

Leader-Based versus Soft Control of Multi-Agent Swarms

Guillaume Sartoretti¹

Received: date / Accepted: date

Abstract We focus on the control of heterogeneous swarms of agents that evolve in a random environment. Control is achieved by introducing special agents: leader and infiltrated (shill) agents. A refined distinction is made between hidden and apparent controlling agents. For each case, we provide an analytically solvable example of swarm dynamics.

Keywords multi-agent swarms · leader-based control · soft control · analytical results · numerical simulations

1 INTRODUCTION

Thorough understanding of swarm dynamics ideally requires simultaneous mastery of analytical, simulation-based and experimental approaches. In recent contributions, this ambitious program has been successfully fulfilled for homogeneous swarms, where the mean-field approach can be used [1]. Mean-field theory cannot be applied to heterogeneous swarms; thus, analytical approaches are very difficult. Nevertheless, for specific dynamics, other approaches can be explored in order to obtain exact results, and the goal of this paper is to unveil such possibilities. Specifically, we focus on swarms controlled by the use of special agents (leader or infiltrated (shill) agents). Leader-controlled swarms are commonly used in robotics, and recent contributions presented joint analytical and simulation illustrations [2]. The study of leadership also plays an important role in ethology, where recent papers presented experimental validation of theoretical results [3]. In parallel,

¹EPFL/IMT/LPM, Ecole Polytechnique Fédérale de Lausanne, Lausanne, Switzerland
(Tel: +41-21-693-7778; E-mail: guillaume.sartoretti@epfl.ch)

due to numerous applications, the control of a swarm via shill agents (referred to as "soft-control" [4]) is becoming very popular in robotics and ethology [5,6]. This paper attempts a classification of the different types of swarm control mechanisms based on special agents. Table 1 summarizes the four types of swarm control, depending on whether the shills/leaders can remain hidden within the swarm. Our classification is presented along with analytically tractable illustrations.

Table 1: Classification diagram of the presented swarm control examples.

Special Agent \ Visibility	Leader	Shill
Apparent	Sect. 2.1	Sect. 3.1, depending on shill's behavior
Hidden	Sect. 2.2	

2 LEADER-BASED CONTROL

The first type of swarm control can be achieved by using a number of **leader** agents (possibly only one). In this case, all regular agents must be aware of the presence of the special agents. Accordingly, the dynamics of the regular agents must incorporate the role of the leader agents a priori. The result is a heterogeneous swarm, where regular agents simultaneously follow their own nominal dynamics, and feel the influence of the leader agents. By adequately controlling the leader agents, one can therefore control the whole group.

Within leader-controlled swarms, we further differentiate between two classes of models. The separation

lies in whether the special agents can be identified as such by an external observer of the swarm's behavior. This observer would be aware of the presence of special agents, but would not be allowed access to the underlying swarm's dynamics. The **visibility** of the leader agents (i.e., whether they can be externally recognized) helps us assess their vulnerability. Since a small number of leader agents control the entire swarm, leaders are the main liability of the group. Depending on the application, hiding the leaders' true nature can therefore effectively protect the swarm.

2.1 Apparent Leaders

Apparent leaders usually stand out in the swarm by their positioning. They act as general landmarks for the regular agents, either stationary or moving, and can therefore easily be identified by an external observer. Swarms of agents controlled by apparent leaders have been studied extensively in the literature, and are mentioned here mostly for completeness.

As a classical illustration, we consider a swarm of mobile robots initially positioned in a single file. Our aim is to allow the whole swarm to move in formation, with the first robot acting as the swarm's leader, while the others follow in its path. To this aim, we let each regular mobile robot act as differential-motor, two-sensor light-attracted Braitenberg vehicle [7]. This vehicle regularly updates the speed of its left (resp. right) motor proportionally to the intensity of the light read by its right (resp. left) sensor. In other words, each of our regular robots consistently adapts its course to drive toward the brightest light source its sensors can detect. By attaching a light at the rear of each robot, each regular robot will be attracted toward the robot directly in front of it in the line. Hence, each agent (except for the last one in the file) acts as a (local) leader for the agent robot in the file. Similarly, each regular robot simultaneously acts as a follower, attracted by the previous robot in the file. This effectively enables us to build a "robot train", whose path the first agent in line (i.e., the leading robot) controls (similar to the column formation in [8]).

