

Towards 3-D Distributed Odor Source Localization: An Extended Graph-Based Formation Control Algorithm for Plume Tracking

Jorge M. Soares, Ali Marjovi, Jonathan Giezendanner, Anil Kodiyan,
A. Pedro Aguiar, António M. Pascoal, Alcherio Martinoli

Abstract—The large number of potential applications for robotic odor source localization has motivated the development of a variety of plume tracking algorithms, the majority of which work in restricted two-dimensional scenarios. In this paper, we introduce a distributed algorithm for 3-D plume tracking using a system of ground and aerial robots in formation. We propose an algorithm that takes advantage of spatially distributed measurements to track the plume in 3-D and lead the robots to the source by integrating three behaviors – upwind movement, plume centering, and Laplacian feedback formation control. We evaluate this strategy in simulation and with real robots in a wind tunnel. For a source close to the ground, results show that a team of robots running our algorithm reaches the source with low lateral error while also tracing the horizontal and vertical plume shape.

I. INTRODUCTION

With the advances in robotics, embedded systems, and chemical sensor research in the last two decades, odor sniffing robots and sensor nodes have grown into a very active research field in which searching for odor sources, mapping odor distribution, and discriminating among chemical substances are some of the most popular topics. A number of interesting applications concerned with airborne and waterborne plumes can be found in the security, safety, military, and medical domains, with canonical examples including the detection of oil spills or wild fires, the localization of landmines in humanitarian demining operations, and the identification of dangerous leaks inside tunnels, mines, or production plants.

The focus of the present study is odor source localization, that is, finding the source of a chemical release using a set of sensing nodes. Chemical release substances can range from molecules (e.g., O_3 , CO_2 , NO_x , SO_2) to droplets (e.g., water vapor) or particles (e.g., fine particles found in smoke). The odor is released by a source and mainly transported by the airflow, forming a three-dimensional odor plume. As the plume travels away from the source, it becomes more diluted due to molecular diffusion and turbulence that mixes the odor molecules with clean air [1].

J. M. Soares, A. Marjovi, J. Giezendanner, A. Kodiyan and A. Martinoli are with the Distributed Intelligent Systems and Algorithms Laboratory, School of Architecture, Civil and Environmental Engineering, École Polytechnique Fédérale de Lausanne (EPFL), 1015 Lausanne, Switzerland.

J. M. Soares and A. M. Pascoal are with the Laboratory of Robotics and Engineering Systems, Instituto Superior Técnico, University of Lisbon, Av. Rovisco Pais, 1049-001 Lisboa, Portugal.

A. P. Aguiar is with the Research Center for Systems and Technology, Faculty of Engineering, University of Porto, Rua Dr. Roberto Frias, 4200-465 Porto, Portugal.

This work was partially funded by Fundação para a Ciência e Tecnologia under project FCT [UID/EEA/5009/2013] and grant SFRH/BD/51073/2010, and by the Swiss National Science Foundation under grant 200021.153310/1.

Molecular diffusion is a slow process whose short-term effect on the plume shape and the odor concentration can be neglected; thus, in outdoor or ventilated indoor environments, the dispersion of odor molecules is dominated by flow turbulence. The odor molecules move downwind due to mean flow velocity U while their net motion is almost random because of small-scale turbulence curls. This causes the smooth odor gradient to break into patches. As the odor is carried downwind, the average concentration within a patch decreases while the average time between successive patches increases.

The odor source localization problem can be divided into three sub-problems or phases [2], although these actions might not be necessarily performed consecutively: (i) odor plume finding [3], (ii) odor plume tracking [4], and (iii) odor source declaration [5]. This paper targets the odor plume tracking problem, considering an airborne plume generated by a single source and transported by the wind. The underlying formation framework can, however, be used in all three phases.

Odor plume tracking has been the focus of many studies in mobile robotics. Gradient ascent [6], bio-inspired [7], probabilistic inference [8], [9], and map-based [10] methods are the main families of solutions, mostly implemented on single-robot systems. Leveraging multiple robots for distributed search has its own peculiarities, many of which have been explored during the last decade. Multi-robot search algorithms proposed in the literature include biasing expansion swarm approaches [11], biased random walk [12], particle swarm optimization [13], gradient climbing techniques [14], infotaxis [15], probabilistic reasoning [16], search through exploration [17], physics-based swarming [18], attraction-repulsion swarming [19], and formation-based algorithms [20], [21].

Although the nature of the odor plumes and of the underlying turbulence phenomena is intrinsically 3-D, in all of these works the problem was reduced to 2-D search under highly simplified environmental conditions. There are but a few previous works studying olfactory problems in 3-D, most relying exclusively on simulations. Kuroki et al. [22] presented an expert system for contaminant mapping based on a genetic algorithm for simulated Unmanned Aerial Vehicles (UAVs) and considering the Gaussian plume/puff model for odor distribution. In a simple simulation environment, Kovacina et al. [23] presented a decentralized swarming algorithm for mapping a chemical cloud. This method considers the restrictions on sensors, computation, and flight envelope that affect aerial robots. Bamberger et

al. [24] proposed a stigmergic potential field approach for coordination between UAVs. Their experiments featured simulated sensors and plumes. Wu et al. [25] developed a stochastic model for plume spikes based on a Poisson counting process, transforming odor measurements into a smooth field that supports a simpler source-seeking behavior.

In a non-aerial 3-D setting, Osório et al. [26] deployed multiple gas sensors vertically on a single wheeled mobile robot to perform 3-D plume tracking using various well-known algorithms (surge-caste, casting, and gradient following). They found that plume tracking with three-dimensional acquisition outperformed a two-dimensional solution by 25% in terms of success rate. A similar setup was used in [27] to enable a map-based approach. Badia et al. [28] used a single blimp equipped with gas sensors for a mine clearance task including chemical mapping. Neumann et al. [29] tackled the odor distribution mapping problem using a single microdrone and tested their odor mapping methodology outdoors and in a wind tunnel. The same author with different colleagues also proposed a source localization approach [30]. Ishida et al. [31] used a handheld probe comprising a 3-D anemometer and six odor sensors to guide a human to a gas source by combining the wind direction and the instant odor gradient. We are unaware of any work in this area that addressed the problem of odor source localization with a distributed and heterogeneous 3-D sensing system.

