

variables, such as fluid velocity and chemical concentration. This paper consists of a formal study of the effectiveness of our novel algorithm, including comparisons with the two most popular alternatives on which it is built. To supplement the discussion of the underlying theory, we include results from software simulations that implement the theoretical scenarios we present, and include realistic elements of measurement discretization.

2 Motivation

The authors' goal is to design a search algorithm that scales well to a large number of robots, ranging perhaps from ten agents to a thousand and beyond. In order to achieve this goal, two things are necessary: a formal theory upon which the algorithm is based, and a suitable task that can be used to test the algorithm. The task of chemical plume tracing has posed problems for a number of years in a variety of manufacturing and military applications. In light of the current national concern with security and the possibility of a chemical terrorist attack, several private and government agencies have expressed interest in updating current techniques used to track hazardous plumes, and improving the search strategies used to locate the toxin emitter [3, 4, 5, 6].

Because the physicomimetics framework relies on application of virtual forces to construct and maintain the robotic lattice, physics is the natural choice for the theoretical foundation of our work. In particular, the well-studied field of fluid physics and mechanics is well-suited for the development and validation of our algorithms.

There is another advantage of using a physics-based foundation. Computational fluid mechanics requires computational meshes for sampling and processing of flow-field variable values. The lattice arrangements that emerge naturally from the physicomimetics framework can be used as computational meshes, thus forming a massively parallel system, capable of performing complex computations in real time, with the added benefit of resilience to failure, and ability to adjust when the environment characteristics change. The natural synergy between the different system components translates directly into an improved performance of the system. For instance, the construction of hexagonal formations requires the least amount of communication and sensor information within the agent control framework [7]; at the same time, a hexagonal lattice was shown [8] to have superior boundary characteristics for solving an important class of fluid mechanics problems.

3 Related Work

Current research in the field has been inspired by biological olfactory systems of lobsters and moths [9, 10, 11, 12]. The base requirement for any system that attempts to trace a chemical plume is of course the ability to sense the presence of the chemical agent, as well as its concentration. The best understood and most widely applied approach is that of *chemotaxis*, which consists of following

a local gradient of the chemical concentration within a plume [13, 14, 11]. While chemotaxis is very simple to perform, it frequently leads to locations of high concentration in the plume that are not the source, such as a corner of a room. Furthermore, we have a proof, which we omit here due to space limitations, that a chemotaxis search strategy is likely to fail near the emitter’s location, due to the fact that for a typical time-varying Gaussian plume density profile, the gradient goes to zero near the distribution’s peak.

To overcome this problem, another common approach, called *anemotaxis*, has been developed. An anemotaxis-driven agent measures the direction of the fluid’s velocity and navigates “upstream” within the plume [15, 14]. Such a strategy is successful in problems where the flow has no large-scale turbulence. In general, we do not have the luxury of assuming this type of airflow. On the contrary, the airflow could have large turbulent eddies that curl and circulate, thus creating a region where traveling upwind will result in a cycle, causing the anemotaxis technique to fail.

Early results from applying the solution of fluid dynamic problems to robotic systems are reported by Keymeulen and Decuyper [16, 17, 18]. In this work, a highly simplified model of fluid flow was used successfully in simulation to navigate a single robot in a semi-dynamic environment; the approach was inspired by the fact that fluid flow is a good model of the iterative, local-to-global route finding task optimization, since the local pressure fields that are responsible for the existence of the stable optimal path are void of local minima. In the development of their approach, Keymeulen and Decuyper relied on the concepts of a fluid *source* and *sink*, which they used to specify the robot’s initial and goal locations. In the present work, we also base our method’s development on these two concepts, and extensively utilize both mathematical and physical properties of these two entities in the verification of our algorithm.

Work by Balkovsky and Shraiman [19] on the subject of statistical analysis of the plume is also relevant. They develop a probability density function having a Gaussian form, and use it to develop a simplified model of the chemical plume, which is then traversed using an algorithm that takes the probability of the source’s location into account. In the development of their algorithm, several assumptions were made regarding the type of the flow that the agent is expected to search. In our work we do not assume a particular flow-field, but rather establish several general categories of fluid flow and prove mathematically that our algorithm performs well in these broad and important categories.