Fig. 1(a) illustrate the initial state of the robots, and 1(b)-1(d) the selected consecutive states of a swarm of 10 robots. In this run, the leading robot is scripted in order to follow a horizontal figure of eight. This path enables us to test the swarm's resistance to turns induced by the leader. Each robot reacts to the change of direction of the robot in front of it in the file with a little delay. These delays add up, leading to a large difference of path between the front and rear robots in the file. If the control parameters of the Braitenberg

dynamics (i.e., forward and rotation speeds) are not adequately chosen, a turn induced by the leading robot can break the cohesion of the swarm.




Fig. 1: (a): Initial state of a swarm of 10 robots (1 leader (in red) and 9 followers), with all robots facing upwards. (b)-(d): Selected consecutive states of the swarm through time. During this run, the leader follows a horizontal scripted figure of eight, while the rest of the train follows in its path.

2.2 Hidden Leaders

Hidden leader agents are more protected than apparent leaders. Hidden leaders are often encountered in swarms of social animals; a prime example is swarms of bees. When swarming, either moving or resting, regular bees regroup around the queen bee (which acts as the leader). Even when the swarm is moving (and thus less tightly packed around the queen), an (untrained) external observer will not be able to tell the queen bee from the regular bees. This allows the swarm to safely move to a new hive location, led by a leader hidden within the swarm and following a similar dynamics.

A close illustration of such a swarm led by a hidden leader can be found within nonlinear filtering (i.e., estimation problems). An estimation problem models the case of an (usually noisy) input signal $X(t)$ from which noisy measures $Z(t)$ are taken (either continuously or at discrete timesteps):

$$\begin{cases} dX(t) = f(X(t)) + \sigma_x dW_x(t), \\ dZ(t) = g(X(t)) + \sigma_z dW_z(t), \end{cases} \quad (1)$$

where f and g are arbitrary functions, $\sigma_{x,z}$ the noise variances of each process, and $dW_{x,z}(t)$ independent White Gaussian Noise (WGN) sources. After a measure $Z(t)$ is obtained at time t , one wishes to obtain the best estimation of $X(t)$ from the set of all previous measures $\mathcal{Z}(t) := \{Z(T) \mid T \leq t\}$. The probability

distribution $\mathbb{P}[X(t) = x | \mathcal{Z}(t)]$ is usually computed (or empirically constructed) to obtain this estimation. The best estimation for $X(t)$ is computed as the expected value $\mathbb{E}[X(t) | \mathcal{Z}(t)]$.

For linear estimation problems (linear f and g functions in Eq. (1)), the Kalman-Bucy filter is well known to be optimal. In this method, the probability distribution $\mathbb{P}[X(t) = x]$ can be exactly computed a priori by propagating the first line of Eq. (1) from the estimation of X at the last timestep. The distribution $\mathbb{P}[Z(t) = z | X(t) = x]$ is then expressed using the second line of Eq. (1). The estimation is finally extracted from the distribution $\mathbb{P}[X(t) = x | Z(t) = z]$, which can be expressed by Bayesian inversion of $\mathbb{P}[Z(t) = z | X(t) = x]$. For nonlinear problems, a recent filtering mechanism uses feedback particles (i.e., agents) to empirically form the distribution $\mathbb{P}[X(t) = x | Z(t) = z]$ [9]. The Feedback Particle Filter (FPF) lets a swarm of N particles follow the same dynamic f as the noisy input signal $Y(t)$ (g , however, is assumed to be linear). The interaction kernel \mathcal{K} lets the agent self-arrange, based on the noisy measures $dZ(t)$ obtained and on the state $\mathbf{X}(t)$ of the whole swarm. The regular agents $X_i(t)$ ($1 \leq i \leq N$), and the leader $Y(t)$, follow the dynamics:

$$\begin{cases} dX_i(t) = f(X_i(t)) dt \\ \quad + \mathcal{K}(X_i(t), \mathbf{X}(t), dZ(t)) + \sigma dW_i(t), \\ \text{leader's} \quad \begin{cases} dY(t) = f(Y(t)) dt + \sigma dW(t), \\ \text{dynamic} \quad \begin{cases} dZ(t) = h Y(t) dt + \sigma_o dW_y(t). \end{cases} \end{cases} \end{cases} \quad (2)$$