To state the problem, consider a heterogeneous group of N mobile sensing nodes positioned in \mathbb{R}^3 . The question we address is how to efficiently track a detected odor plume to its source. The sensing nodes have different mobility constraints, depending on whether they are ground or aerial vehicles, and are equipped with sensors that measure both odor concentration and airflow direction. Given the scales, we assume each node is able to communicate with every other node and has line-of-sight relative localization capabilities. No central base station exists and therefore the nodes must act autonomously.

The main contribution of the present study is to depart from the classical 2-D spatial approach and add the third dimension using a distributed robotic system comprising ground and aerial robots. This makes it possible to localize odor sources that are not on the ground plane, provided that the plume is detectable by some of the vehicles. Furthermore, we believe that migrating from conventional ground robots to a distributed system capable of accessing various locations in 3-D space will lead to novel results in odor plume tracking, especially considering that odor distribution is an inherently 3-D phenomenon.

We approach the problem by adapting our previously proposed graph-based Laplacian formation method [20], [32], [33] to a 3-D system consisting of ground and aerial robots. The 2-D version of our method works on two basic principles: i) under stable wind conditions, the wind direction provides a strong hint as to the relative position of the source; and ii) having multiple robots distributed in the environment provides us with enough data to allow a computationally simple algorithm to track a plume in spite of its patchiness.

Starting from these principles, we implement a method based on three components: formation control, upwind movement, and plume centering. A Laplacian feedback controller allows us to organize the robots in an arbitrary formation that may change over time to better adapt to the plume. Wind direction and odor concentration sensors on each robot allow us to bias the base movement in the incoming wind direction and to center the entire formation within the plume based on the differential readings of robots in different points of the formation.

We have shown our 2-D algorithm to be efficient in tracking a plume both in simulation and using Khepera IV robots in a wind tunnel with reproducible laminar flow conditions [20]. The main contribution of this paper is the extension of the framework to a 3-D formulation, involving additional expressions for height control and formation scaling in the vertical direction. We present simulation results using the Webots high-fidelity robotic simulator with a custom odor and wind modeling plugin, as well as real-robot validation results obtained in the same wind tunnel using a set of ground robots and an additional node capable of 3-D movement through its anchoring to a Cartesian robot inside the wind channel.

The paper is organized as follows: Section II outlines our method, starting with an introduction to graph-based formation control and then describing our plume tracking algorithm and implementation; Section III explains our experimental scenarios, both for simulation and wind tunnel experiments, and discusses the results obtained; finally, Section IV summarizes our work and draws conclusions and plans for future research.

II. PROPOSED METHOD

Our method is based on three main components, each yielding an individual movement vector: formation control (\mathbf{u}_f), upwind movement (\mathbf{u}_w), and plume centering (\mathbf{u}_c). The desired direction movement of each vehicle is computed as a weighted sum of each component. This section describes each component in detail.

A. Background on graph-based formations

Laplacian formation control is a solution to the problem of guiding a group of individual mobile robots to a desired spatial formation. In this approach a graph $\mathcal{G} = (V, E)$ is defined in which vertices V correspond to controlled mobile robots and edges E correspond to inter-robot communication and relative positioning links. Built upon basic linear algebra, a stable solution to the formation control problem in two dimensions [34] is given by

$$\dot{\mathbf{x}} = -(\mathcal{L} \otimes I_2)(\mathbf{x} - \mathbf{b}), \quad (1)$$

where $\mathcal{L} = \mathcal{I} \cdot \mathcal{W} \cdot \mathcal{I}^T$ (called the Laplacian matrix) is obtained from the incidence matrix \mathcal{I} that defines the edges of \mathcal{G} and the optional diagonal weight matrix \mathcal{W} used to tune the weights assigned to the edges in the graph. I_2 is simply a 2×2 identity matrix (for the case of a two-dimensional plane). The (x, y) absolute position vector for all robots is given

by \mathbf{x} , and the desired offsets of each robot to the formation centroid are given by the bias matrix \mathbf{b} .

In the above formulation, the absolute position of each robot may be trivially replaced with the relative position with respect to the neighboring robots. In this case, (1) yields the goal position in the local coordinates of the robot. Most robotic platforms do not measure relative positions in x and y but, instead, use polar exteroceptive sensors (e.g., infrared, LASER) to obtain the range and bearing to neighboring robots.

Finally, the aforementioned movement goal for each robot i is given by

$$\begin{aligned} \dot{x}'_i &= - \sum_{j=1}^N \mathcal{L}_{i,j} (x'_j - b_j^x) \\ \dot{y}'_i &= - \sum_{j=1}^N \mathcal{L}_{i,j} (y'_j - b_j^y) \end{aligned}, \quad (2)$$

where b_j^x and b_j^y are respectively the desired biases in the x and y dimension of robot j and $\mathcal{L}_{i,j}$ corresponds to the entry (i, j) of the Laplacian matrix \mathcal{L} . It is mathematically proven and experimentally shown that this Laplacian control law makes the mobile robots converge to the desired formation [35].

