



ECOLE
POLYTECHNIQUE
DE BRUXELLES

UNIVERSITÉ LIBRE DE BRUXELLES

Ecole Polytechnique de Bruxelles

IRIDIA - Institut de Recherches Interdisciplinaires
et de Développements en Intelligence Artificielle

Information Transfer in a Flocking Robot Swarm

Eliseo FERRANTE

Promoteur de Thèse:
Prof. **Marco DORIGO**

Co-Promoteur de Thèse:
Prof. **Mauro BIRATTARI**
Prof. **Ali Emre TURGUT**

Thèse présentée en vue de l'obtention du titre de
Docteur en Sciences de l'Ingénieur

Année académique 2012-2013

Abstract

In this dissertation, we propose and study methods for information transfer within a swarm of mobile robots that coordinately move, or flock, in a common direction. We define information transfer as the process whereby robots share directional information in order to coordinate their heading direction. We identify two paradigms of information transfer: explicit information transfer and implicit information transfer.

In explicit information transfer, directional information is transferred via communication. Explicit information transfer requires mobile robots equipped with a communication device. We propose novel communication strategies for explicit information transfer, and we perform flocking experiments in different situations: with one or two desired directions of motion that can be static or change over time. We perform experiments in simulation and with real robots. Furthermore, we show that the same explicit information transfer strategies can also be applied to another collective behavior: collective transport with obstacle avoidance.

In implicit information transfer, directional information is transferred without communication. We show that a simple motion control method is sufficient to guarantee cohesive and aligned motion without resorting to communication or elaborate sensing. We analyze the motion control method for its capability to achieve flocking with and without a desired direction of motion, both in simulation and using real robots. Furthermore, to better understand its underlying mechanism, we study this method using tools of statistical physics, showing that the process can be explained in terms of non-linear elasticity and energy-cascading dynamics.

To my dad

Acknowledgments

Differently from some other PhD students, my adventure at IRIDIA started in Italy rather than in Belgium, in what I remember being December 2007. Prof. Marco Dorigo had come to visit some friends in Milan and decided to give me the chance to have an interview. We therefore invaded one of the professors offices in Politecnico di Milano, did the interview, and few days later I was in Brussels presenting my master thesis. Two weeks later, my PhD had already started.

I am extremely thankful to Prof. Dorigo for giving me the chance to join the IRIDIA family. He gave me excellent conditions for doing research outside my home country in an environment that still made me feel home. I also thank him for the supervision he has been giving me through these years and for his continuous support in all the other aspects of researcher life. I would like to thank also Prof. Mauro Birattari, who also supervised my thesis and other projects I completed during my period at IRIDIA. The frequent interactions I had with him made me learn a lot. Among other important aspects, he made me aware of the importance of scientific communication, and taught me many skills I now use every day to make my work convincing and bullet-proof. He also fostered my leadership skills, by allowing me to supervise bachelor and master students for all the four years I've spent in Brussels. Last but absolutely not least, my deepest thanks go to Prof. Ali Emre Turgut (Alio), also my co-supervisor and, more importantly, one of my best friends. I owe to Alio a lot: he not only followed me in all the aspects of my thesis, but also in all the other aspects of my life since I met him. He has been there in every moment, being them happy or "lunatic" ones, during late evenings in the lab or lunch/dinner at the weird guy or at Eskxi (I wish I learned how to spell this). My deepest thanks go also to Dr. Cristián Huepe, an amazing scientist that taught me a lot, made me work on extremely interesting homeworks, had with me many uncountable skype discussions, and made me feel a statistical physics apprentice. Thank to all of you for making this dissertation possible.

I should now express my gratitude to all the IRIDIA members I've encountered during my adventure. This is maybe the hardest part, as the probability to forget somebody is very high. Indeed, I joined the lab during a "generation shift", while

the “old” (or the “very old” as I upgraded myself to “old” now) IRIDIANS were leaving. Among them, I remember and I would like to thank Christos, Elio, Shervin, Anders and Alex, for the short but still useful chats I had with them that helped me to “bootstrap” my experience in the lab. I should now thank Cardo, the oldest of the old IRIDIANS (but not of the very old), who gifted me with his friendship, three versions of the simulator (trying the third one at the time of writing), many collaborations and for allowing me to graduate before him. I also would like to thank Giovanni, or Gianni the first (henceforth G1), who made me steal the previous sentence from his acknowledgments :-), for his incredible administration of the IRIDIA server, for the incredible cooperativeness exhibited during all these years and in particular during the Dark Ages of Lausanne. Thanks also to Arnucci, for teaching me the brutal ways to say NO, for the administration of the cluster (first period) and for all the technical stuff he taught me (together with Cardo). Thanks to Mario Montes and Nithigno, for all the friendly time together and for starting together the first scientific collaboration I ever had in IRIDIA. Thanks to all the other scientists in IRIDIA who had a collaboration with me: to Bambrile (Manu), for the big adventure on the swarm robotics survey, for the one on collective transport and for the short period spent together in Lausanne (OK, Bambri, I know for you it was much longer :D); to Westy (Alessandro), for the incredible work and passion he put on the flocking project we did together and on the development of the vision tools I used in all my flocking projects thereafter; to Lucia, for being my first master student and for the amazing work she did. Thanks to all the other old and new IRIDIANS: to Prasanna, for all the fun we had together; to Giacomotto d.C. (yes you belong to IRIDIA too :P), for all dinners, musical and other topics-related chats and for coming to eat horse meat in Santeramo; to Gabri, for sharing similar passions with me about both science and life; to Giancy, for sharing terronian blood with me and for helping me moving to Leuven despite he was relying only on one knee; to Maria, Giancy’s girlfriend, for the heel lekkere etenjes she cooked (LASAGNA!) and for the languages tutorials; to the other “naturalized IRIDIAN” such as Rachele, Teresa, Elisa and so on; to G2, for the time spent with me and Alio in weird guy’s restaurant and not; to Sabri, for making me feel useful/knowledgeable with her C-related questions; to Leonarcio, for the few but rewarding Xbox 360 sessions at my place (with or without G2). Thanks also to the very new IRIDIANS: to Vito (well, in a way you are new or renewed :-)); to Leslie, Turaggio, Lorenzo, DJ, Gaetano, Francuccio, Romano and all the others that might have joined IRIDIA after I left Brussels: I strongly believe you all are worthy successors. Thanks also to those who visited IRIDIA, Brussels, or Belgium only for a short period: to Brizzukka, his antitapas, “fasolatas”, cotechinos (even the rotten

ones) and the many “champions”, “warriors” and “marvels” sessions we had while he was my couch guest; to the already mentioned Giacomotto; to Mattia; to Sambo & Matteo Scossotti; to Barbara & Saretta; ... and also to you, yes you, that I for sure forgot. Special thanks to all those that helped me improve the thesis: Prof. Dorigo, Mauro, Alio, G1, Ale, DJ, Giancy, Arne, Manu, Gabri, and G2.

In May 2012 I moved to another wonderful city, Leuven, where I had the chance to meet other amazing people that helped me growing both professionally and personally. I first would like to thank Prof. Tom Wenseleers, for allowing me to join his unique lab and to work (and mainly party) with other unique people on a new exciting project. I also thank him for introducing me for the first time to biology and to its (to me) completely different way of thinking. Thank you Edgar, for your senior advices on my new and also old projects and I wish you the best in your soon to be started career. Thanks also to the other people of the Socio² lab: Jelle, Annetch, Ricardo, Cintia, Sofie, Loren, An, Fred, Dries, Sharon, Kristof, Anneleen, Karolien, Prof. Billen, Conny. Thanks also to Davide C, Ann and Cusco. You made me feel at home from day 1. I would like to thank also all the new friend, met either in the ILT Dutch lessons or in the Oude Markt (where else?). Thank you Frenghie, Dami, Vladje, Bruna, Rocco, Umberto, Marco, Silvia, etc ... and to you that I’m also forgetting. Very special thanks go to Xenia: you are very important to me and you show you care for me a lot. This helped me setting straight the priorities in my life, and gave me back motivation and life energy that have been essential in getting to the end of this adventure.

