

Physicomimetics Positioning Methodology for Distributed Autonomous Systems

William M. Spears, Diana F. Spears, and Dimitri Zarchitsky

Department of Computer Science

University of Wyoming

Laramie, WY 82071

wspears@cs.uwyo.edu

Abstract: *This paper presents a physics-based framework for the distributed control of a mobile swarm of simple robots tasked with localizing a source of a toxic chemical plume. The framework, called physicomimetics, is a robust control scheme built on local interactions between the vehicles, making it highly scalable, adaptive, and cost effective. The chemical plume-tracing task discussed here is an example of a problem where vehicle collaboration improves performance of the system, by acting as a distributed computational mesh.*

Keywords: physicomimetics; artificial physics; chemical plume tracing.

Introduction

The focus of our research is to design and build rapidly deployable, scalable, adaptive, cost-effective, and robust networks of autonomous distributed vehicles. This combines sensing, computation and networking with mobility, thereby enabling deployment, self-organization, and reconfiguration of the multi-vehicle collective. Our objective is to provide a scientific, yet practical, approach to the design and analysis of distributed sensor systems.

The general purpose for deploying tens to hundreds of such vehicles can be summarized as “volumetric control.” Volumetric control means monitoring, detecting, 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. The result is enhanced situational awareness and the potential for rapid and appropriate response. Our objective is to design fully automated, coordinated, multi-vehicle sensor systems.

Background and Motivation

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. A vehicle's sensors perceive the world, including other vehicles, and a vehicle's effectors make changes to that vehicle and/or the world, including other vehicles. 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.

Our approach to the distributed control of autonomous systems is called “physicomimetics” or “artificial physics” (AP). We use the term “artificial” (or virtual) because although we are motivated by natural physical forces, we are not restricted to them. Although the forces are virtual, vehicles *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. The vehicle's effectors must allow it to respond to this perceived force.

There are two potential advantages to this approach. First, in the real physical world, collections of small entities yield surprisingly complex behavior from very simple interactions between the entities. Thus, there is a precedent for believing that complex control is achievable through simple local interactions. This is required for very small vehicles, since their sensors and effectors will necessarily be primitive. Two, since the approach is largely independent of the size and number of vehicles, the results scale well to larger vehicles and larger sets of vehicles.

Three primary emphases distinguish the AP framework from others that are related: minimality, ease of implementation, and run-time efficiency. First, AP formations are achieved with a minimal set of sensors and sensor information. Our rationale is that this will:

1. reduce overall vehicle cost,
2. enable physical embodiment with small vehicles, and
3. increase vehicle stealth if sensing is active.

Second, theoretical results translate directly into practical advice on how to set system parameters for desired system performance. This makes the robotic implementation straightforward and its deployment rapid. Third, AP is computationally efficient [1].

This paper is organized as follows. First, we present the general AP framework. Second, using AP, we illustrate how vehicles can self-organize into hexagonal lattices, which are used as sensor networks. Then we address the “chemical plume tracing” (CPT) task, which traces a toxic plume to its source emitter. *In our approach, each robot acts as a computational fluid dynamics grid point, thereby*



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

enabling the robotic lattice to behave as a distributed computer for fluid computations. Although robots only gather information about the fluid locally, the desired direction of movement for the overall lattice emerges in the aggregate, without any global control.

Physicomimetics Approach

In our physicomimetics framework, 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, and the system acts as a molecular dynamics simulation.

To create hexagonal 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 $r > R$. The only effector is to be able to move with velocity v . To ensure that the force laws are local in nature, robots have a visual range of only $1.5R$. Also, due to the discrete-time nature of the model, it is important to define a maximum force F_{\max} that can be obtained. This force law works well for constructing hexagonal lattices of robots [1-4].

Formation Control Using Virtual Forces

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 also attempts to sense some 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 focus of this project is the physical embodiment of the physicomimetics framework on a team of robots. For our experiment, we built seven robots. The 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, $G < 308$ 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 formation and never showing clustering.