Research by Parunak and Brueckner [11] makes a case for analysis of the self-organization property in multi-agent systems from the standpoint of entropy and the Second Law of Thermodynamics. They develop an analogy between entropy in the context of a system’s energetic quality and informational disorder, and show how understanding and management of system entropy can be used to analyze a multi-agent system. They illustrate the idea by solving an agent coordination problem with the use of simulated randomly-diffusing pheromones. Our work complements their thermodynamic approach by looking at the conserva-

tion properties of matter, and improves it by providing a more realistic model of information flow within a system.

A promising approach to tracking and localizing a target with soft real-time constraints is discussed in Horling, et. al. [20]. The major contribution of their work is a radar network capable of operating under real-world conditions with realistic restrictions of noisy communication channels, limited sensory capabilities, and restricted computational power. The system however, only allows for fixed sensors and makes use of partially centralized sector and target manager agents, introducing local points of failure. In our approach, decisions are made in a fully decentralized manner, improving robustness of the entire system. In addition, our framework places no restriction on mobility of either the plume or the tracking agents.

Also of interest is the work of Polycarpou et al. [21], where the notion of artificial potential fields is used to find the goal object (an attractor) while avoiding obstacles (repellents). In order to apply potential fields, they create a map of the environment and the agents then are able to compute virtual forces based on the knowledge of the environment. However, such global maps are costly to build and mapping errors are a significant problem. The strategy we are proposing does not require environment mapping, and works well with the local information obtained in a highly distributed manner by the agents.

4 Computational Fluid Dynamics

Our approach makes use of the methods and concepts developed in the context of computational fluid dynamics (CFD), so a brief review of the relevant material will be useful. Flow of fluids is governed by three fundamental laws: the conservation of mass, conservation of momentum (Newton's Second Law), and the conservation of energy [22, 23]. There is also an equation that captures turbulent effects [24], but for simplicity we omit it here. Collectively, these equations are known as the Governing Equations. Equations that describe theoretical inviscid flows are also known as the Euler equations, while the more complex real viscous flows are described by the Navier-Stokes equations. These equations come in several forms, but we will focus on the conservation form, which is based on the time analysis of a differential volume spatially fixed in the flow field [23]. For instance, the simplest equation, the conservation of mass, is written as

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

Here, ρ denotes the mass density of the chemical, \mathbf{V} is the fluid's velocity (collectively, ρ and \mathbf{V} are known as the *flow-field variables*), and t denotes time. For any real flow of practical interest, an analytical solution of the Governing Equations is impossible to obtain, due to the inherent non-linearity of the fluid dynamic systems. Thus, one CFD approach replaces the continuous partial derivatives with the corresponding discretized finite-difference approximations, and computes the unknown flow-field variables using a computational grid which

spans the region of interest. Our algorithm takes advantage of the lattice formations formed by our robotic agents to simulate the computational grid, thereby allowing the agents to perform a sophisticated analysis of the flow and make navigational decisions based on this analysis.

Other discretization methods, of which *finite-volume* and *finite-element* are best known, are also applicable to the AP-driven robotic lattices. However, in this paper, we only make use of the finite-difference discretization method because of its simple derivation from the Taylor-series expansion of partial derivatives [22]. For brevity and greater focus, we also ignore the interesting problem of boundary conditions, and focus on a theoretically limitless domain. Since we are interested in the problem of emitter localization, this simplification does not have a significant impact on the solution, as long as the region in which plume tracing is performed does not have walls nor obstacles. This limitation will be addressed in the later stages of our research.

The work presented in the following sections deals with the development of our physics-based solution to the chemical plume tracing task. It assumes a lattice of mobile agents with a limited, local view of the plume. The early theoretical results have been verified in simulation, and more complex flow configurations are currently being investigated.