I’ve dedicated this thesis to my dad, who has been supporting all my life adventures and he continued to do so until he could. I would like to thank you and I really wished you were reading these lines and coming at my PhD defense. I would like to thank my mom for the same reasons, and the rest of my family who is always welcoming me home. Thanks also to the other friends in Santeramo and Milan who make me always feel as I never left the country. Thank you Carles, Lelliwood, Cugi, Tostis, Peppino, Atariello, Nicola, and all the others. Finally, I acknowledge the support of Sony, Microsoft, Nintendo, Square Enix, Bioware, Ubisoft, and many other companies who developed many wonderful consoles and games that helped me getting through even the hardest moments. I’m sure I’m not mentioned everybody. This is only because of space reasons, and not because I care less about you. Many many thanks to everybody.

Contents

Abstract	iii
Acknowledgments	vii
Contents	xiv
1 Introduction	1
1.1 Original contributions and related publications	3
1.2 Other publications not related to this dissertation	5
1.3 Dissertation outline	8
2 Background and related work	11
2.1 Swarm intelligence	11
2.1.1 Naturally evolved swarm intelligence systems	12
2.1.2 Artificially engineered swarm intelligence systems	14
2.1.3 Characteristics and properties of swarm intelligence systems	15
2.2 Swarm robotics	16
2.2.1 Collective behaviors	17
2.2.2 Design methods	19
2.2.3 Modeling methods	20
2.3 Information transfer in biological systems	21
2.3.1 Signal-based information transfer	22
2.3.2 Cue-based information transfer	23
2.4 Flocking	23
2.4.1 Flocking in biology	25
2.4.2 Flocking in statistical physics	27
2.4.3 Flocking in control theory	29
2.4.4 Flocking in robotics	31

3	Self-organizing flocking method	35
3.1	Explicit and implicit information transfer	35
3.2	Flocking control method	36
3.3	Cohesion via proximal control	39
3.4	Explicit information transfer via alignment control	40
3.4.1	Heading communication strategy (HCS)	41
3.4.2	Information-aware communication strategy (ICS)	41
3.4.3	Self-adaptive communication strategy (SCS)	41
3.5	Implicit information transfer via motion control	43
3.5.1	Magnitude dependent motion control (MDMC)	43
3.5.2	Magnitude independent motion control (MIMC)	44
3.5.3	Wheel speed computation	44
4	Explicit information transfer via communication	45
4.1	Problem description	45
4.2	Metrics	47
4.3	Experimental setup	48
4.3.1	Simulation experimental setup	48
4.3.2	Real robot experimental setup	52
4.4	One stationary goal direction	56
4.4.1	Simulation	56
4.4.2	Real robots	59
4.5	One non-stationary goal direction	59
4.5.1	Simulation	59
4.5.2	Real robots	59
4.6	Two stationary goal directions	60
4.7	Two non-stationary goal directions	62
4.7.1	Simulation	62
4.7.2	Real robots	68
4.8	Summary and Discussion	69
5	Explicit information transfer in collective transport	73
5.1	Problem description	73
5.2	Explicit information transfer and collective transport	74
5.2.1	Social mediation	75
5.2.2	Collective transport	76
5.3	Experimental setup	78
5.4	Results	79

5.5	Summary and Discussion	80
6	Implicit information transfer via motion Control	83
6.1	Problem description	83
6.2	Metrics	84
6.3	Experimental setup	86
6.3.1	Simulation experimental setup	86
6.3.2	Real robot experimental setup	90
6.4	Hybrid explicit-implicit information transfer	90
6.4.1	Results in the MIMC-MIMC case	90
6.4.2	Results in the MIMC-MDMC case	91
6.5	Implicit information transfer without informed robots	94
6.5.1	Simulations	94
6.5.2	Real robots	97
6.6	Implicit information transfer with informed robots	97
6.6.1	Simulations	97
6.6.2	Real robots	100
6.7	Summary and discussion	101
7	Statistical physics analysis of implicit information transfer	105
7.1	Problem description	107
7.2	Active Elastic Sheet model	108
7.3	Numerical dynamics and stationary solutions	109
7.4	Convergence dynamics	114
7.4.1	Energy cascading mechanism	114
7.4.2	Dependence on system size	116
7.5	Linear stability analysis	118
7.6	Exploring Active Elastic Sheet dynamics	119
7.6.1	Dynamics of a ring-shaped configuration	119
7.6.2	Propagation of perturbations	120
7.6.3	Heterogeneous self-propulsion speeds	121
7.7	Summary and discussion	122
8	Conclusions and future work	125
8.1	Contributions	125
8.2	Future work	126

A	The robot and the simulation platform	129
A.1	The foot-bot and the implementation	129
A.1.1	The foot-bot	129
A.1.2	Flocking with foot-bots	131
A.1.3	Collective transport with foot-bots	132
A.2	The ARGoS simulator	132
	Bibliography	148

Chapter 1

Introduction

Robots are nowadays an essential gear in guaranteeing that our society remains functional. Many tasks faced by humans are today performed by robots. Examples can be found in many sectors of industry, from food processing to advanced technology and car manufacturing. These tasks are characterized by a set of well-defined operations that, when carried out by robots, can be performed at a faster pace.

There are however other tasks for which robotics does not provide yet a complete solution but for which research is still ongoing. Examples of such tasks are large area exploration and data gathering, surveillance (e.g. of facilities, of crowds, of large human events), search and rescue, mine-clearance, monitoring of natural disasters (i.e., floods, fires) and of accidents (i.e., plume of toxic gas tracking, concentration of pollutants measuring), and many more. All these tasks are carried out in a non-structured environment which might change over time and require a degree of flexibility and autonomy that cannot be found in industrial robots. Furthermore, given their space and time extension, these tasks are more efficiently tackled by large groups of autonomous robots rather than by a single robot with complex hardware: groups of robots can cooperate to cover larger areas, can work in parallel to increase efficiency and can exhibit redundant characteristics to increase fault tolerance. A promising approach for controlling a large groups of autonomous robots is represented by *swarm robotics*.

Swarm robotics is the study and the design of collective behaviors for swarms of autonomous robots that can be used to tackle tasks that require high degree of autonomy and flexibility ([Şahin, 2005](#); [Brambilla et al., 2013](#)). In swarm robotics, the desired collective behavior results from local interactions between the robots themselves and between the robots and the environment. Swarm robotics collective behaviors are flexible (with respect to different environments), robust (to robot failures)

and scalable (to different swarm and problem sizes). These three properties are promoted by the working principles of swarm robotics, that are self-organization and local interaction as opposed to centralized coordination and global communication.

The research presented in this dissertation is motivated by the pursue of design methods for displacing large number of very simple robots in space. The work presented here summarizes some significant steps in the context of self-organized flocking, which studies the cohesive and coordinated motion of a group of robots. Flocking is usually studied by different research fields, including but not limited to biology, statistical physics, control theory and robotics. This dissertation presents novel contributions in the robotics (Chapter 4, 5 and 6) and in the statistical physics (Chapter 7) fields.

Coordinating the movement of a swarm of robots can be advantageous compared to relocating the robots one by one. For instance, a group of robots moving together can have a much increased sensing range compared to a single robot, and this might lead to improved performance in search and rescue or data acquisition tasks to be performed at a particular location in the environment. Additionally, a swarm of flying or underwater robots flocking together can produce an energy efficient motion, in the same way as reported by studies in biology for shoals of fish (Hoare et al., 2000) or flocks of birds (Newton, 2010). In dangerous applications, such as demining or in the military domain, a swarm of robots flocking together can lead to increased robustness, whereby the incapacitation of one or few of the robots can increase the level of alert of the other robots in the flocking group in order to minimize further losses and to enable the continuation of the mission. Finally, studying the principles underlying collective motion can also lead to applications outside robotics, such as the possibility to realize realistic and stunning computer simulations for the videogames or for the movie industry, as pioneered by Reynolds (1987).

