achieved with an increase in the swarm size. A clue to the difference in localization performance between chemotaxis and anemotaxis lies in the fact that anemotaxis does not share plume observations amongst neighboring vehicles, but simply navigates upstream in the local airflow. 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, and benefiting from the multiplicity of the sensor nodes.

VII. 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 chemical plume tracing task is a complex application implemented using a collection of simple robots, with emergent behavior under AP control. Our contribution to this area is the development of a new plume-tracing algorithm, called fluxotaxis, that is based on theoretical principles of fluid flow and utilizes AP-constructed lattices. We showed that our theoretically-founded fluxotaxis algorithm is able to demonstrate statistically and practically significant gains in localization performance over the two most popular alternatives, even in an environment with obstacles.

We showed that a challenging real-world problem, characterized by a complicated set of non-linear interactions, can in fact be solved in a fully distributed manner by using simple, inexpensive mobile sensor platforms. The effectiveness of our solution comes from the firm foundation in Newtonian physics. Our ability to carry out accurate mathematical analysis of the system resulted in the design and construction of a *predictable* autonomous swarm, with an exceptionally low number of control parameters. Furthermore, by extending the physical model from vehicle formation control to sensor information management, we created a novel search and rescue system that provides an innovative and theoretically sound way to identify the location of a chemical source. Our experimental work also produced performance metrics suitable for validating swarm-based CPT algorithms.

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 [19]. 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.

REFERENCES

- [1] 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)*, pages 2593–2598, 2002.
- [2] 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.
- [3] 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.
- [4] 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.
- [5] J. A. Fax and R. M. Murray. Information flow and cooperative control of vehicle formations. In *Proceedings of the 2002 IFAC World Congress*.
- [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] F. W. Grasso. Invertebrate-inspired sensory-motor systems and autonomous, olfactory-guided exploration. *Biol. Bull.*, 200:160–168, 2001.
- [8] 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.
- [9] 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.
- [10] 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.
- [11] S. T. Kazadi. Extension of plume tracking behavior to robot swarms. In *Proceedings of the SCI2003 Conference*, 2003.
- [12] O. Khatib. Real time obstacle avoidance for manipulators and mobile robots. In *International Journal of Robotics Research*, volume 5, pages 90–98, 1986.
- [13] M. Koehl et al. Lobster sniffing: Antennule design and hydrodynamic filtering of information in an odor plume. *Science*, 294:1948–1951, 2001.
- [14] M. Polycarpou, Y. Yang, and K. Passino. Cooperative control of distributed multi-agent systems. *IEEE Control Systems Magazine*, 2001.
- [15] 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*, 2002.
- [16] 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.
- [17] D. Spears, D. Zarzhitsky, and D. Thayer. Multi-robot chemical plume tracing. In *Proceedings of the Third International Workshop on Multi-Robot Systems*, 2005.
- [18] 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.
- [19] 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.
- [20] 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.
- [21] 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.