Since each measure $dZ(t)$ obtained from the signal $Y(t)$ is known by each agent, the swarm can be seen as heterogeneous: The noisy input acts as the leader, while the particles act as regular agents [10]. The regular agents try and arrange themselves optimally around the leader, knowing only its noisy instantaneous position. The leader controls the position of the swarm, since the regular agents always self-arrange around the leader's position. The leader also controls the variance of the regular agents' positions around itself, by tuning the parameter σ_o . The larger the noise added to the leader's actual position before it is revealed to the regular agents, the wider the swarm's dispersion around its leader. Agents closer to the leader feel a weaker influence of the self-arranging mechanism (since they are already well arranged). Therefore, the leader and its close neighbors nearly follow the same dynamics, precluding an external observer from spotting the leader from even thorough observation of the swarm's behavior. Fig 2 shows how the leader is able to steer the whole swarm (2(a)), while remaining well hidden among the other agents (2(b)-2(c)).

The results of the FPF applied on a nonlinear filtering problem are usually not analytically tractable. How-

ever, for a class of nonlinear problems where f is derived from Weber Parabolic Cylinder functions, we have shown that the mean and variance of $\mathbb{P}[X(t) = x | Z(t) = z]$ can be explicitly written [10]. This enables us to analytically study the leader's influence on the spreading of the swarm around its position, regarding the noise variance σ_o .

Several shortcomings can be observed with this type of control. First, the leader's influence must explicitly feature in the regular agents' dynamics. Therefore, when one wants to control a specific swarm, its dynamics must be constructed in order to encompass the presence of leader agents. Second, the small number of leaders simplifies the control of a large swarm, but limits its resistance to single agent failures. If a leader does not perform as expected, the control of the whole swarm is weakened.

3 SOFT CONTROL

The second controlling mechanism consists of inserting special agents in a swarm. The infiltrated agents (called **shill** agents) are recognized by the other agents as regular agents. Control of these infiltrated agents can lead to the control of the whole swarm, provided one possesses extensive knowledge of the regular agents' dynamics. This type of control is often referred to as *soft control*, since the swarm is not aware of being infiltrated. The influence of the shill agents is not explicitly included in the dynamic of each agent. Shills control the other




Fig. 2: (a): Positions of $N = 1000$ agents (red) following Eq.(2), arranging around one leader (black). (b)-(c): Distribution of the agents (blue) around the leader (red) at times $t = 5$ and $t = 10$. The function f is the Weber Parabolic Cylinder function with parameter $B = 0.49$ (nonlinear case, see [10]), and $\sigma = \sigma_o = h = 1$.

agents only by taking advantage of the agents' interactions.

Since skills are infiltrated within an unwitting swarm, their actions can positively influence the swarm's dynamics. This constructive soft control can help the swarm fulfill its objectives, by adding the possibility of a real-time external control. Alternatively, soft control can also be used destructively – for example, it may break the swarm cohesion or drive the swarm toward an incorrect direction. Since a skill must control the regular agents in a subtle way, in order to remain unnoticed within the swarm, the number of skill agents must usually be selected according to the number of regular agents.

We introduce a similar distinction between apparent and hidden skill agents. An example in which skills must be apparent to control a swarm can help assess the robustness of the swarm's dynamics. In those cases, an infiltrated swarm can be immediately recognized as such from the outside, precluding destructive skills to quietly take control of the swarm. Cases in which the soft control can be achieved using hidden skills however, help protect the special agent's identity. In the same manner as the leaders, skill agents can be a liability since the whole swarm can be controlled through their single influence. However, since agents in a soft-controlled swarm do not recognize their skills, the swarm can resume functioning on its own (uncontrolled) in their absence.

3.1 Apparent or Hidden Skills

Similar to apparent leaders, apparent skills stand out in a swarm of regular agents. For example, if we consider a swarm of flocking agents constantly driven toward their barycenter. A skill can then steer the group toward a chosen direction by changing its position, therefore changing the global barycentric position. If the skill is too fast, it can be driven outside the swarm and become apparent. An external observer would recognize the skill's action as the cause of the swarm's global movement.