In this dissertation, we put particular emphasis on the study of information transfer in self-organized flocking. Information transfer is a key process for the understanding of flocking (Sumpter et al., 2008), but also in general for the design of algorithms and methods for displacing large groups of agents or robots. In flocking, information transfer can be defined as the process whereby individuals share information about their orientation or their desired direction of motion.

In the remaining of this chapter we summarize the main contributions of this dissertation by describing the publications directly linked to its main topic: information transfer (Section 1.1). We then list our other scientific publications not related to this dissertation (Section 1.2). We conclude by describing how the chapters of this dissertation are organized (Section 1.3).

1.1 Original contributions and related publications

In this dissertation, we identify two possible mechanisms of information transfer in flocking: explicit or implicit. In *explicit information transfer*, robots equipped with communication devices can share directional information; in *implicit information transfer*, simpler robots that do not communicate explicitly are also able to self-organize into flocking.

In Chapter 4, we present experiments on explicit information transfer. The contents of Chapter 4 are based on the following publications. In the first, we developed the first information transfer algorithm that can be used in presence of a single static or changing goal direction:

- **E. Ferrante**, A. E. Turgut, N. Mathews, M. Birattari, and M. Dorigo. Flocking in stationary and non-stationary environments: A novel communication strategy for heading alignment. In *Parallel Problem Solving from Nature PPSN XI*, volume 6239 of *Lecture Notes in Computer Science*, pages 331-340. Springer, Berlin, Germany, 2010.

In an extension, we considered the case where a second goal direction, with higher priority with respect to the first, is present. Here, the swarm needs to keep cohesion while moving either in the first direction or in the second direction when present:

- **E. Ferrante**, A. E. Turgut, A. Stranieri, C. Pinciroli, M. Birattari, and M. Dorigo. A self-adaptive communication strategy for flocking in stationary and non-stationary environments. *Natural Computing*, 2013. In press.

We additionally considered the case in which two goal directions with equal priority are present:

- **E. Ferrante**, S. Wenjie, A. E. Turgut, M. Birattari, M. Dorigo., and T. Wenseleers. Self-organized flocking with conflicting goal directions. In *Proceedings of the 12th European Conference on Complex Systems (ECCS 2012)*, Berlin, Germany, 2012. Springer.

In Chapter 5, we show how the same information transfer mechanism studied in Chapter 4 can be employed in another collective behavior: collective transport. The results of this work have been used in the Swarmanoid project (see Section 1.2) and have been published in:

- **E. Ferrante**, M. Brambilla, M. Birattari, and M. Dorigo. Socially-mediated negotiation for obstacle avoidance in collective transport. *International Symposium on Distributed Autonomous Robotics Systems (DARS-2010)*, volume 83

of Springer Tracts in Advanced Robotics, pages 571-583. Springer, Berlin, Germany, 2013.

- **E. Ferrante**, M. Brambilla, M. Birattari, and M. Dorigo. “Look out!”: Socially-Mediated Obstacle Avoidance in Collective Transport. In *Swarm Intelligence: 7th International Conference, ANTS 2010*, volume 6234 of *Lecture Notes in Computer Science*, pages 572-573. Springer, Berlin, Germany, 2010.

Implicit information transfer experiments are presented in Chapter 6. The contents of the chapter is largely based on the following publication:

- **E. Ferrante**, A. E. Turgut, C. Huepe, A. Stranieri, C. Pinciroli, and M. Dorigo. Self-organized flocking with a mobile robot swarm: a novel motion control method. *Adaptive Behavior*, 20(6):460-477, 2012.

In the same chapter, we also consider the case in which the swarm is heterogeneous. Here, a part of the swarm uses explicit information transfer while the rest uses implicit information transfer. The results of this work have been published in:

- A. Stranieri, **E. Ferrante**, A. E. Turgut, V. Trianni, C. Pinciroli, M. Birattari, and M. Dorigo. Self-Organized flocking with a heterogeneous mobile robot swarm. In *Proceedings of ECAL 2011*, pages 789-796, Cambridge, Massachusetts, 2011. MIT Press.

In the context of implicit information transfer, we performed an analysis from the perspective of statistical physics. The analysis is described in Chapter 7 and is reported in the following two articles:

- **E. Ferrante**, A. E. Turgut, M. Dorigo, and C. Huepe. Elasticity-driven collective motion in active solids and active crystals, 2013. Submitted to *Physical Review Letters*. Technical Report <http://arxiv.org/abs/1301.2620>
- **E. Ferrante**, A. E. Turgut, M. Dorigo, and C. Huepe. Collective motion dynamics of active solids and active crystals, 2013. *New Journal of Physics*, 2013. In press.

Finally, we published a paper that can be considered a short summary of Chapter 4 and 6:

- **E. Ferrante**, A. E. Turgut, C. Huepe, M. Birattari, M. Dorigo., and T. Wenseleers. Explicit and implicit directional information transfer in collective motion. In *Artificial Life 13*, volume 13, pages 551-552. MIT Press, 2012.

1.2 Other publications not related to this dissertation

Here, we present our other scientific contributions done during the doctoral studies that are not related to information transfer in flocking and, as such, are not in the framework of this dissertation. The first important contribution is an extensive review of the swarm robotics literature. This review has been published in:

- M. Brambilla, E. Ferrante, M. Birattari, and M. Dorigo. Swarm robotics: A review from the swarm engineering perspective. *Swarm Intelligence*, 7:1-41, 2013.

Our remaining publications can be divided in four categories.

Machine learning The first category contains two contributions performed in the context of machine learning.

Our first contribution belongs to reinforcement learning. Reinforcement Learning (Kaelbling et al., 1996) considers an agent (i.e., a robot) that learns a behavior based on a reward obtained from the environment. We consider a sequence of learning tasks, and we tackle the *transfer learning* problem, whereby part of the solution of a task is transferred to the subsequent task in order to speed-up learning. We devised a generic method that allows automatic transfer of knowledge pertaining to the environment and to the objective. The results of this study have been published in:

- E. Ferrante, A. Lazaric, and M. Restelli. Transfer of task representation in reinforcement learning using policy-based proto-value functions. In Proceedings of the 7th international joint conference on autonomous agents and multiagent systems - Volume 3, AAMAS 08, pages 1329-1332, Richland, SC, 2008. International Foundation for Autonomous Agents and Multiagent Systems.

Our second contribution within machine learning was performed in the context of evolutionary swarm robotics. We developed a method for the automatic design of swarm robotics collective behaviors that is based on the concept of grammatical evolution, a variant of genetic programming:

- E. Ferrante, E. Duéñez Guzmán, A. E. Turgut, and T. Wenseleers. Geswarm: Grammatical evolution for the automatic synthesis of collective behaviors in swarm robotics. In Proceedings of the fifteenth international conference on Genetic and evolutionary computation conference companion (GECCO 2013). ACM, New York, NY, USA, 2013. In press.

Opinion dynamics The second category contains works related to opinion dynamics, which studies the dynamics of opinion spread in a group of individuals ([Castellano et al., 2009](#)). In a first publication, we developed a simple opinion dynamics model based on particle swarm optimization ([Kennedy and Eberhart, 1995](#)). The model was used to achieve collective decision making (path selection) in a swarm of robots:

- M.A. Montes de Oca, **E. Ferrante**, N. Mathews, M. Birattari, and M. Dorigo. Optimal collective decision-making through social influence and different action execution times. In Proceedings of the Workshop on Organisation, Cooperation and Emergence in Social Learning Agents of the European Conference on Artificial Life (ECAL 2009). 2009.