5 Our Fluxotaxis Algorithm

The RHS of (1) represents the divergence of mass flux within the differential volume. Divergence plays a key role in the proposed algorithm; it is therefore helpful to briefly review the basics. Divergence is a convenient way to quantify the change of a vector field in space. Although our approach is applicable to 3D geometries, for greater simplicity, we express the mass flux divergence in 2D Cartesian coordinates as

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

where

$$\mathbf{V} = u \hat{\mathbf{i}} + v \hat{\mathbf{j}} \quad (3)$$

and $\hat{\mathbf{i}}$ and $\hat{\mathbf{j}}$ are unit vectors in the x and y coordinate directions, respectively. If at some spatial point location P , $\nabla \cdot (\rho \mathbf{V}) > 0$, then it is said that point P is a source of $\rho \mathbf{V}$, while $\nabla \cdot (\rho \mathbf{V}) < 0$ indicates a sink of $\rho \mathbf{V}$. It helps to point out that the product $\rho \mathbf{V}$ is called the *mass flux* [23], and represents the time rate of change of mass flow per unit area; dimensional analysis shows that $\rho \mathbf{V}$ is simply mass/(area · time). The role of this quantity in the CPT task can be better understood with the aid of the Divergence Theorem [25] from vector calculus:

$$\int_W \nabla \cdot (\rho \mathbf{V}) dW = \oint_S (\rho \mathbf{V}) \cdot d\mathbf{S} \quad (4)$$

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 a source (e.g., emitter) will have a positive mass flux divergence, while a control volume containing a sink will have a negative mass flux divergence. This result serves as our basic criterion for theoretically identifying a chemical emitter. To the best of our knowledge, previous criteria for emitter identification are purely heuristic, e.g., [14]. Our method is the first with a solid theoretical basis.

Furthermore, this result is also the basis of our novel plume tracing algorithm, which we call *fluxotaxis*. With fluxotaxis, the robotic lattice will compute the local divergence of mass flux, and will follow its gradient (the direction of steepest increase). Mathematically, the gradient being followed is:

$$\nabla(\nabla \cdot \rho \mathbf{V}) = \nabla\left(u \frac{\partial \rho}{\partial x} + \rho \frac{\partial u}{\partial x} + v \frac{\partial \rho}{\partial y} + \rho \frac{\partial v}{\partial y}\right) \quad (5)$$

Each individual robot independently calculates this flux gradient (5). Due to the virtual cohesive forces holding the lattice together, the whole lattice will move in the flux gradient direction determined by the majority (with no explicit voting).

From (2) it is clear that the fluxotaxis algorithm combines information about both velocity and chemical density, and the fact that it also encapsulates the notion of mass flux, as demonstrated in (4), provides assurance that we will find the emitter as opposed to a local density maximum. The following section presents several formal proofs in support of this statement.

6 Fluxotaxis Theory

Our ultimate objective is to invent a foolproof mathematical formula that the robotic lattice can use to guide it to a chemical source. To date, the fluxotaxis formula is our best candidate, although it is not foolproof. With our objective in mind, we are currently beginning an in-depth study of the strengths and weaknesses of the fluxotaxis technique. Through such an analysis, we anticipate discovering a variant of the fluxotaxis method that will satisfy our objective.

In this section, we prove a sequence of lemmas that begin to elucidate the strengths of the fluxotaxis strategy as a local guide to the location of the chemical emitter. In subsequent papers, we will also explore and rectify its weaknesses. Here, we present initial versions of lemmas that have restrictive (albeit realistic) assumptions; future versions will relax these assumptions. We limit ourselves to lemmas because the final theorem is the complete navigation strategy that we intend to develop. Each of the following lemmas looks at a realistic scenario and demonstrates the performance of a fluxotaxis-managed, 1D robotic swarm.

All of these lemmas assume a local coordinate system shared by all of the robots in the robotic lattice. Such a shared coordinate system is achievable via local communication accompanied by coordinate transformations [2, 26]. The lemmas in this section assume a single coordinate axis for simplicity; generalization to 2D is expected to be straightforward, due to symmetries, and has already been verified in software simulations.