The light 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 not computed by individual robots and there is no global controller – it is an emergent property of the collective.*

Chemical Plume Tracing (CPT) Task

To illustrate how the physicomimetics framework is used in practice, we present its application to the problem of finding the source of a toxic airborne plume, which may form after an accident at an industrial facility or at an urban area after a deliberate terror attack. We assume a single, continuous chemical emitter and a moderate air current that causes the expansion and movement of the plume.

Almost all prior CPT solutions derive from two biomimetic approaches: *chemotaxis*, or density gradient following, and *anemotaxis*, or upwind following [5, 6]. These approaches yield satisfactory results in simple scenarios; however, they fail when the fluid flow is turbulent, or the search area contains obstacles and corners [7, 8]. Generally, they are difficult to formalize, and little theory is available to assist in the design and validation of plume-tracing systems based on these methods.

To address the shortcomings of these approaches, the authors developed a physics-based plume-tracing algorithm




Figure 2. Sample execution sequence of a CPT algorithm. Denser chemical concentrations are drawn with darker shades. The lattice, denoted by a set of seven black ‘x’ marks, begins in the top left corner, and traces the plume to its source in the bottom right corner. The wind field is omitted to improve clarity of the images.

called *fluxotaxis* [9], which takes advantage of the physicomimetics framework to form an adaptive sensor grid for measuring *flow variables*, such as chemical density and fluid velocity, and then computes the gradient of the divergence of mass flux, expressed mathematically as:

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

where u and v are fluid velocity in the x and y coordinate directions, and ρ is the chemical density. The continuous derivatives are approximated by the discretized equivalents, *with the lattice operating as a distributed computational mesh*. The fluxotaxis algorithm is based on the physics of fluid flow; it combines the strengths of both chemo- and anemotaxis, uses local information from the AP-maintained lattice, and a rigorous mathematical analysis complements the theoretical and empirical results for the AP framework.

CPT Algorithm Simulation Experiments

To study the behavior and to evaluate the performance of the different CPT strategies, we constructed a sophisticated simulator, which includes a realistic flow-field solver based on the work of Farrell, *et al.* [10], providing the “*forward solution*” of the CPT problem for the transport of the chemical plume from the emitter and into the environment. The other simulator component produces the “*inverse solution*,” which implements the three previously described CPT strategies. For the purpose of this study, we assume seven identical robots, similar to our setup in the lab, all using the physicomimetics approach to self-organize into a hexagonal formation. The robots then use strictly local interactions to maintain the formation and to move toward goal locations computed by a given CPT algorithm. Figure 2 shows grayscale screenshots of a plume-tracing sequence, with the developing plume and a moving lattice.

We measured the performance of the CPT algorithms with two parameters: a real-valued emitter proximity index (higher is better), and a Boolean emitter enclosure test, called “*success*.” The selection of these evaluation metrics is driven by their relevance to the CPT task: for success, we would like the robotic swarm to localize the emitter by surrounding it (so that it can be extinguished). For the less favorable case, when the exact location of the source cannot be determined, the robotic lattice should be as close to the emitter as possible. Note that the anemotaxis and chemotaxis algorithms do not have termination or source identification criteria; reference implementations in the literature [5, 6, 8, 10] tend to rely on heuristic rules, which vary with a given problem. In contrast, due to its theoretical foundation in fluid mechanics [9], the fluxotaxis algorithm provides an explicit method for chemical source detection, and is thus capable of terminating its execution upon discovery of the chemical source. For the sake of fairness in the comparison of the experimental results, we disable the source detection ability of fluxotaxis, and instead let each algorithm execute for a fixed number of steps, which we determined empirically to be sufficient for a lattice to reach the chemical emitter.

Experimental results are summarized in Fig. 3 and Table 1, for 35 different plumes with flow regimes varying from stable and laminar to transient and turbulent. The tracing area is 10,000 sq. ft., and each robot is approximately one square foot in size. The minimal distance between center-points of each neighboring robot in the lattice is set at 1.5 feet. Our experiments vary the maximal radial expansion factor of the lattice from 5 to 15, resulting in an inter-vehicle local communication range of about 8 to 23 feet. For each plume scenario, the performance metrics are averaged over 200 randomly selected starting locations.




Figure 3. Performance of the CPT algorithms versus the maximum radial expansion factor over 35 plumes, with 200 random starting locations per plume.