In a follow-up study, two other mechanisms for opinion dynamics were studied: the majority and the expert rules. The results of this study have been published in:

- M.A. Montes de Oca, **E. Ferrante**, N. Mathews, M. Birattari, and M. Dorigo. Opinion dynamics for decentralized decision-making in a robot swarm. In Marco Dorigo et al., editors, LNCS 6234. Swarm Intelligence. 7th International Conference, ANTS 2010, pages 251-262. Springer, Berlin, Germany, 2010.

The majority rule was then further analyzed. We developed a test case scenario in swarm robotics involving a large swarm forming several groups of three robots. Each group had to perform collective transport and choose the shortest path in a double-bridge environment. This study also included an analytical model and was published in:

- M.A. Montes de Oca, **E. Ferrante**, A. Scheidler, C. Pinciroli, M. Birattari, and M. Dorigo. Majority-rule opinion dynamics with differential latency: A mechanism for self-organized collective decision-making. *Swarm Intelligence*, 5(3-4):305-327, 2011.

Other two opinion dynamics rules were developed. The first, called the k-unanimity rule, resulted in a yet to be submitted article and in one publication in a video proceedings:

- A. Scheidler, A. Brutschy, E. Ferrante, and M. dorigo. The k-unanimity rule for self-organized decision making in swarms of robots. 2011. Technical Report TR/IRIDIA/2011-023.
- A. Brutschy, A. Scheidler, **E. Ferrante E.**, M. Dorigo, and M. Birattari. "Can ants inspire robots?" Self-organized decision making in robotic swarms. In *Intelli-*

gent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on, pages 4272-4273, Los Alamitos, CA, 2012. IEEE Computer Society Press.

The second, based on a simple exponential smoothing rule, was published in:

- M. A. Montes de Oca, **E. Ferrante**, A. Scheidler, and L. F. Rossi. Binary consensus via exponential smoothing. In R. Colbaugh et al., editors, Proceedings of the Second International Conference on Complex Sciences: Theory and Applications (COMPLEX 2012). 2013. In press.

Group selection In the third category, we place one publication in which we present a protocol able to select a single or a group of robots in a larger swarm of mobile robots using an aerial robot:

- N. Mathews, A. L. Christensen, **E. Ferrante**, R. O'Grady, and M. Dorigo. Establishing spatially targeted communication in a heterogeneous robot swarm. In Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1 - Volume 1, AAMAS '10, pages 939-946, Richland, SC, 2010. International Foundation for Autonomous Agents and Multiagent Systems.

Swarmanoid The remaining category includes work that has been performed in the context of the Swarmanoid project¹. Swarmanoid was a Future and Emerging Technologies project funded by the European commission, which came to an end in November 2010. The project objective was the study and implementation of a distributed heterogeneous robotics platform composed of three types of robot: the eye-bots, which can fly and attach to the ceiling; the hand-bots, which can climb shelves and manipulate objects; the foot-bots, which can navigate on the ground and connect to objects or other robots.

The first publication presents the main results of the project:

- M. Dorigo, D. Floreano, L.M. Gambardella, F. Mondada, S. Nolfi, T. Baaboura, M. Birattari, M. Bonani, M. Brambilla, A. Brutschy, D. Burnier, A. Campo, A.L. Christensen, A. Decugnière, G. Di Caro, F. Ducatelle, **E. Ferrante**, A. Förster, J. Martinez Gonzalez, J. Guzzi, V. Longchamp, S. Magnenat, N. Mathews, M. Montes de Oca, R. O'Grady, C. Pinciroli, G. Pini, P. Réturnaz, J. Roberts, V. Sperati, T. Stirling, A. Stranieri, T. Stützle, V. Trianni, E. Tuci, A.E. Turgut, F. Vaussard. Swarmanoid: a novel concept for the study of heterogeneous robotic swarms. IEEE Robotics & Automation Magazine, 20(4), 2013. In press.

¹Swarmanoid project, <http://www.swarmanoid.org>, February 2013

The Swarmanoid project was successful and ended with an integrated demonstrator showing cooperative foraging in a three-dimensional complex environment². A video of the demonstrator has been published in:

- M. Dorigo, D. Floreano, L.M. Gambardella, F. Mondada, S. Nolfi, T. Baaboura, M. Birattari, M. Bonani, M. Brambilla, A. Brutschy, D. Burnier, A. Campo, A. L. Christensen, A. Decugnière, G. A. Di Caro, F. Ducatelle, **E. Ferrante**, A. Förster, J. Martinez Gonzales, J. Guzzi, V. Longchamp, S. Magnenat, N. Mathews, M. A. Montes de Oca, R. O' Grady, C. Pinciroli, G. Pini, P. Rétornaz, J. Roberts, V. Sperati, T. Stirling, A. Stranieri, T. Stützle, V. Trianni, E. Tuci, A. E. Turgut and F. Vaussard. Swarmanoid, the movie. In AAAI-11 Video Proceedings. AAAI Press, San Francisco, CA, 2011. Winner of the "AAAI-2011 Best AI Video Award".

For the project, we contributed to the development of a simulation tool called ARGoS. The ARGoS simulator, which we describe in Section A.2, has been used for executing most of the experiments presented in this dissertation. ARGoS is described in the following two publications:

- C. Pinciroli, V. Trianni, R. O'Grady, G. Pini, A. Brutschy, M. Brambilla, N. Mathews, **E. Ferrante**, G. Di Caro, F. Ducatelle, T. Stirling, Á. Gutiérrez, L. M. Gambardella, M. Dorigo. ARGoS: a Modular, Multi-Engine Simulator for Heterogeneous Swarm Robotics. In Proceedings of the 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS' 11), pages 5027-5034. IEEE Computer Society Press, Los Alamitos, CA, 2011.
- C. Pinciroli, V. Trianni, R. O'Grady, G. Pini, A. Brutschy, M. Brambilla, N. Mathews, **E. Ferrante**, G. Di Caro, F. Ducatelle, M. Birattari, L. M. Gambardella, M. Dorigo. ARGoS: a Modular, Parallel, Multi-Engine Simulator for Multi-Robot Systems. *Swarm Intelligence*, 6(4):271-295, 2012.

1.3 Dissertation outline

The rest of the dissertation is organized as follows.

Chapter 2 introduces the context of our research. It gives an overview of what swarm intelligence and swarm robotics are, presents the biological background on information transfer and reviews the flocking literature from different points of view.

²Swarmanoid: the movie, <http://www.youtube.com/watch?v=M2nn1X9X1ps>, January 2013

Chapter 3 describes the main methods used for information transfer in flocking. It describes the principles and methods used to achieve cohesion, the communication strategies used for explicit information transfer, and the implicit information transfer mechanisms.

Chapter 4 presents the results obtained with explicit information transfer. In all the experiments presented in the chapter, we assume that only few robots are aware of the desired goal direction to be followed. We present results of experiments with stationary and non-stationary goal directions, with one and two goal directions.

Chapter 5 presents an application of the principles behind explicit information transfer to another collective behavior: collective transport. We show how a mechanism similar to the one used in Chapter 4 can be used to achieve collective transport of an object to a goal location in an environment containing obstacles to be avoided.

Chapter 6 presents the results obtained in experiments with flocking with implicit information transfer. The chapter first presents experiments with hybrid explicit and implicit information transfer. It then focuses on implicit information transfer.

Chapter 7 analyzes implicit information transfer from another perspective than swarm robotics: statistical physics. We use numerical and analytical tools to further analyze the properties of flocking with implicit information transfer and to better understand the driving forces underlying the mechanism.

Chapter 8 concludes the dissertation with a summary of the main contributions and with several proposals for future research directions.

Finally, Appendix A describes the simulated and real-robot platforms used to carry out the experiments presented in this dissertation that are related to swarm robotics.

Chapter 2

Background and related work

In this chapter, we introduce the context in which the research described in this dissertation has been carried out.