6.1 Fluxotaxis in Constant Velocity

Constant Velocity Lemma 1. Assume that the following conditions hold:

1. Chemical plume has a general Gaussian distribution $\rho(x) = \kappa e^{-(x-c)^2}$, centered at $x = c$.
2. Lattice position x_0 is such that $x_L < x_0 < x_R$, where x_L, x_R are solutions to $\partial^2 \rho(x)/\partial x^2 = 0$ (see Fig. 1); this implies that $\partial^2 \rho(x)/\partial x^2 < 0$ in the region of interest.
3. \mathbf{V} is constant in magnitude throughout the flow, except right at the emitter ($x = c$), and is an outward radial vector.




Fig. 1. The Gaussian chemical density distribution and the radial outflow velocity profile used in the Constant Velocity Lemma 1. The shaded area indicates the region where plume tracing is carried out by the fluxotaxis agents, and the arrow at $x = c$ marks the location of the chemical emitter

W.l.o.g., assume the existence of P_{emt} and P_{far} such that P_{emt} is closer to the emitter than P_{far} . Then execution of one step of the fluxotaxis algorithm implies that the agent lattice moves closer to the emitter, or equivalently

$$\left[u \frac{\partial \rho}{\partial x} + \rho \frac{\partial u}{\partial x} \right]_{\text{far}} < \left[u \frac{\partial \rho}{\partial x} + \rho \frac{\partial u}{\partial x} \right]_{\text{emt}} \quad (6)$$

Proof. The problem is symmetric with respect to the emitter's location ($x = c$); thus it is sufficient to prove the case where $x_L < P_{\text{far}} < P_{\text{emt}} < c$. Because \mathbf{V} is constant, $\partial u/\partial x = 0$, and (6) simplifies to

$$\left[u \frac{\partial \rho}{\partial x} \right]_{\text{far}} < \left[u \frac{\partial \rho}{\partial x} \right]_{\text{emt}}$$

Since u is a negative constant, the inequality can be simplified to

$$\left[\frac{\partial \rho}{\partial x} \right]_{\text{far}} > \left[\frac{\partial \rho}{\partial x} \right]_{\text{emt}}$$

Grouping like terms gives

$$0 > \left[\frac{\partial \rho}{\partial x} \right]_{\text{emt}} - \left[\frac{\partial \rho}{\partial x} \right]_{\text{far}}$$

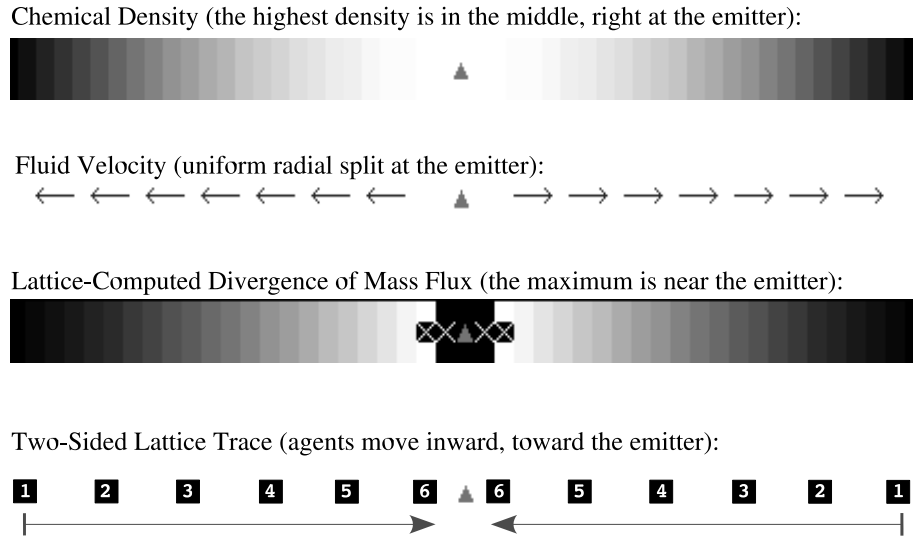


Fig. 2. Simulation results for the Constant Velocity Lemma 1. Individual agents are shown as black boxes with the white \times in the middle, and the time trace of the two independent agent lattices is shown with boxed numbers indicating the location of the lattice at a given time step. The Lemma holds for any initial lattice configuration, and fluxotaxis successfully locates the chemical emitter