B. Adaptive 3-D formation control

For a 3-D formation, the governing equation that yields the 2-D movement vector stated in (2) continues to apply with the straightforward addition of an equivalent expression for the z axis and new biases b_j^z . The first component of the algorithm, formation control, thus becomes

$$\mathbf{u}_f = - \begin{bmatrix} \sum_{j=1}^N \mathcal{L}_{i,j} (x'_j - b_j^x) \\ \sum_{j=1}^N \mathcal{L}_{i,j} (y'_j - b_j^y) \\ \sum_{j=1}^N \mathcal{L}_{i,j} (z'_j - b_j^z) \end{bmatrix}. \quad (3)$$

As stated, (3) drives a group of robots to the formation specified by the bias vectors but does not guarantee a particular orientation of the formation. Since our goal is to place robots at specific locations within the plume section, we replace the b_j vector in (3) with a dynamic bias vector β_j given by

$$\beta_j = R_z(\theta_w) \begin{bmatrix} s_{uw} & 0 & 0 \\ 0 & s_{cw} & 0 \\ 0 & 0 & s_v \end{bmatrix} \begin{bmatrix} b_j^x \\ b_j^y \\ b_j^z \end{bmatrix}, \quad (4)$$

where $R_z(\theta_w)$ denotes a rotation matrix about the z -axis going through the robot center by the azimuthal angle of the wind velocity θ_w . This rotates the bias vectors, and therefore the formation, to keep it oriented towards the wind.

In order to better track the plume, we implement an additional mechanism to adjust the formation biases based on the odor concentration readings. The scaling factors s_{uw} , s_{cw} , and s_v in (4) allow us to dynamically scale the formation to adapt to the plume shape. We choose to keep a constant

upwind scaling (s_{uw}), and adjust s_{cw} and s_v according to the laws

$$\dot{s}_{cw} = k_{cw}((c_l + c_r) - c_c) \quad (5)$$

$$\dot{s}_v = k_v(0.66c_t - c_b), \quad (6)$$

where c_l is the mean concentration reported by the N_l robots on the left side of the plume (i.e., whose coordinate in the upwind–crosswind frame is negative) and c_r is analogous for the robots on the right side. Likewise, the quantities c_t and c_b used for vertical scaling are the average concentrations measured by the top and bottom robots. The center concentration c_c is the mean concentration reported by robots in the center (i.e., the robots with null crosswind bias), of particular use to scaling. The process gains k_{cw} and k_v affect the dynamic response of the controllers but not their points of convergence.

The goal of (5) is to match the sum of the side concentrations to the center concentration so that as the plume narrows or widens, the formation does too. The goal of (6) is to match the odor concentration measured by robots on the ground and in the air, i.e., encompassing the plume with the two groups, not only allowing us to determine the source height but also the vertical shape of the plume. Equation (5) results in a mean bottom concentration c_b that follows $\frac{2}{3}c_c$ of the bottom center concentration; since we work with a single centered aerial robot, we adjust (6) accordingly. For a vertically symmetrical formation, this factor would be 1.

It is worth noting that the positions used to assign measurements are the steady-state positions in the formation specification and not the current positions of each robot. This may result in erroneous concentration measurements during the initial convergence but, as the resulting effect impacts all robots similarly, the system spontaneously recovers.

C. Upwind urge

Equation (2) drives the robots to a desired formation around the center of mass of the group; once the robots arrive at their desired positions, they stop. However, to find the source of a gaseous distribution, the robots must move in the direction of the source. Because advection (flow transport) is the main factor of displacement of odor patches in our environment, we implement upwind movement for each robot in the form

$$\mathbf{u}_w = R_z(\theta_w) \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}. \quad (7)$$

This expression projects an upwind movement vector in the local coordinates and implicitly assumes that the wind is approximately parallel to the ground plane. While it is straightforward to modify the algorithm in a way that eliminates this constraint, wheeled robots are unable to track a plume away from the ground plane; such a scenario would require a team consisting entirely of robots moving in 3-D space.

D. Plume centering

In addition to staying in formation and moving towards upwind direction, the robots must remain in the plume in order to eventually reach the source. Although the plume is patchy, the time-averaged odor concentration will, in principle, decrease as we move away from its center. Therefore, we include a plume centering component that drives the formation cross-wind towards the side with the highest concentrations and is given by

$$\mathbf{u}_c = R_z(\theta_w) \begin{bmatrix} 0 \\ -u_c^{max} + \frac{2u_c^{max}}{1+e^{-(c_l-c_r)/k_l}} \\ 0 \end{bmatrix}, \quad (8)$$

where u_c^{max} is the maximum crosswind force that can be requested by the algorithm and the constant k_l is proportional to the dynamic range of the sensor and is manually tuned to get a smooth response.

Equation (8) introduces a generalized logistic response on the differential odor reading, with \mathbf{u}_c reflecting the magnitude of misalignment of the robot relative to the center line of the plume. If the concentration reported by the robots on the left and on the right is balanced (i.e., $c_l = c_r$), then the formation is close to centered and \mathbf{u}_c approaches 0.

In a scenario with ground robots, it makes no sense to include a vertical centering component: centering, by definition, would require the entire formation to move up or down, and this is not physically possible with this setup. Vertical adjustments are, for these reasons, solely within the scope of the dynamic bias control.

E. Final movement vector

The final movement goal for each robot $\mathbf{u} = [u_x, u_y, u_z]'$ combines the desired movement vector of each component in the weighted sum

$$\mathbf{u} = k_w \mathbf{u}_w + k_c \mathbf{u}_c + k_f \mathbf{u}_f \quad (9)$$

and points, in the local frame, towards the direction the robot should move in so as to better optimize a specified balance of the three behaviors. The weights k_w , k_c , and k_f allow us to adjust the relative influence of each component, although in this particular work we retain uniform weights $k_w = k_c = k_f = 1$.

The aerial robot is holonomic in \mathbb{R}^3 and can therefore track the commanded \mathbf{u} directly. For the differential drive ground robot, the reference \mathbf{u} needs to be handed off to a lower-level controller to compute appropriate linear and angular speeds. We implement the same controller used in [20].

III. EVALUATION AND RESULTS

The proposed method has been implemented and evaluated in simulation and in a real wind tunnel. We now present our experimental setup, scenarios, and results.