In Section 2.1, we briefly introduce swarm intelligence. We define the research scope of swarm intelligence and the properties of swarm intelligence systems.

In Section 2.2, we introduce the main context of this research: swarm robotics. In doing so, we briefly review the collective behaviors that have been studied in the literature and the methods used to design and analyze such systems. This review is not intended to be exhaustive: the interested reader can refer to [Brambilla et al. \(2013\)](#) for a more complete review of swarm robotics.

In Section 2.3, we introduce the topic of information transfer putting particular emphasis on examples in biological systems. As with the chapters in this dissertation, we distinguish between signal based or explicit and cue-based or implicit information transfer.

Finally, in Section 2.4, we review the literature related to self-organized flocking, the main collective behavior studied in this dissertation. Self-organized flocking has been object of study for many disciplines since three decades. Mainly, flocking has been studied in biology, statistical physics, control theory, and robotics. In Section 2.4, we review representative contributions within these four disciplines.

2.1 Swarm intelligence

Swarm intelligence studies self-organizing, decentralized collective behaviors exhibited by systems composed of many locally-interacting individuals. Swarm intelligence focuses on both naturally evolved (Section 2.1.1) and artificially engineered (Section 2.1.2) systems. Both types of systems are characterized by several interesting and useful properties (Section 2.1.3).

2.1.1 Naturally evolved swarm intelligence systems

The most striking examples of swarm intelligence systems are represented by group of animals exhibiting complex behaviors (Camazine et al., 2001) (Figure 2.1). Social insects often manifest complex and astonishing collective behaviors. For example, many species of ants (Figure 2.1a) are organized in very complex societies. Parts of their colonies are allocated to environment exploration: they look for resources and collectively select the best or the closest resource to their nest (Beckers et al., 1990). Other parts of the colony are instead capable of collectively building very complex nests with tunnels and chambers for processing resources, growing food, and dumping garbage. These tasks are performed without any centralized decision-making mechanism. All the decisions are made in a purely distributed fashion, with individuals sensing each other via odor or interacting locally and indirectly via pheromone, a chemical compound that ants deposit in the environment. Hence, all individuals act locally with no clear picture of what is going on at the level of the colony. Other examples of complex behaviors in social insects are exhibited by honey bees (Figure 2.1b), which are for instance capable of regulating very precisely the temperature of their hive by modifying their own individual temperature (Stabentheiner et al., 2010).

Between humans and social insects we find many other species that exhibit another collective behavior: coordinated motion (or, as we will refer to in this dissertation, flocking). Examples of coordinated motion can be observed in birds (such as starlings - Figure 2.1c) and in fish (Figure 2.1d). These animals can aggregate and move together in very large groups that can expand, contract, change shape, escape from predators, all as if they were a single super organism (Krause and Ruxton, 2002). In these groups, there is no external guidance nor leader, although some individuals might possess better information and have developed a way to directly and indirectly transfer this information to the rest of the group (Couzin et al., 2005). Coordinated motion is also observed in other animals. Examples are: locusts that, by means of cannibalistic instincts, developed coordinated motion based on escape and pursuit behaviors (Bazazi et al., 2008); mammals such as sheep able to pursue coordinated motion initiated by few individuals trained to recognize food (Pillot et al., 2011); and human pedestrian crowds that have shown to exhibit global patterns such as lane formations without centralized coordination or explicit negotiation (Moussaïd et al., 2011). Coordinated motion is the main collective behavior studied in this dissertation.

From the biological perspective, the evolution of collective behaviors and cooperation in social insects is studied from the perspective of inclusive fitness (Hamilton,




Figure 2.1: Examples of naturally evolved swarm intelligence systems: (a) a swarm of army ants; (b) a bee hive; (c) a starlings flock; (d) a fish shoal. Sources: (a) CC BY 2.0 - Geoff Gallice - http://commons.wikimedia.org/wiki/File:Army_ants.jpg (b) CC BY 2.0 - Chas Redmond - http://commons.wikimedia.org/wiki/File:Bees_cluster.jpg (c) CC BY 2.0 - Jeremy Bolwell - <http://www.geograph.org.uk/photo/3347688> (d) CC BY 2.0 - Matt Kieffer - <http://www.flickr.com/photos/mattkieffer/5423074048/>.

1964). According to this theory, the evolution of cooperation and swarm behaviors is promoted in these biological systems by having all individuals in a colony related to each other. This is because cooperation in spite of selfishness can be evolutionary favoured when individuals share the same genes. Thus, the emergence of self-organizing behaviors is perhaps not surprising in social insects when seen from the inclusive fitness perspective. However, inclusive fitness theory cannot be applied directly to explain biological systems such as those that inspired this dissertation: bird flocks and fish schools. In these systems, groups are composed of non relatives, which normally causes individuals to act out of self interest only. Thus, any group-level behaviors is a side result of individual and selfish behaviors. One of the main

theory that explains the evolution of coordinated motion-type behaviors is known as the *selfish herd effect* (Hamilton, 1971), which states that individuals act as to have other individuals between them and some predator, thereby resulting in aggregation. Another theory able to explain flocking is the so called *information centre hypothesis* (Ward and Zahavi, 1973), which states that some animals benefit from being and moving together in order to follow other individuals that are foraging, thereby maximizing their success to find profitable food patches. Although we do not advocate specifically for this last theory, we do believe that the capability to transfer and exploit information in a swarm is of a key importance for the realization of collective motion and self-organization in general in artificial swarms.

2.1.2 Artificially engineered swarm intelligence systems

Inspired by the natural systems described above, artificially engineered swarm intelligence systems have been developed to solve complex real-world problems. Examples of such systems are ant colony optimization, or ACO (Di Caro and Dorigo, 1998) and particle swarm optimization, or PSO (Kennedy and Eberhart, 1995). ACO and PSO algorithms are inspired by two naturally evolved swarm intelligence systems: ant colonies and bird flocks. These algorithms are used to tackle difficult optimization problems by exploiting the idea of having many software agents searching for a solution in parallel: In ACO, good solutions are promoted by accumulation of artificial pheromone; in PSO, virtual individuals flock together in the abstract solution space influenced by random fluctuations and the good solutions found so far.

We give here two examples of problems solved by artificially engineered swarm intelligence systems: *routing* and *clustering*. Routing, is a distributed problem that is intrinsic in telecommunication networks. In routing, each node must decide to which other node each packet should be sent. Ant Net (Di Caro and Dorigo, 1998) is a distributed algorithm inspired by the behavior of social ants which was shown to have superior results compared to other algorithms. Clustering is a problem related to data analysis which consists in grouping data in order to have similar-to-each-other items together in the same cluster and different-from-each-other items in different clusters. Ant-based clustering algorithms (Handl and Meyer, 2007), also inspired by the behavior of social ants, tackle the problem as follows: virtual ants move randomly in the abstract space of data items, they pickup data items when encountered, and probabilistically drop them close to other data items, with a dropping probability that increases with increasing similarity between the carried items and the encountered one.

Other artificially engineered swarm intelligence systems, more important for the

context of this dissertation, are represented by swarm robotics systems, presented in more details in Section 2.2.

It is important to note that, differently from naturally evolved swarm, which could evolve to maximize either individual or group performance (see Section 2.1.1), artificial swarms are always engineered with the objective to maximize the group performance in the task at hand.

2.1.3 Characteristics and properties of swarm intelligence systems

Swarm intelligence systems are driven by self-organization, which refers to a process through which a system composed of many components reaches a coherent and interpretable macroscopic spatio-temporal pattern or state. This pattern is reached spontaneously, without the need of a leader or external influence, and without requiring the individual components to have information about the macroscopic pattern itself. Self-organization is believed to be induced by two mechanisms: *positive feedback* and *negative feedback* (Garnier et al., 2007). Positive feedback is the mechanism whereby microscopic *fluctuations*, for example due to the stochastic dynamics of the components, are amplified within the system. Positive feedback can help the system to reach one of its possible stable macroscopic states, but can potentially destabilize the system as well. Negative feedback acts as a counterbalancing mechanism for positive feedback, by dampening some fluctuations and biasing the system towards stability.