This is true because, by assumption 2,

$$0 > \frac{\partial^2 \rho}{\partial x^2} \quad \square$$

Results of a software simulation for this lemma are shown in Fig. 2. In the figure, light-colored areas denote large values, and dark-colored areas correspond to small values. The location of the chemical emitter is marked by the triangle symbol. The initial positions of two separate agent lattices are at the outer edges of the environment, to the left and right of the emitter. During execution of the fluxotaxis algorithm, each agent (shown as a black box with a white \times in the middle) computes the divergence of the mass flux using (2), with the partial derivatives replaced by the second-order accurate central difference approximation [27]. This value is recorded by the simulator for analysis purposes, and is displayed along with the final agent positions in the screen shot. Observe that the resulting divergence “landscape” has a global peak which coincides with the location of the emitter, and does not have any local maxima that could trap or mislead the agents. There is a small gap in the computed divergence plot near the emitter because the agents had terminated their search upon reaching the emitter. Each simulated agent (the black box) corresponds to one of the reference points (P_{emt} or P_{far}) in the Lemma’s proof and, just as in the Lemma,



Fig. 3. Agent lattice coordinate axis orientation and the chemical source location in the Divergence Lemma 1

there are two agents per lattice. In this simulation, both agent lattices correctly moved toward the emitter in the center. In the proof of the Constant Velocity Lemma 1, we only considered the case where the lattice was to the left of the emitter; however, a similar proof can be given for the symmetric case, where the lattice starts out on the right side of the emitter, and the simulation in Fig. 2 demonstrates that the algorithm works regardless of the initial position of the agent lattice with respect to the chemical source.

6.2 Fluxotaxis at Source and Sink

Divergence Lemma 1. Fluxotaxis technique will advance the agent lattice toward a chemical source.

Proof. As before, assume a general Gaussian chemical plume distribution. W.l.o.g., assume the existence of two points P_{emt} and P_{far} , such that P_{emt} is closer to the source than P_{far} (see Fig. 3). Two cases result, based on the orientation of the lattice coordinate axis. (\mathbf{V} is at the bottom of Fig. 3, below the axis.)

Case I assumes that the direction of the lattice coordinate axis is opposite to the direction of the fluid flow, and thus

1. $\partial^2 u / \partial x^2 \geq 0$
2. $\partial u / \partial x > 0$; thus $0 \geq u_{\text{emt}} > u_{\text{far}}$
3. $\partial^2 \rho / \partial x^2 \leq 0$
4. $\partial \rho / \partial x > 0$ and therefore $\rho_{\text{emt}} > \rho_{\text{far}}$

We need to prove that the agent will move toward the source, or

$$\left[u \frac{\partial \rho}{\partial x} + \rho \frac{\partial u}{\partial x} \right]_{\text{far}} < \left[u \frac{\partial \rho}{\partial x} + \rho \frac{\partial u}{\partial x} \right]_{\text{emt}} \quad (7)$$

Assumptions 1 and 3 imply

$$\left[\frac{\partial u}{\partial x} \right]_{\text{far}} \leq \left[\frac{\partial u}{\partial x} \right]_{\text{emt}} \quad \text{and} \quad \left[\frac{\partial \rho}{\partial x} \right]_{\text{far}} \geq \left[\frac{\partial \rho}{\partial x} \right]_{\text{emt}}$$

Together with assumptions 2, 4, and algebraic rules, Case I holds. \square

Case II is with the lattice coordinate axis in the same direction as the fluid flow, so that both u_{emt} and u_{far} are non-negative (see Fig. 3), and the previous assumptions become

1. $\partial^2 u / \partial x^2 \leq 0$
2. $\partial u / \partial x > 0$; thus $0 \leq u_{\text{emt}} < u_{\text{far}}$
3. $\partial^2 \rho / \partial x^2 \leq 0$
4. $\partial \rho / \partial x < 0$ and therefore $\rho_{\text{emt}} > \rho_{\text{far}}$