A. Simulations

To evaluate the performance of the proposed method, we designed a simulated controlled environment in Webots [36], configured and calibrated to approximate real world characteristics of the environments, sensors, and actuators.

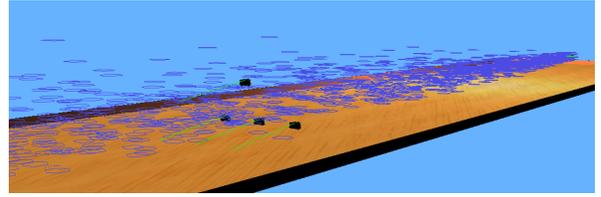


Fig. 1. Webots simulation setup. Four (three ground and one aerial) robots track the odor plume, with odor filaments drawn in blue.

1) *Setup*: The simulation arena, shown in Fig. 1, is a $20 \times 4 \times 4$ m³ volume. The odor simulation is based on the filament model proposed by Farrell et al. [37] and generates an intermittent 3-D plume similar to the one observed in the wind tunnel. Odor is modeled as a set of filaments, each containing a constant amount $s = 8.3 \times 10^9$ of molecules. At each time step, the position of a filament is updated according to a uniform wind field and a stochastic process consisting of a vector of three independent Gaussian random variables. More details of the odor dispersion simulation in Webots are provided in [32].

The virtual odor source releases 100 filaments per second with an initial width of 10 cm and initial position distributed over the circular area of the source. The odor concentration measured by a robot is the sum of the contribution of all filaments, which decays exponentially with the increasing distance to the center of a filament. In order to better cope with the patchiness of the plume, samples are run through a sliding window filter that outputs the highest amongst the 50 most recent readings, a window of 1.6 s.

2) *Robots and scenario*: We use a four-robot team consisting of three ground robots and one moving in 3-D. For simplicity, we use the built-in model of the Khepera III [38] for all robots, including the aerial one. As the Khepera is a wheeled robot and incapable of vertical movement, we employ a *supervisor* controller to freely displace the aerial robot in 3-D. This is a good approximation of the real setup in the wind tunnel, detailed in Section III-B.1.

The range, bearing, and elevation mechanism is based on infrared emitters. Measurements are subject to Gaussian noise drawn from independent random variables with standard deviation of 0.1 rad in bearing and elevation and 10% in distance, and are locally converted to (x, y, z) coordinates.

The source is placed at the 1 m downwind mark, centered, at varying heights. The chosen odor source heights for different experimental runs are 0.0, 0.25, 0.5, and 1.0 m. The robots start approximately 14 m downwind. We use a triangle formation with the ground robots in a crosswind line and the aerial robot moving behind the ground flock; the default biases for the four robots are, in meters,

$$\begin{aligned} b_x &= [0.0 & 0.0 & 0.0 & -0.5] \\ b_y &= [-0.5 & 0.0 & 0.5 & 0.0] \\ b_z &= [0.0 & 0.0 & 0.0 & 0.15] \end{aligned} \quad (10)$$

where the first column represents the robot on the left side, the second column the center robot, the third column the robot on the right side, and the last column the aerial robot.

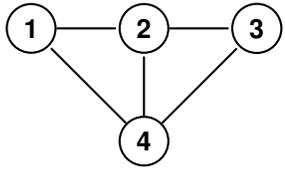


Fig. 2. Formation graph used in the experiments. Node labels correspond to the indexes in the bias vectors, with node 4 being the aerial robot.

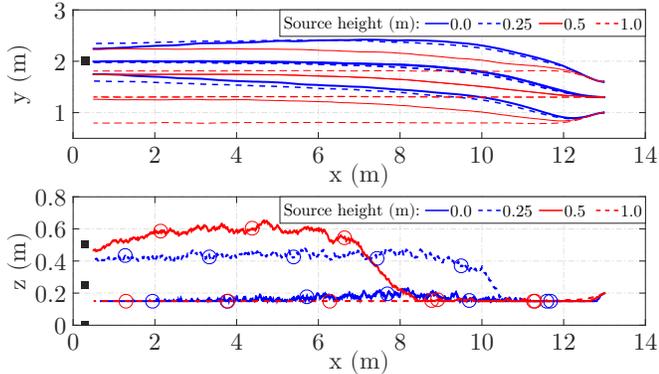


Fig. 3. (Top) Plot of the horizontal trajectory of the three ground robots for different plume heights (m). (Bottom) Vertical trajectory of the aerial robot in the same experiments. The black squares on the left mark the position of the source, and circles show the robot positions every 40 s.

To approximate the real setup, where care must be taken not to occlude the view of the overhead cameras, the aerial robot must always float at a minimum height of 15 cm and lag at least 50 cm behind the ground robots.

The graph used for the Laplacian formation control is shown in Fig. 2. Robots take into account their relative position to all teammates in line of sight.

3) *Results*: Fig. 3 presents the trajectories followed by the robot on four discrete simulation runs for different source heights. In all runs, the robots start in the same positions, with the formation shifted 0.5 m off-center. The robots start on the right-hand side of the plot and move left towards the source.

The ground tracks are shown in the top plot. We can see that, for low plumes (at $h = 0$ m and $h = 0.25$ m), the robots quickly move to the plume center and reach the end of the tunnel with the formation aligned to the plume source. In fact, the only clearly observable difference between the two scenarios is the formation scaling, due to differences in the center concentration measured by the top robot. As the source height increases, the tracking behavior deteriorates. With $h = 0.5$ m the formation still moves towards the plume center and reaches the end approximately 0.2 m off-center, but no centering effort is observed when the source is at $h = 1$ m. As we only use one aerial robot, the centering behavior is determined by the side robots on the ground. Thus, in this case where the ground robots are unable to detect the plume, there can be no lateral adjustments.

In the bottom plot, we can see the aerial robot moving along the top edge of the plume. For a ground plume, the aerial robot remains at its lowest defined z bias of 0.15 m.