Swarm intelligence systems have intrinsic properties that make them unique and often preferable to monolithic and centralized systems.

Scalability refers to the ability to perform well for increasing swarm sizes. In swarm intelligence systems, scalability is promoted by the exclusive use of local interactions among the components of the system. This makes the number of interactions per individual constant independently of the swarm size and, consequently, the total number of interactions scales linearly with respect to the swarm size.

Parallelism refers to the potential of a swarm to perform multiple tasks in parallel using multiple individuals at the same time, improving the overall performance (Clark and Mangel, 1986). Examples can be found in nature: for instance, colonies of leaf-cutting ants (Anderson and Ratnieks, 2000) are shown to perform different activities in parallel, such as cutting leaves, collecting leaf pieces, storing them, etc. . .

Robustness refers to the capability of the system to perform well even in presence of disturbances in the environment. This is achieved through several factors, such as: redundancy, that is, the presence of many individuals that can provide the same functionality; multiplicity of sensing, that increase the overall signal-to-noise ratio; and decentralization, that is, the absence of a leader and of a central decision point.

Flexibility refers to the capability of the system to change in response to changes in the environment. Once more we find examples in social ants: ant colonies, after having found the shortest path to a food source, are shown to adapt and choose another path when the previous one is no longer the shortest ([Bert Hölldobler and Wilson, 1990](#)).

To obtain these properties in artificial systems, particular care should be given to the design of the interactions among the system components. Such a design is not simple, as we advocate in [Brambilla et al. \(2013\)](#), and sometimes it is delegated to automated methods ([Trianni and Nolfi, 2011](#)), as we briefly discuss in Section 2.2.2.

2.2 Swarm robotics

Swarm robotics has been defined in [Şahin \(2005\)](#) as:

“the study of how a large number of relatively simple physically embodied agents can be designed such that a desired collective behavior emerges from the local interactions among agents and between the agents and the environment.”

The main characteristics of swarm robotics systems are the following:

- they consist of a large number of autonomous robots;
- sensing and communication among the robots is purely local;
- robots do not rely on global information and centralized control;
- the desired collective behavior is obtained through self-organization, robot-to-robot and robot-environment interactions.

In [Brambilla et al. \(2013\)](#), we compiled a survey in which we classify the literature according to what collective behavior is studied and according to what design or analysis method is used. In the following, we use the same axes and provide a very

concise summary of the swarm robotics literature. In Section 2.2.1, we briefly review some collective behaviors studied in swarm robotics. In Section 2.2.2 and Section 2.2.3, we briefly review the main design and modeling methods used in swarm robotics, respectively. The main goal of this overview is to frame the work presented in this dissertation in the broader context of swarm robotics.

2.2.1 Collective behaviors

Swarm robotics studies many collective behaviors that, alone or combined, can be used to tackle complex problems. In Brambilla et al. (2013), we divide collective behaviors in three main categories: i) spatially-organizing behaviors, which organize and distribute robots and objects in space; ii) navigation behaviors, able to organize and coordinate the movement of the swarm; iii) and collective decision-making and task allocation, able to make the swarm agree on a common decision or divide the swarm into two or more groups based on their decision. In this section, we only report few representative examples drawn from each category: aggregation, pattern formation, collective exploration, collective decision-making/task allocation, and coordinated motion¹.

Aggregation is a spatially-organizing behavior used to group all the robots of a swarm at the same location in the environment. Aggregation is perhaps the simplest collective behavior, but it is hard to obtain only through swarm intelligence principles. Aggregation is a necessary precondition for other collective behaviors such as pattern formation and flocking, which require robot meeting at the same location. Authors studying aggregation take inspiration from the behavior observed in bacteria, cockroaches, bees, fish, and penguins (Camazine et al., 2001).

The most challenging type of aggregation problem in swarm robotics is the one in which robots should aggregate at a random location, that is, no information, cues or landmarks are available to decide where to aggregate. In this case, methods such as probabilistic finite-state machines (Minsky, 1967) or artificial evolution (Nolfi and Floreano, 2004) are used to perform aggregation. These methods are both explained in Section 2.2.2. Other works also deal with the case in which cues (such as shelters (Campo et al., 2011) or gradients (Bodi et al., 2012)) are present in the environment. In this case, authors used mainly probabilistic finite-state machines, and took inspiration from the behavioral models

¹The other collective behaviors analyzed in Brambilla et al. (2013) are: chain formation, self-assembly and morphogenesis, object clustering and assembling, collective transport, collective fault detection, group size regulation, and human-swarm interaction.

of animals such as cockroaches and bees.

Pattern formation is another spatially-organizing behavior used to organize robots in a regular and repetitive manner. Pattern formation is inspired by biological processes such as the formation of chromatic patterns in animals (Meinhardt, 1982) or physical processes such as crystal formation (Langer, 1980). The most common approach to pattern formation is the virtual physics-based design, explained in Section 2.2.2. Pattern formation is a necessary component of coordinated motion, the main collective behavior considered in this dissertation.

Collective exploration is a navigation behavior used by robots to explore an environment, find resources, and navigate efficiently. A source of inspiration for this behavior are social insects such as ants (Camazine et al., 2001). Typically, exploration of the environment is realized through area coverage, whereby robots form an interconnected static or dynamic network covering areas of interest in an environment. Area coverage is in turn realized by means of virtual physics-based design (Howard et al., 2002).

Navigation in the environment is instead a collective behavior executed after area coverage and building on its outcome. Navigation is realized by means of communication techniques and finite-state machines (Ducatelle et al., 2011).

Collective decision making and task allocation refers to the capability of robots to influence each other when making choices. Two common possible outcomes of this process are either consensus to the same decision or specialization to two or more possible alternatives. Examples of consensus decision-making in nature can be found in many animals, although the typical source of inspiration for swarm robotics are cockroaches (Amé et al., 2006). Examples of task allocation are instead more often found in animals organized in castes, such as ants and bees (Camazine et al., 2001).

Diverse methods are used for consensus decision-making (Campo, 2011), such as probabilistic finite-state machines (Garnier et al., 2005) and methods based on statistical-physics (Montes de Oca et al., 2011). The design methods instead used for task allocation are mostly probabilistic finite-state machines (Liu et al., 2007; Brutschy et al., 2012), pioneered by the seminal work by Bonabeau et al. (1997).

Coordinated motion or flocking is the main collective-behavior studied in this dissertation. Therefore, we review the flocking literature more in details than for other behaviors in a dedicated section (Section 2.4). The design method to achieve flocking is also described in details in Chapter 3.

2.2.2 Design methods

Swarm robotics still lacks rigorous design methods. We believe this is due to the intrinsic nature of swarm robotics systems: given a desired macroscopic objective, the derivation of the microscopic behaviors and interaction rules is not obvious. We summarize here the main design methods for swarm robotics and, as done in [Brambilla et al. \(2013\)](#), we classify them in two categories: hand-coded design methods and automatic design methods.

Hand-coded design methods use an iterative process in which the microscopic behavior is implemented and improved until the desired macroscopic objective is achieved.

One of the existing hand-coded design paradigms consists in using *finite-state machines* ([Minsky, 1967](#)). Often, these finite-state machines are probabilistic (PFMSs). In this case, the state transitions are governed by probabilities, as it is the case also for the behavioral models of social insects such as ants ([Bonabeau et al., 1997](#)) or cockroaches ([Amé et al., 2006](#)). In few other cases, such as morphogenesis ([O’Grady et al., 2009a](#)), finite-state machines are deterministic, that is, state transitions are not governed by probabilities.