The agent will turn around and move toward the source if (7) holds. From assumption 1 we conclude

$$\left[\frac{\partial u}{\partial x} \right]_{\text{far}} < \left[\frac{\partial u}{\partial x} \right]_{\text{emt}}$$

Similarly, assumption 3 yields

$$\left[\frac{\partial \rho}{\partial x} \right]_{\text{far}} < \left[\frac{\partial \rho}{\partial x} \right]_{\text{emt}}$$

Algebraic application of the remaining assumptions shows that (7) holds. \square

Chemical Density (the highest density is in the middle, right at the emitter):



Fluid Velocity (radial flow speeds up away from the emitter):



Lattice-Computed Divergence of Mass Flux (the maximum is near the emitter):



Two-Sided Lattice Trace (agents move inward, toward the emitter):



Fig. 4. Simulation of a fluxotaxis-driven lattice (represented by black boxes) in the vicinity of a chemical source from the Divergence Lemma 1. The time trace, denoted by the numbered boxes, shows the location of each of the two different agent lattices at sequential time steps in the simulation. Both lattices correctly converge on the true location of the chemical emitter

Software simulation of this Lemma’s configuration for both cases is shown in Fig. 4. As before, the fluxotaxis-driven lattice (represented by black boxes marked with the white \times symbol) begins at the outer edges of the simulated world, and moves in toward the emitter, denoted by the triangle in the center.




Fig. 5. Location of the chemical sink and the two possible agent coordinate axis orientations in the Divergence Lemma 2

The direction of motion is determined by the gradient of the divergence of the mass flux, which is computed locally by each agent using a central difference approximation of the partial derivatives in (2), and as can be seen from the divergence plot, has the maximum value near the emitter's location. Similar to the previous simulation, the divergence value right at the emitter is not computed by the lattice, since the search terminates as soon as the emitter is found. Two fluxotaxis lattices are shown in the screen shot, and as expected, both of them successfully navigate toward the chemical source. As this figure illustrates, the initial position of a lattice with respect to the emitter does not impede the agents' ability to correctly localize the emitter.

Divergence Lemma 2. Fluxotaxis-controlled agents will move away from a chemical sink (see Fig. 5).

Proof. As before, assume a general Gaussian chemical plume distribution. W.l.o.g., assume the existence of two points P_{snk} and P_{far} , such that P_{snk} is closer to the sink than P_{far} (see Fig. 5). To prove that the agents will move away from the sink, we must show

$$\left[u \frac{\partial \rho}{\partial x} + \rho \frac{\partial u}{\partial x} \right]_{\text{snk}} < \left[u \frac{\partial \rho}{\partial x} + \rho \frac{\partial u}{\partial x} \right]_{\text{far}} \quad (8)$$

Two cases result, based on the orientation of the lattice coordinate axis. (\mathbf{V} is at the bottom of Fig. 5, below the axis.)

Case I occurs when the lattice coordinate axis points in the opposite direction to the fluid flow, so that both u_{snk} and u_{far} are negative (see Fig. 5). For this case, the assumptions are

1. $\partial^2 u / \partial x^2 \geq 0$
2. $\partial u / \partial x < 0$; thus $0 \geq u_{\text{snk}} > u_{\text{far}}$
3. $\partial^2 \rho / \partial x^2 \leq 0$
4. $\partial \rho / \partial x < 0$ and therefore $\rho_{\text{snk}} > \rho_{\text{far}}$

The agent will continue moving away from the sink if (8) is true. From assumption 1 we observe that

$$\left[\frac{\partial u}{\partial x} \right]_{\text{snk}} \leq \left[\frac{\partial u}{\partial x} \right]_{\text{far}}$$

Chemical Density (the highest density is in the center, but the emitter is absent):



Fluid Velocity (radial flow slows down near the center):



Lattice-Computed Divergence of Mass Flux (the maximum is at the outer edges):



Two-Sided Lattice Trace (agents move outward, away from the center of the sink):