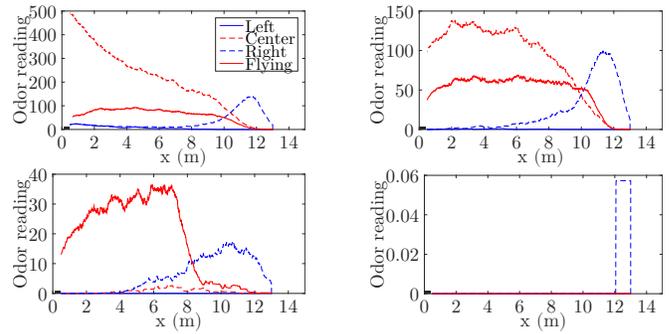


Fig. 4. Plots of each robot’s odor readings for four different odor source heights (top to bottom, left to right): 0.0 m, 0.25 m, 0.50 m, 1.0 m. For readability, the signals were smoothed over 500 samples and the y axes use different scales.

Progressively higher plumes lead to the robot climbing higher. For a source at 0.5 m the ground robots stop being exposed to the plume as they approach the source, and the aerial robot compensates by dropping. Finally, for the top height of 1 m, the plume is not detected by any robot (readings below 0.1 in a 0–1000 range) and the formation merely moves upwind.

The above results are confirmed by Fig. 4, which shows the odor readings for each robot over the course of the experiment. In the interest of readability, we plot a smoothed signal (using a 500-sample moving average filter) instead of the noisy original signal; for an example of the actual measurements the robots use when making decisions, we refer the reader to Fig. 7.

It is clear from these results that our approach fails when plumes are too high up in the air. This is to be expected, as our solution depends on ground robots that cannot move upwards and are therefore unable to trace the plume. For plumes at ground level, the solution is effective but there are limited gains from having an aerial robot in the formation.

In order to assess the performance impact of introducing 3-D tracking at plume heights where all robots can detect the plume, we evaluated our algorithm with and without the aerial robot. For each configuration and plume height (0.10 m, 0.15 m, 0.20 m, 0.25 m and 0.30 m), we ran 100 simulations and measured the final absolute cross-wind distance error between the formation center and the source location. The error is only calculated along the horizontal lateral direction so as to produce comparable results. The final lateral errors in millimeters, presented in Table I, show that the 4-robot formation consistently outperforms the solution featuring only ground robots in this range, yielding lower mean error with lower standard deviation. Furthermore, the 3-D algorithm is more resilient to increased plume heights, its performance only beginning to degrade when the source reaches 0.25 m.

B. Wind tunnel experiments

The wind tunnel experiments followed our initial simulation trials and were meant to validate the effectiveness of the algorithm using real robots and sensors. The algorithm was tested in our wind tunnel using a team of Khepera IV robots, one of which mounted on a 3-D traversing system to emulate

TABLE I

MEAN ABSOLUTE FINAL CROSSWIND ERROR (LATERAL DISTANCE FROM THE FLOCK CENTER TO THE SOURCE) FOR DIFFERENT SOURCE HEIGHTS. EACH ENTRY (IN MILLIMETERS, ALONG WITH STANDARD DEVIATIONS) IS THE RESULT OF 100 SIMULATION RUNS.

Aerial robot	Source Height (mm)				
	100	150	200	250	300
With	2.5±1.8	2.5±1.7	2.4±1.8	3.3±2.5	5.2±3.8
Without	3.1±2.3	3.3±2.4	3.8±3.4	4.9±3.5	6.6±4.6



Fig. 5. The wind tunnel setup (left), the emulated aerial robot mounted on the traversing system (center), and a Khepera IV robot augmented with odor board and the anemometer board (right).

an aerial robot.

1) *Setup*: The wind tunnel test channel, shown in Fig. 5 (left), has a usable area of $4 \times 18 \text{ m}^2$ and a height of 1.9 m. We are able to generate wind speeds of up to 5 m/s, more than sufficient for emulating realistic scenarios for odor source localization experiments.

The tunnel is equipped with a 3-axis traversing system, seen in Fig. 5 (middle). The system receives commands in the form of goal coordinates and is able to move along the entire length of the tunnel and position instruments precisely. In our experiments, a robot is mounted on the traversing system, facing forward, and used to emulate a UAV.

Absolute positioning inside the tunnel is provided by a 6-camera overhead tracking system running SwisTrack [39]. Our setup is able to track the poses of multiple robots at approximately 20 Hz.

To generate the odor plume, we use an A15-A ethanol release bubbler with a 1.2 l/min air pump and position the outflow hose at the desired source location.

2) *Robots and scenario*: Shown in Fig. 5 (right), the Khepera IV robots [40] are augmented with an odor sensing board and an anemometer board (explained in detail in [41]). Instead of using infrared ranging, for which we only have 2-D capable hardware, the absolute positions received from SwisTrack and the traversing system are merged in an external interface application, and the resulting range, bearing, and elevation data are sent to the robots over the network.

The aerial robot does not actuate its motors, and instead sends commands to an auxiliary program that operates the traversing system. Due to safety and mechanical restrictions on the vertical movement of the traversing system, our operation range extends from a bottom height of 0.23 m to a top height of 0.53 m. The limited range impacts our ability to test the algorithm: any source too close to the floor or above a certain height results in saturation of the vertical control.