The other existing hand-coded design method is based on *virtual physics*. In virtual physics-based design, each robot is considered to be a virtual particle subject to virtual forces exerted by other robots or by the environment. Virtual forces can be used, for example, to guide robots towards a target, to make them avoid obstacles ([Khatib, 1986](#)) or other robots ([Reif and Wang, 1999](#)), or to organize them into patterns or formations ([Spears et al., 2004](#)). The design methodology described in Chapter 3 belongs, in part, to this category. Other recently-proposed hand-coded design methodologies are amorphous computing ([Bachrach et al., 2010](#)) and property-based design ([Brambilla et al., 2012](#)).

Automatic design methods are able to automatically produce the microscopic behaviors and interactions through the optimization of a function that evaluates the performance at the swarm level. The two main design paradigms that have been used in swarm robotics are reinforcement learning ([Kaelbling et al., 1996](#)) and evolutionary robotics ([Nolfi and Floreano, 2004](#)).

In reinforcement learning, an agent learns a behavior through trial-and-error interactions with the environment and by receiving positive and negative feedback from the environment. This feedback is typically a measure of the performance of the individual behavior of the robot. Unfortunately, in reinforcement learning it is hard to evaluate directly the swarm-level behavior. This and other

important challenges (thoroughly analyzed in [Brambilla et al. \(2013\)](#)) make reinforcement learning not directly applicable to swarm robotics, and only limited success has been reported ([Matarić, 1998](#); [Panait and Luke, 2005](#)).

In evolutionary robotics, methods borrowed from evolutionary computation ([Goldberg, 1989](#)) are applied to robotics. In evolutionary robotics, evolutionary computation algorithms are used to maximize a fitness function, which can either be the individual or swarm performance ([Waibel et al., 2009](#)), although typically the swarm performance is chosen. For this reason, evolutionary robotics can be applied directly to swarm robotics ([Trianni and Nolfi, 2011](#)), and more evidence of success has been reported ([Baldassarre et al., 2003](#); [Ampatzis et al., 2008](#); [Sperati et al., 2011](#)). A drawback of evolutionary robotics is that it typically uses artificial neural networks ([Beer and Gallagher, 1992](#)) as the representation for individual robot's behaviors. Despite being very general, artificial neural networks are very difficult to reverse engineer. Thus, it is difficult to understand and to extract the principles responsible for the self-organized behavior once this has been evolved.

2.2.3 Modeling methods

In swarm robotics, modeling is very useful to predict whether given properties of the designed collective behavior hold or not. Unfortunately, a unique theoretical framework for modeling swarm robotics systems does not exist yet, but studies borrow diverse modeling techniques from several other disciplines. In swarm robotics, modeling can be performed at two different levels: the microscopic, that focuses on the individuals and on the interactions among them, and the macroscopic level, that focuses on the swarm as a whole. On the one hand, at the microscopic level, modeling is performed using numerical simulations, since mathematical methods can hardly be applied to systems composed of more than three or four components. On the other hand, at the macroscopic level, some mathematical methods do exist. We distinguish the latter into deterministic and stochastic modeling.

Deterministic modeling is usually based on control theory, differential equations or model checking. Control theory modeling typically assumes that robots interact with each other without noise and are connected through a static or a dynamic graph topology. Using control theory, it is possible to prove whether the swarm will eventually (i.e., asymptotically) reach the desired macroscopic state. Control theory has been used, for example, to prove convergence in aggregation ([Gazi and Passino, 2004](#)), foraging ([Liu and Passino, 2004](#)), and task

allocation (Hsieh et al., 2008). In a few cases, control theory has also been used to derive the optimal control rules to achieve the desired properties (Hsieh et al., 2008). Control theory models are based on some assumptions which are unfortunately violated in swarm robotics systems. Other deterministic models include ordinary differential and difference equations (Montes de Oca et al., 2011; O'Grady et al., 2009b), partial differential equations (Berman et al., 2009), and model checking (Winfield et al., 2005).

Stochastic modeling is usually based on stochastic difference or differential equations or on probabilistic model checking. One of the first stochastic model ever proposed in swarm robotics is based on rate and master equations (Martinoli et al., 1999). Rate and master equations can be used to define a macroscopic model given an initial microscopic model expressed as a PFSM. Rate equations have been used to model many behaviors such as clustering (Martinoli et al., 1999), foraging (Lerman and Galstyan, 2002), and stick-pulling (Martinoli et al., 2004). The main limitation of rate equations is their incapability to explicitly model space and time.

A modeling framework that explicitly models space and time has been introduced in Hamann and Wörn (2008) and is based on the Fokker-Planck equation. The Fokker-Planck equation is a stochastic differential equation used to model the time evolution of a continuous probability distribution, such as the distribution of robots in the physical space. The model has been used to study emergent taxis (Hamann and Wörn, 2008), aggregation (Schmickl et al., 2009), and area coverage (Prorok et al., 2011). More recently, probabilistic model checking has been used to prove properties of swarm robotics systems (Brambilla et al., 2012; Massink et al., 2012).

2.3 Information transfer in biological systems

Information transfer is a concept that unifies our understanding of collective behaviors of social animals such as flocking (Sumpter et al., 2008). In nature, evolution applies selection pressure for effective information transfer mechanisms, since those lead to better environment exploration strategies: individuals have more chances of finding resources if other individuals in the group share this information with them (Sumpter, 2010).

In this section, we briefly describe some of the information transfer mechanisms that have evolved in social animals. Inspired by Sumpter (2010), we divide these mechanisms in two categories: signal-based and cue-based information transfer.

This taxonomy corresponds to the main taxonomy used in this dissertation: explicit versus implicit information transfer.

2.3.1 Signal-based information transfer

Signal-based information transfer mechanisms are often observed in social insects as well as in other animals.

In ant colonies, *pheromone* is an effective signaling mechanism used to mark the route from a food source to the nest (Bert Hölldobler and Wilson, 1990). After having found the food source, few ants informed about the food source return to the nest while leaving pheromone on their way back. The pheromone acts effectively as a signal, as it is detected and followed by other ants that were not informed about the food source. Those ants will also drop pheromone, thus reinforcing the signal (positive feedback).

Another signal-based information transfer mechanism is used by honeybees and is called *waggle dance* (Seeley, 2010). With waggle dance, honey bees can make different types of collective decisions, such as selecting one or more flower for foraging or selecting one nest location to relocate. In the first case, to transfer information to other individuals, few informed individuals called scout perform the waggle dance, which is a figure eight pattern whose direction and duration are correlated to the direction and the distance from the bee hive to the nectar, respectively. Non-informed bees will follow the instructions encoded in the dance and, in some cases, also find food using odor and visual cues (Riley et al., 2005). After the food source is found, previously non-informed individuals will also start performing the waggle dance, thus reinforcing the signal. In case the swarm needs to decide a nest site to relocate, it is very important that consensus is achieved to only one location. This is ensured by two mechanisms: first, the waggle dance also encodes the quality of the nest site, thus higher-quality locations will be advertised faster than lower-quality locations; second, a “piping signal” is used to inform bees advertising lower-quality location that they should cease doing so.

Explicit signaling can be found also in other species. For example, mountain gorillas use vocalization to coordinate switching between daily activities such as resting, traveling or feeding (Stewart and Harcourt, 1994). Another example is represented by naked mole rats, who leave odor trails and make chirping noises upon finding food and returning to the nest (Judd and Sherman, 1996).

In all these examples, the information transfer mechanism is based on explicit signals. In biology, a signal is defined as “an act or structure that alters the behavior of another organism, which evolved because of that effect and which is effective

because the receiver's response has also evolved" (Smith and Harper, 2003). One of the main reasons why it is important to draw a distinction between signal-based and cue-based information transfer is that signals are associated to a cost, that can be the cost to develop the signaling organ or the energy spent in emitting the signal (Guilford and Dawkins, 1991). This distinction is true also for the two paradigms used in this dissertation: by having robots using explicit information transfer, we also incur in a cost, which can be a manufacturing cost for producing the necessary hardware or an energy cost for maintaining communication.