Fig. 6. Simulated performance of the fluxotaxis algorithm within the chemical sink from the Divergence Lemma 2. As stated in the proof and visualized in the last time-step diagram, the robust fluxotaxis method forces the agent lattice out of the sink, even if the lattice starts out directly in the center of the sink, where the chemical concentration is at a local maximum. The robust physical foundation of the fluxotaxis algorithm allows it to outperform the simpler chemotaxis CPT strategy

Likewise, assumption 3 implies

$$\left[\frac{\partial \rho}{\partial x} \right]_{\text{snk}} \geq \left[\frac{\partial \rho}{\partial x} \right]_{\text{far}}$$

The remaining assumptions with algebraic simplification prove that (8) is true. \square

Case II is when the direction of fluid flow and the lattice coordinate axis are the same, so that

1. $\partial^2 u / \partial x^2 \leq 0$
2. $\partial u / \partial x < 0$; thus $0 \leq u_{\text{snk}} < u_{\text{far}}$
3. $\partial^2 \rho / \partial x^2 \leq 0$
4. $\partial \rho / \partial x > 0$ and therefore $\rho_{\text{snk}} > \rho_{\text{far}}$

From assumptions 1 and 3 we conclude that

$$\left[\frac{\partial u}{\partial x} \right]_{\text{snk}} \leq \left[\frac{\partial u}{\partial x} \right]_{\text{far}} \quad \text{and} \quad \left[\frac{\partial \rho}{\partial x} \right]_{\text{snk}} \leq \left[\frac{\partial \rho}{\partial x} \right]_{\text{far}}$$

Algebraic simplification using assumptions 2 and 4 proves Case II. \square

Simulation results for this lemma are presented in Fig. 6. Confirming the theoretical results just obtained, the high-density chemical build-up in the center of the environment does not fool the fluxotaxis algorithm, which correctly avoids the local spike in the density by directing the agents (again represented by black boxes) to the outer edge of the tracing region, where as can be seen from the divergence plot, the maximum mass flux divergence occurs. *The Divergence Lemma 2 proves that a fluxotaxis-driven agent lattice will escape from a sink.* However, a simple chemotaxis strategy is easily fooled by sinks, since by definition of a sink, $\partial\rho/\partial x > 0$ going into the sink. The fluxotaxis scheme is more robust in this case because it looks at the second order partial of ρ , and also takes the divergence of velocity into account. This simulation provides an example of how effectively the fluxotaxis technique merges the chemotaxis and anemotaxis CPT methods into a physically sound algorithm with valuable self-correcting properties.

7 Summary and Future Work

In this paper, we presented a new chemical plume tracing algorithm called fluxotaxis, that combines key strengths of chemotaxis and anemotaxis - the two most popular plume tracing methods. We showed that the fluxotaxis algorithm has been developed from the fundamental physical principles of fluid flow, and that it is able to overcome a major flaw of chemotaxis. We also built a formal mathematical tool set that we will employ to further improve the algorithm. In particular, we plan to soon extend the basic fluxotaxis approach outlined here to handle turbulent eddies, thus overcoming a major flaw of anemotaxis. To experimentally confirm our theoretical results, we will implement the algorithm on a massively distributed system of simple robotic agents currently under development for the task of toxic chemical plume emitter localization.

The most important contribution of our work is the development of a mobile robotic swarm control algorithm that can be analyzed with formal methods, such that the agents' behavior can now be mathematically predicted and guaranteed. Some of our work is relevant to the design and evaluation of artificial worlds, as it develops and refines methods for emulation of real-world physics in a simulated environment. The distributed nature of the CFD computations performed by the robotic swarm may also be of interest to the community. The contribution of this research is interdisciplinary and has a wealth of applications in domains other than the chemical plume tracing we discussed in this paper.