We use the same formation shape and graph structure as in

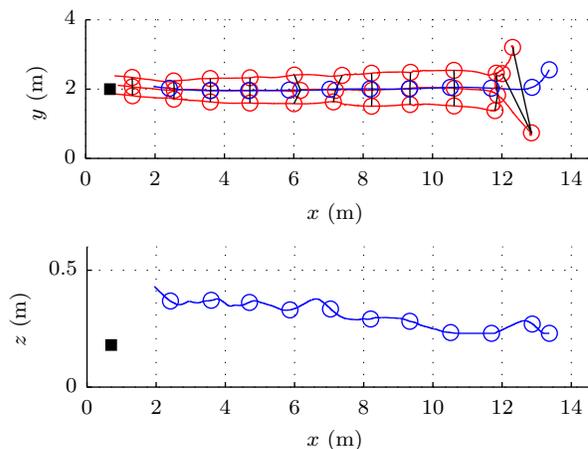


Fig. 6. Top and side view of the trajectories described using the 3-D controller with the odor source placed at $h = 0.18 \text{ m}$. The ground robots are drawn in red and the aerial robot in blue. Circles and black lines are plotted connecting the robot positions every 40 seconds, and the odor source is represented by the black square (size not to scale) on the left side. The trajectories are smoothed over a 1.5 s moving window to suppress SwisTrack noise.

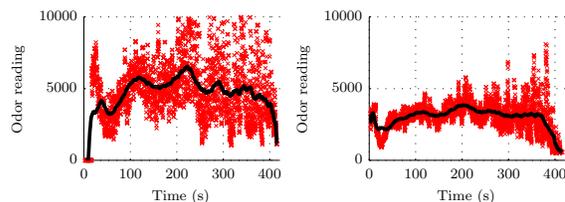


Fig. 7. Odor readings for top robot (left) and bottom robots (right) over time, corresponding to the experiment shown in Fig. 6. The individual readings are plotted along with a 500-sample moving average for readability.

simulation. The three ground robots are in a line formation, with the aerial robot centered but now trailing 1 m behind so as not to obstruct the camera view. The chosen odor source heights for this setup are 0.0, 0.18, 0.37, and 0.5 m, laterally centered in the tunnel (at $y = 2$), near its inlet.

3) *Results*: Fig. 6 is an example that shows the trajectories followed by the robots from their random starting positions (on the right) to the odor source (on the left). The source in this experiment was placed at a height of 0.18 m. The trajectories are smoothed using a sliding window of 1.5 s to remove SwisTrack noise and more accurately match the real paths.

On the ground plane, the robots are quick to align with the source and trace the plume center, although they can be seen slightly drifting on the final 1 m stretch. Because the source is placed higher than the sensors, robots cease to be able to detect the plume as they get close to the source and turn based on spurious environmental readings.

This final concentration drop can be seen on Fig. 7, which shows the concentration readings over time (left-to-right evolution in the time plot corresponds to right-to-left evolution in the position plot), both in raw form and smoothed over a 40-second sliding window for easier reading. The left plot shows the top robot readings and the right plot shows the

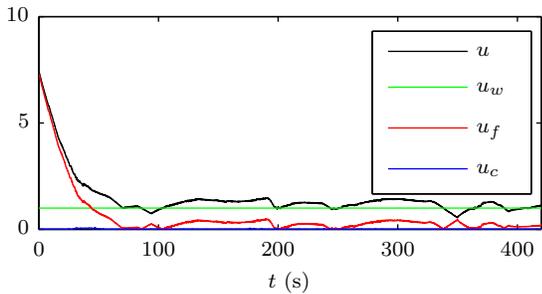


Fig. 8. Influence of each component on the movement of the aerial robot during the experiment shown in Fig. 6. The norm of each vector is plotted over time.

averaged ground robot readings.

The concentration plots also help understand the vertical track, in which we see the aerial robot starting from the lower limit and progressively climbing. For a source height of 0.18 m, the robot remains low for the first 2 m of upwind movement. Afterwards, it starts detecting a higher concentration and climbs to trace the plume, reaching the end at approximately 0.42 m and enclosing the source. The bottom concentration tracks the $\frac{2}{3}$ of the top concentration, as designed.

Fig. 8 shows the influence of each algorithm component on the final $x - y$ movement vector of the aerial robot in the same experiment. As the robots start in random positions, there is a first phase in which formation control dominates while driving the robots to their designated spots; after the formation is achieved, the magnitude of \mathbf{u}_f decreases to under 0.5 (arbitrary units). Plume centering acts through minor adjustments that suffice to steer the group and track the plume. Throughout the experiment, the upwind movement vector \mathbf{u}_w remains constant (in norm, although not in angle) and is the dominant force in the movement of the robot.

The robots adapt their trajectories as long as at least one senses some odor patches. Fig. 9 shows the trajectory of the robots in another example run with a source at 0.37 m. The odor source is elevated but still detectable, so the aerial robot attempts to move up to match the odor concentrations of the ground robots. Due to the range limit of the traversing system, the robot saturates at the maximum height.

The odor readings in Fig. 10 show that, in this experiment, the aerial robot senses significantly higher concentrations than the ground robots, playing an important role in tracking the plume. Working with a single aerial robot, however, the above-ground concentration readings only yield information about the vertical profile of the plume and the plume centering performance degrades as the ground concentration drops. Having two or more flying robots would make it possible to get a better sense of the full 3-D profile and steer the team towards the source even at times when the plume does not reach the ground.

IV. CONCLUSIONS

We presented a distributed algorithm for multi-robot 3-D odor plume tracking using a Laplacian feedback formation

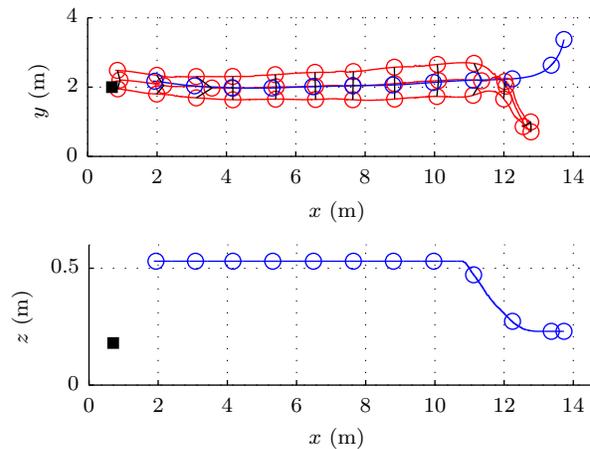


Fig. 9. Top and side view of the trajectories of the robots with the odor source placed at $h = 0.37$ m. The color code is the same used in Fig. 6.