We studied a similar model in which a skill agent is able to drive a stationary swarm of flocking agents toward a selected direction [11]. The skill can be analytically proven to remain hidden or to quit the group depending on its level of turbulence. In this model, a swarm of N Brownian agents ($1 \leq i \leq N$) diffuse on \mathbb{R} following the dynamics of a Hybrid Atlas Model (HAM) [12]:

$$\begin{cases} dY_i d = \left(\sum_{k=1}^N g_k 1_{Q_k(i)} \{Y d\} + \gamma_i + \gamma \right) dt \\ \quad + \sigma_i dW_i(t), \\ Y_i(0) = y_i, \end{cases} \quad (3)$$

with $\sigma_i \in \mathbb{R}$ the respective noise variances of the N independent WGN sources $dW_i(t)$. In the HAM, each agent's drift is constructed from a global barycentric drift γ , an agent-based drift γ_i , and a rank-based drift g_i . The rank-based drift depends on the agent's position within the swarm. The first agent has the largest position $Y_i(t)$, and the last one the smallest $Y_i(t)$. The agent-specific and barycentric drifts are time-independent, whereas the rank-based drift is updated constantly with time.

Ichiba et al. showed that a HAM swarm achieves flocking iff each agent is constantly driven toward the barycenter of the swarm [12]. This flocking conditions translates into the set of constraints

$$\sum_{k=1}^l [g_k + \gamma_{p(l)}] < 0, \quad (4)$$

for $l \in \{1, \dots, N-1\}$ and all possible permutations $p = (p(1), \dots, p(N)) \in \Sigma_N$ – i.e., $N!(N-1)$ constraints.

Without loss of generality, we can always assume that $\sum_{k=1}^N [g_k + \gamma_k] = 0$. In other words, the average barycentric speed of the swarm is governed only by γ , which we will assume vanishes (stationary swarm). We assume that agent-specific drifts γ_i vanish $\forall i$, and let $g_i = -g$ ($1 \leq i \leq N-1$) and $g_N = (N-1)g$ (with $g \in \mathbb{R}$). This dynamics fulfills the flocking conditions of Eq. (4). This choice of rank-based drifts means that every agent is pushed back with a drift $-g$, except for the last agent which is strongly pushed forward with a drift $(N-1)g$. In our flocking swarm, we then introduce our skill agent, agent 1 (without loss of generality), following the same HAM dynamics but driven by a Ballistic Noise source $dZ(t)$:

$$dZ(t) = \beta \tanh[\beta Z(t)] dt + dW(t), \quad Z(0) = 0. \quad (5)$$

Despite its nonlinear nature, this specific noise source can be reduced to a Brownian motion with constant drift $\pm\beta$ [13]. In other words, one realization of the ballistic noise consists of an initial random draw deciding whether the constant drift will be $+\beta$ or $-\beta$, with equal probability. The rest of the realization will then just be that of a Brownian motion with constant drift: $dZ(t) = \pm\beta + dW(t)$.

Therefore, our ballistic skill agent follows the same HAM dynamics as its regular fellows, but with a constant extra drift $\pm\beta$ (depending on the realization). The skill steers the (initially stationary) swarm toward by inducing an average barycentric drift $\pm\frac{\beta}{N}$. For small β , flocking is still achieved within the swarm, and the skill will be hidden while still steering the swarm. However, for large β , the skill agent extracts itself from the swarm, and diverges from the other agents. The skill therefore becomes apparent, obviously being the reason for the swarm's movement to an external observer.

Our dynamic is asymmetric in rank-based drifts: The last agent is strongly pushed forward, whereas the others are gently pushed backward. Therefore, the shill's visibility also depends on the realization of its ballistic noise source. When it gets a $-\beta$ extra drift, the shill most likely always is ranked last in the swarm, and thus strongly pushed forward toward the rest of the swarm. Conversely, during realization with an extra $+\beta$ drift, the shill is most likely ranked first, and thus only gently driven toward the other agents. Therefore, three outcomes arise:

1. The shill can always remain hidden.
2. The shill is hidden for its $-\beta$ realizations, but apparent for the $+\beta$ ones.
3. The shill is always apparent.