References

1. Spears, W., Gordon, D.: Using artificial physics to control agents. In: Proceedings of the IEEE Conference on Information, Intelligence, and Systems (ICIIS'99). (1999)
2. Spears, W., Spears, D., Hamann, J., Heil, R.: Distributed, physics-based control of swarms of vehicles. In: Autonomous Robots. Volume 17(2-3). (2004)

3. Board on Atmospheric Sciences and Climate: Tracking and Predicting the Atmospheric Dispersion of Hazardous Material Releases: Implications for Homeland Security. National Academies Press (2003)
4. Hsu, S.S.: Sensors may track terror's fallout. *Washington Post* (2003) A01
5. Cordesman, A.H.: Defending america: Asymmetric and terrorist attacks with chemical weapons (2001) report produced by the Center for Strategic and International Studies (CSIS).
6. Caldwell, S.L., D'Agostino, D.M., McGeary, R.A., Purdy, H.L., Schwartz, M.J., Weeter, G.K., Wyrsh, R.J.: Combating terrorism: Federal agencies' efforts to implement national policy and strategy (1997) Congressional report GAO/NSIAD-97-254, produced by the United States General Accounting Office.
7. Spears, W., Heil, R., Spears, D., Zarzhitsky, D.: Physicomimetics for mobile robot formations. In: Proceedings of the Third International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-04). Volume 3. (2004) 1528–1529
8. Carlson, E.S., Sun, H., Smith, D.H., Zhang, J.: 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.
9. Ball, P.: Odour-tracking trick sniffed out. *Nature* (2002)
10. Adam, D.: Wing scents. *Nature* (2000)
11. Parunak, H.V.D., Brueckner, S.: Entropy and self-organization in multi-agent systems. In: Proceedings of the International Conference on Autonomous Agents (AGENTS'01). (2001) 124–130
12. Koenig, S., Liu, Y.: Terrain coverage with ant robots: A simulation study. In: Proceedings of the International Conference on Autonomous Agents (AGENTS'01). (2001) 600–607
13. Sandini, G., Lucarini, G., Varoli, M.: Gradient driven self-organizing systems. In: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'93). (1993)
14. Chemical plume tracing. In Cowen, E., ed.: *Environmental Fluid Mechanics*. Volume 2. Kluwer (2002)
15. Hayes, A., Martinoli, A., Goodman, R.: Swarm robotic odor localization. In: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'01). (2001)
16. Decuyper, J., Keymeulen, D.: A reactive robot navigation system based on a fluid dynamics metaphor. In Schwefel, H.P., Männer, R., eds.: *Parallel Problem Solving from Nature (PPSN I)*. Volume 496 of *Lecture Notes in Computer Science*., Springer-Verlag (1991) 356–362
17. Keymeulen, D., Decuyper, J.: 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., IEEE Computer Society Press (1994) 378–385
18. Keymeulen, D., Decuyper, J.: The stream field method applied to mobile robot navigation: a topological perspective. In Cohn, A.G., ed.: *Proceedings of the Eleventh European Conference on Artificial Intelligence (ECAI'94)*, John Wiley and Sons (1994) 699–703
19. Balkovsky, E., Shraiman, B.: Olfactory search at high reynolds number. *National Academies Press* **99** (2002) 12589–12593

20. Horling, B., Vincent, R., Mailler, R., Shen, J., Becker, R., Rawlins, K., Lesser, V.: Distributed sensor network for real time tracking. In: Proceedings of the International Conference on Autonomous Agents (AGENTS'01). (2001) 417–424
21. Polycarpou, M.M., Yang, Y., Passino, K.M.: Cooperative control of distributed multi-agent systems. IEEE Control Systems Magazine (2001)
22. Tannehill, J.C., Anderson, D.A., Pletcher, R.H.: Computational Fluid Mechanics and Heat Transfer. Taylor and Francis (1997)
23. Anderson, J.D.: Computational Fluid Dynamics. McGraw-Hill, Inc. (1995)
24. Versteeg, H.K., Malalasekera, W.: An introduction to computational fluid dynamics: the finite volume method. Longman Scientific and Technical (1995)
25. Hughes-Hallett, D., Gleason, A.M., et al., W.M.: Calculus: Single and Multivariable. John Wiley and Sons (1998)
26. Craig, J.J.: Introduction to robotics: mechanics and control. Addison-Wesley Publishing Company, Inc. (1989)
27. Faires, J.D., Burden, R.: Numerical Methods. Brooks/Cole - Thomson Learning (2003)