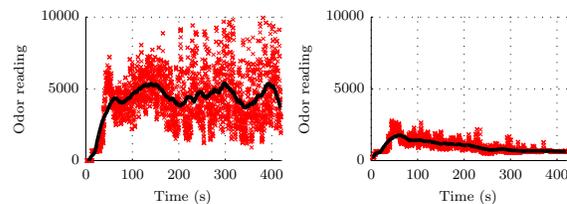


Fig. 10. Odor readings for top robot (left) and bottom robots (right) over time, corresponding to the experiment shown in Fig. 9.

controller. Our approach uses an heterogeneous group of ground and aerial robots to follow a chemical plume to its source based on spatially distributed odor concentration and wind flow measurements.

We provided qualitative and quantitative results for the performance of the algorithm, obtained in a high-fidelity simulation featuring wind and odor modelling. Systematic simulations show that our 3-D algorithm can outperform a similar solution using only ground robots, reducing the final lateral error by up to 36% for source heights at which both the ground and aerial robots are able to detect the plume. Furthermore, the algorithm is able to trace the plume shape in the vertical direction too, providing additional information to the user.

In addition to the performance evaluation obtained in simulation, we also presented real-world results gathered in a wind tunnel, which validated the functionality of the algorithm in a real setup. Despite the constraints of the environment, which prevent us from doing a full quantitative assessment, we experimented with different plume heights and showed that the robots successfully followed the plume to its source.

Our work is an early foray into multi-robot 3-D plume tracking and uses an idealized aerial robot. In the future, we expect to replace our traversing system setup with a quadrotor and study how to cope with and exploit the turbulent airflow generated by the robot. We also intend to expand the number of aerial vehicles, allowing for differential measurements at different heights, as well as experiment with homogeneous

aerial formations. Finally, we intend to move towards more realistic conditions, both within the constraints of the wind tunnel and in outdoor settings.