Analytical investigation of the model allowed us to express the thresholds values for β , marking the transition between each of these outcomes [11].

Fig. 3 shows the two types of realizations, when an initially stationary swarm is infiltrated by a ballistic shill for the second outcome. The shill's visibility depends on the realization of its noise source (outcome 2.). In the first realization ($-\beta$ drift), the shill is able to induce a negative average drift to the swarm while remaining flocked with the regular agents. But in the second realization where it gets a $+\beta$ extra drift, the turbulent shill becomes apparent, unable to remain within the swarm.




Fig. 3: Two realizations of the HAM dynamics of Eq.(3) with a swarm of $N = 9$ regular agents and 1 shill. Here, $\beta = 1.5$, $\sigma = g = 1$, which corresponds to the outcome 2. (i.e., the shill's visibility depends on the realization of its noise source). The upper realization shows the $+\beta$ realization, where the shill (black) becomes apparent from the regular agents (green). The lower trajectories show a $-\beta$ realization, in which the shill (black) steers the swarm of regular agents (red) while remaining flocked.

4 CONCLUSION

Depending on the application, we believe that our classification helps select the appropriate control mechanism. The inherent need for robustness of the resulting control mechanism helps dictate the visibility of

the special agents. In parallel, the desired "subtlety" of the control type can help choose between soft and hard control. Although we believe our classification to be exhaustive, whether further refinements are needed remains an open question at this stage.

ACKNOWLEDGEMENTS

We thank M.-O. Hongler for numerous discussions in the writing of this paper.

The author is funded by the Swiss National Science Foundation.

References

1. G. Sartoretti, Hongler M.-O., M. Elias de Oliveira and F. Mondada. Decentralized self-selection of swarm trajectories: From dynamical system theory to robotic implementation. *Swarm Intelligence*, vol. 8(no. 4):329-351, 2014.
2. M. Michini, H. Rastgoftar, M. A. Hsieh, and S. Jayasuriya. Distributed formation control for collaborative tracking of manifolds in flows. In *Proceedings of the American Control Conference*, pages 3874-3880, 2014.
3. B. Collignon, J. L. Deneubourg, and C. Detrain. Leader-based and self-organized communication: Modelling group-mass recruitment in ants. *Journal of theoretical biology*, 313:7986, 2012.
4. Jing Han, Ming Li, and Lei Guo. Soft control on collective behavior of a group of autonomous agents by a shill agent. *Journal of Systems Science and Complexity*, 19(1):5462, 2006.
5. Deneubourg J. L. H. Bleuler Gribovskiy A., Halloy J. and F. Mondada. Towards mixed societies of chickens and robots. In *Proceedings of the International Conference on Intelligent Robots and Systems* pp. 4722-4728, 2010.
6. J. Krause, A. F. T. Winfield, and J. . Deneubourg. Interactive robots in experimental biology. *Trends in Ecology and Evolution*, 26(7):369375, 2011.
7. V. Braitenberg. *Vehicles: Experiments in Synthetic Psychology*. MIT Press, Cambridge, MA, 1984.
8. K. N. Krishnanand and D. Ghose. Formations of minimalist mobile robots using local-templates and spatially distributed interactions. *Robotics and Autonomous Systems*, 53(3-4):194213, 2005.
9. T. Yang, P. G. Mehta, and S. P. Meyn. Feedback particle filter. *IEEE Transactions on Automatic Control*, 58(10):2465-2480, 2013.
10. G. Sartoretti, M.-O. Hongler, and R. Filliger. The estimation problem and heterogeneous swarms of autonomous agents. In *Proceedings of the Stochastic Modeling Techniques and Data Analysis International Conference*, 2014.
11. G. Sartoretti and M.-O. Hongler. Soft control of swarms: Analytical approach. In *Proceedings of the International Conference on Agents and Artificial Intelligence*, volume 1, pages 147-153, 2013.
12. Banner A. I. Karatzas Ichiba T., Papathanakos V. and R. Fernhold. Hybrid atlas model. *Ann. Appl. Prob.* 21(2) 609-644, 2011.
13. R. Filliger Hongler M. O. and P. Blanchard. Soluble models for dynamics driven by a super-diffusive noise. *Physica A* 370, 301-355, 2006.