Table 1. CPT algorithm performance averaged over 35 plumes, 200 restarts, and 10 radial expansion factors, showing mean \pm standard deviation.

Algorithm	Proximity	Success
Anemotaxis	0.6843 ± 0.0291	0.0667 ± 0.0362
Chemotaxis	0.6745 ± 0.0049	0.3184 ± 0.0132
Fluxotaxis	0.9235 ± 0.0089	0.7460 ± 0.0250

Discussion: the stability of the results across a wide range of radial expansion factors confirms the robustness of the AP control framework. Stable performance in spite of the changing dynamics of the system is a characteristic aspect of the physicomimetics-controlled system. Each CPT algorithm requires satisfaction of two conflicting goals: a stable formation for accurate sensing, and a moving, dynamic formation for rapid response to a changing flow. Physicomimetics is able to meet both of these constraints, which is evident from the consistently good performance of the fluxotaxis strategy. Note that the analytically designed fluxotaxis algorithm outperforms the heuristic chemo- and anemotaxis algorithms. Because the fluxotaxis strategy incorporates local knowledge about fluid flow, it is able to use the combined limited sensor information from the robotic swarm with greater efficiency and effectiveness than the other strategies. This allows fluxotaxis to finish the CPT task faster and more reliably, and to take full advantage of the distributed system architecture.

Summary

In this paper, we presented the physicomimetics framework for distributed autonomous systems, and showed that it is robust, cost-effective, and easy to deploy. We provided an

example of an analytical analysis of the system that provides guidance on setting system parameters. This analysis provides the means to go directly from theory to application, reducing cost and time of deployment. To illustrate one practical use of the physicomimetics framework, we showed how seven vehicles, controlled by the framework, locate in simulation a source of toxic plume. The plume-tracing algorithm we developed from the principles of fluid physics is easy to formalize, fitting naturally within the physicomimetics methodology for distributed control of inexpensive robotic vehicles.

References

- Spears, W., Spears, D., Heil, R., Kerr, W., and Hettiarachchi, S., An Overview of Physicomimetics. *Lecture Notes in Computer Science – State of the Art Series*, vol. 3342, 2005.
- Spears, W., Spears, D., Hamann, J., Heil, R. Distributed, Physics-Based Control of Swarms of Vehicles. *Autonomous Robots*, vol. 17(2-3), 2004.
- Gordon, D., Spears, W., Sokolsky, O., and Lee, I., Distributed Spatial Control, Global Monitoring and Steering of Mobile Physical Agents. *IEEE International Conference on Information, Intelligence, and Systems*, November, 1999.
- Spears, W., and Gordon, D., Using Artificial Physics to Control Agents. *IEEE International Conference on Information, Intelligence, and Systems*, 1999.
- Hurtado, J.E., Robinett, R.D., Dohrmann, C.R, and Goldsmith, S.Y., “Distributed Sensing and Cooperating Control for Swarms of Robotic Vehicles,” *International Conference on Control and Applications (IATED)*, 1998.
- Hayes, A., Martinoli, A., and Goodman, R., “Swarm Robotic Odor Localization,” *International Conference on Intelligent Robots and Systems (IROS)*, 2001.
- Feddema, J.T., Robinett, R.D., and Byrne, R.H, “An Optimization Approach to Distributed Controls of Multiple Robot Vehicles,” *International Conference on Intelligent Robots and Systems (IROS)*, 2003.
- Cui, X., Hardin, C.T., Ragade, R.K., and Elmaghraby, A.S., “A Swarm Approach for Emission Sources Localization,” *International Conference on Tools with Artificial Intelligence (ICTAI)*, 2004.
- Zarzhitsky, D., Spears, D., Thayer, D., and Spears, W., “Agent-Based Chemical Plume Tracing Using Fluid Dynamics,” *Lecture Notes in Computer Science*, vol. 3228, Springer-Verlag, 2004.
- Farrell, J.A., Murlis, J., Li, W., and Carde, R.T., “Filament-Based Atmospheric Dispersion Model to Achieve Short Time-Scale Structure of Odor Plumes,” *Environmental Fluid Mechanics*, vol. 2, 2002.