REFERENCES

- [1] P. Roberts and D. Webster, *Turbulent diffusion*. ASCE Press, Reston, Virginia, 2002.
- [2] R. Russell, D. Thiel, R. Deveza, and A. Mackay-Sim, "A robotic system to locate hazardous chemical leaks," in *Proc. IEEE Int. Conf. on Robotics and Automation*, 1995, pp. 556–561.
- [3] A. Marjovi and L. Marques, "Optimal swarm formation for odor plume finding," *IEEE Trans. on Cybernetics*, vol. 44, no. 12, pp. 2302–2315, 2014.
- [4] W. Li, J. Farrell, and R. Cardé, "Tracking of fluid-advected odor plumes: strategies inspired by insect orientation to pheromone," *Adaptive Behavior*, vol. 9, no. 3-4, p. 143, 2001.
- [5] F. Li, Q.-H. Meng, J.-W. Sun, S. Bai, M. Zeng, et al., "Single odor source declaration by using multiple robots," in *American Institute of Physics (AIP) Conference Proceedings*, vol. 1137, 2009, p. 73.
- [6] A. Lilienthal and T. Duckett, "Experimental analysis of gas-sensitive braitenberg vehicles," *Advanced Robotics*, vol. 18, no. 8, pp. 817–834, 2004.
- [7] A. J. Rutkowski, R. D. Quinn, M. Willis, et al., "A sensor fusion approach to odor source localization inspired by the pheromone tracking behavior of moths," in *Proc. IEEE Int. Conf. on Robotics and Automation*, 2007, pp. 4873–4878.
- [8] J. A. Farrell, S. Pang, and W. Li, "Plume mapping via hidden Markov methods," *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, vol. 33, no. 6, pp. 850–863, 2003.
- [9] M. Vergassola, E. Villermaux, and B. Shraiman, "Infotaxis as a strategy for searching without gradients," *Nature*, vol. 445, no. 7126, pp. 406–409, 2007.
- [10] A. J. Lilienthal, M. Reggente, M. Trincavelli, J. L. Blanco, and J. Gonzalez, "A statistical approach to gas distribution modelling with mobile robots—the Kernel DM+V algorithm," in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 2009, pp. 570–576.
- [11] X. Cui, C. Hardin, R. Ragade, and A. Elmaghraby, "A swarm approach for emission sources localization," in *Proc. IEEE Int. Conf. on Tools with Artificial Intelligence*, 2004, pp. 424–430.
- [12] C. Lytridis, G. Virk, Y. Rebour, and E. Kadar, "Odor-based navigational strategies for mobile agents," *Adaptive Behavior*, vol. 9, no. 3-4, pp. 171–187, 2001.
- [13] L. Marques, U. Nunes, and A. de Almeida, "Particle swarm-based olfactory guided search," *Autonomous Robots*, vol. 20, no. 3, pp. 277–287, 2006.
- [14] A. Marjovi, J. Nunes, P. Sousa, R. Faria, and L. Marques, "An olfactory-based robot swarm navigation method," in *Proc. IEEE Int. Conf. on Robotics and Automation*, 2010, pp. 4958–4963.
- [15] J. Masson, M. B. Bechet, and M. Vergassola, "Chasing information to search in random environments," *Journal of Physics A: Mathematical and Theoretical*, vol. 42, no. 43, p. 434009, 2009.
- [16] Y.-X. Wu, Q.-H. Meng, Y. Zhang, and M. Zeng, "A novel chemical plume tracing method using a mobile sensor network without anemometers," in *Mechanical Engineering and Technology*, 2012, pp. 155–162.
- [17] A. Marjovi and L. Marques, "Multi-robot olfactory search in structured environments," *Robotics and Autonomous Systems*, vol. 52, no. 11, pp. 867–881, 2011.
- [18] D. Zarzhitsky, D. Spears, and W. Spears, "Swarms for chemical plume tracing," in *Proc. of the IEEE Swarm Intelligence Symposium*, 2005, pp. 249–256.
- [19] A. Marjovi and L. Marques, "Swarm robotic plume tracking for intermittent and time-variant odor dispersion," in *Proc. European Conf. on Mobile Robots*, 2013, pp. 379–384.
- [20] J. M. Soares, A. P. Aguiar, A. M. Pascoal, and A. Martinoli, "A distributed formation-based odor source localization algorithm: design, implementation, and wind tunnel evaluation," in *Proc. IEEE Int. Conf. on Robotics and Automation*, 2015, pp. 1830–1836.
- [21] —, "An algorithm for formation-based chemical plume tracing using robotic marine vehicles," in *Proc. MTS/IEEE Oceans*, 2016.
- [22] Y. Kuroki, G. S. Young, and S. E. Haupt, "UAV navigation by an expert system for contaminant mapping with a genetic algorithm," *Expert Systems with Applications*, vol. 37, no. 6, pp. 4687–4697, 2010.
- [23] M. Kovacina, D. Palmer, G. Yang, R. Vaidyanathan, et al., "Multi-agent control algorithms for chemical cloud detection and mapping using unmanned air vehicles," in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, vol. 3, 2002, pp. 2782–2788.
- [24] R. J. Bamberger, D. P. Watson, D. H. Scheidt, and K. L. Moore, "Flight demonstrations of unmanned aerial vehicle swarming concepts," *Johns Hopkins APL Technical Digest*, vol. 27, no. 1, pp. 41–55, 2006.
- [25] W. Wu, D. Chang, and F. Zhang, "A bio-inspired robust 3D plume tracking strategy using mobile sensor networks," in *Proc. IEEE Annual Conference on Decision and Control*, 2013, pp. 4571–4578.
- [26] L. Osório, G. Cabrita, and L. Marques, "Mobile robot odor plume tracking using three dimensional information," in *Proc. European Conf. on Mobile Robots*, 2011, pp. 165–170.
- [27] M. Reggente and A. Lilienthal, "The 3D-Kernel DM+V/W algorithm: using wind information in three dimensional gas distribution modelling with a mobile robot," in *Proc. IEEE Sensors*, 2010, pp. 999–1004.
- [28] S. Badia, U. Bernardet, A. Guanella, P. Pyk, and P. Verschure, "A biologically based chemo-sensing UAV for humanitarian demining," *Int. Journal of Advanced Robotic Systems*, vol. 4, no. 2, pp. 187–198, 2007.
- [29] P. P. Neumann, S. Asadi, A. J. Lilienthal, M. Bartholmai, and J. H. Schiller, "Autonomous gas-sensitive microdrone: wind vector estimation and gas distribution mapping," *IEEE Robotics & Automation Magazine*, vol. 19, no. 1, pp. 50–61, 2012.
- [30] P. P. Neumann, V. H. Bennetts, A. J. Lilienthal, M. Bartholmai, and J. H. Schiller, "Gas source localization with a micro-drone using bio-inspired and particle filter-based algorithms," *Advanced Robotics*, vol. 27, no. 9, pp. 725–738, 2013.
- [31] H. Ishida, K. Yoshikawa, and T. Moriizumi, "Three-dimensional gas-plume tracking using gas sensors and ultrasonic anemometer," in *Proc. IEEE Sensors*, 2004, pp. 1175–1178.
- [32] J. M. Soares, A. P. Aguiar, A. M. Pascoal, and A. Martinoli, "A graph-based formation algorithm for odor plume tracing," in *Proc. Int. Symp. on Distributed Autonomous Robotics Systems*, ser. Springer Tracts in Advanced Robotics, vol. 112, 2014, pp. 255–269.
- [33] T. Lochmatter, E. Göl, I. Navarro, and A. Martinoli, "A plume tracking algorithm based on crosswind formations," in *Proc. Int. Symp. on Distributed Autonomous Robotics Systems*, ser. Springer Tracts in Advanced Robotics, vol. 83, 2010, pp. 91–102.
- [34] M. Mesbahi and M. Egerstedt, *Graph theoretic methods in multiagent networks*. Princeton University Press, 2010.
- [35] S. A. Goyal, "A framework for graph-based distributed rendezvous of nonholonomic multi-robot systems," PhD Thesis 5845, École Polytechnique Fédérale de Lausanne, 2013.
- [36] O. Michel, "Webots: professional mobile robot simulation," *International Journal of Advanced Robotic Systems*, vol. 1, no. 1, pp. 39–42, 2004.
- [37] J. A. Farrell, J. Murlis, X. Long, W. Li, and R. T. Cardé, "Filament-based atmospheric dispersion model to achieve short time-scale structure of odor plumes," *Environmental Fluid Mechanics*, vol. 2, no. 1-2, pp. 143–169, 2002.
- [38] A. Prorok, A. Arfire, A. Bahr, J. R. Farserotu, and A. Martinoli, "Indoor navigation research with the Khepera III mobile robot: an experimental baseline with a case-study on ultra-wideband positioning," in *Proc. Int. Conf. on Indoor Positioning and Indoor Navigation*, 2010.
- [39] T. Lochmatter, P. Roduit, C. Cianci, N. Correll, J. Jacot, and A. Martinoli, "SwisTrack – a flexible open source tracking software for multi-agent systems," in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 2008, pp. 4004–4010.
- [40] J. M. Soares, I. Navarro, and A. Martinoli, "The Khepera IV mobile robot: performance evaluation, sensory data, and software toolbox," in *Proc. Iberian Robotics Conference*, ser. Advances in Intelligent Systems and Computing, vol. 417, 2015, pp. 767–781.
- [41] T. Lochmatter, "Bio-inspired and probabilistic algorithms for distributed odor source localization using mobile robots," PhD Thesis 4628, École Polytechnique Fédérale de Lausanne, 2010.