2.3.2 Cue-based information transfer

Signal transmission is not the only mechanism whereby information can be transferred in animals. Some other animals exploit cues exhibited "unintentionally" by the informed individuals while these are executing their normal activities.

One example of such mechanism is found in bats. Some uninformed bats are known to ascertain the existence of a novel type of food by just looking at or smelling another individual selfishly eating that food (Ratcliffe and ter Hofstede, 2005). In this case, only the presence of food is transferred, and not the position. In other species, such as hooded crows (Sonerud et al., 2001), also the location is implicitly transferred by having non-informed individuals following the informed ones. Another example include starlings, which use information about food owned by their flock-mates to decide when to start searching for more food (Templeton and Giraldeau, 1996).

Differently from signal-based information transfer, each of these mechanisms evolved just because of a byproduct correlation between the mechanism itself and the act of foraging for a resource. Differently from signal-based information transfer, in cue-based information transfer there is no direct cost associated to the behavior responsible for transferring the information. As already said, this difference is also presented in the categorization used in this dissertation: In implicit information transfer, robots will not use any signaling or communication to share directional information, thus incurring in no energy costs associated to communication and in no additional cost in producing specialized hardware.

2.4 Flocking

Flocking is a phenomenon that can be widely observed in nature. In flocking, a groups of tens up to several thousands animals such as fish or birds (Couzin et al., 2002) move and maneuver together as if they were a single creature.

Flocking did not receive much attention from fields outside biology until the mid '80s, when Reynolds (1987) published a computer model of flocking that is considered seminal. Reynolds was one of the first to propose simple behaviors based on local sensing rules that, combined, could realize flocking with artificial agents. Although his main goal was only to generate realistic computer animation of flocking, later the model turned out to be more general as it became used in domains other than computer graphics. Reynolds work could be the reason why *flocking* is often used in place of the more general *collective motion* term: his work was inspired from the flocking behavior of birds, and many works afterwards re-used the term “flocking”.

In Reynolds (1987), a *flock* was formally defined as:

“... a group of objects that exhibits a general class of polarized, non-colliding, aggregate motion.”

His flocking model was based on local interactions among individuals rather than on a centralized controller that calculates all the trajectories of all individuals. He proposed the following three simple behaviors:

- *Separation*: Each individual avoids collisions with its neighbors, as shown in Figure 2.2a.
- *Alignment*: Each individual matches its velocity to the average of its neighbors, as shown in Figure 2.2b. Velocity was considered as a vectorial quantity, composed of a direction and a magnitude.
- *Cohesion*: Each individual moves towards the center of mass of the relative positions of its neighbors, as shown in Figure 2.2c.

Reynolds' algorithm is very general and can be easily applied to design flocking behaviors in many domains. However, there is a major implementation concern related to how to achieve sensing in physical world applications. The algorithm, in fact, assumes that each individual has the knowledge of the velocity of its neighbors, and that this information is free of noise. These two assumptions are quite unrealistic but they have been relaxed in subsequent work: in statistical physics, where noise was introduced and fully investigated (Vicsek et al., 1995); in robotics, where Turgut et al. (2008a) showed that only the orientation of the neighbors, that is, only the direction of the velocity vector, is needed to achieve flocking (Turgut et al., 2008a).

In the rest of this section, we review the literature on flocking. Flocking is a multi-disciplinary topic, and has been subject of research in many areas. We review




Figure 2.2: The three behaviors of Reynolds' flocking model. (a) Separation behavior. (b) Alignment behavior. (c) Cohesion behavior. Arrows indicate the steering direction of the individual due to each behavior. Reproduced from Reynolds (2008)

the literature related to flocking as studied by these four research areas: biology (Section 2.4.1), statistical physics (Section 2.4.2), control theory (Section 2.4.3) and robotics (Section 2.4.4).

2.4.1 Flocking in biology

Within biology, we identify two categories of studies on flocking: in the first, we place work that focused on the ultimate evolutionary factors promoting flocking; in the second, we place work that focused on the proximate mechanisms underlying flocking.

Ultimate factors that favor flocking

Many studies have confirmed the importance of flocking for the survival of many species. According to these studies, increased survival can be attributed to different reasons: First, the probability of being killed or eaten by a predator when part of a large group is lower than the probability of that event when being alone, since the individuals at risk are mainly those at the periphery of the group (Partridge, 1982); second, having a large group size can confuse predators sensing capabilities (Moyle and Cech, 2004); third, some studies have suggested that having many individuals searching for food is superior to having a single individual due to increased sensing capabilities and information sharing (Pitcher et al., 1982).

Apart from increasing the survivability, other motivations for flocking have also been proposed. For instance, flocking enables a more energy efficient movement of flying animals. Birds exploit the streamlines induced by other individuals in the front for increasing their thrust and decreasing the energy consumption necessary to fly long distances (Ballerini et al., 2008).

[Simpson et al. \(2006\)](#) analyzed the factors at the basis of the marching behavior of crickets. The authors found that cannibalism is the main drive of flocking in this case: Crickets try to move towards other crickets to eat them, while running away from others to avoid being eaten. The collective motion of crickets has been analyzed also in statistical physics by [Romanczuk et al. \(2009\)](#).

Proximate mechanisms that favor flocking

Studies have identified several individual-level mechanisms responsible for flocking. A microscopic model of the behavior of army ants was developed by [Couzin and Franks \(2002\)](#). The model included obstacle-avoidance and pheromone-tracking components. From the results, the authors concluded that ants tend to move in a common direction to avoid collisions with other ants while roaming in the environment.

[Couzin et al. \(2002\)](#) proposed the zone model to closely model the behavior of fish shoals and bird flocks. The model is similar to Reynolds', but the individuals interact with each of their neighbors using only one of the three rules at a time: depending on its distance from the focal individual, a neighbor can be in the zone of repulsion (closest), zone of orientation (intermediate) or zone of attraction (farthest), determining which rule to use.

In a subsequent study, which we consider very important for this dissertation, [Couzin et al. \(2005\)](#) studied the effect of introducing informed individuals, aware of a certain *goal direction* to be followed, in a larger group where other members are not informed about that direction. He found that few informed individuals are enough to lead the group, and that other individuals do not need to know who the informed individuals are. Furthermore, the proportion of informed individuals needed to lead the group decreases as a function of the group size. The framework of informed and non-informed individuals has also been recently studied mathematically by [Yu et al. \(2010\)](#).

Other determining factors to the emergence of flocking behavior are studied by [Buhl et al. \(2006\)](#). The authors suggested that desert locusts form clusters and move in ordered fashion when above a critical density. The authors also modeled the system with a one-dimensional model, showing that the predictions of the model are in agreement with the experiments ([Czirok et al., 1999](#); [Czirok and Vicsek, 2000](#)).

[Ballerini et al. \(2008\)](#) studied the nature of the interactions in a flock of starlings. By equipping birds with GPS trackers, the authors showed that starlings interact with a fixed number of neighbors which are not determined by their metric distance but by their topological distance. These results were checked against a model de-

rived from the one by [Vicsek et al. \(1995\)](#). Using the model, the authors determined that individuals interacting with topologically-close neighbors can more easily separate and rejoin under the attack of a predator compared to individuals interacting with metric-distance neighbors.

More recently, [Cavagna et al. \(2010\)](#) found that correlations between the velocities of the starlings have a scale-free nature. This suggests that, although individuals interact only locally, cascading effect can still create long-range correlations which make the flock very responsive to environmental perturbations such as the presence of predators.

2.4.2 Flocking in statistical physics

In statistical physics, flocking is studied in the context of searching for a theory of self-propelled collective motion applicable to a wide range of systems. Often, the approach is to consider simple models, such as point-particle systems where actuation or sensing noise plays a role similar to that of temperature in standard physical systems. The resulting dynamics are analyzed using tools from nonequilibrium statistical physics.

We describe briefly below some of the most representative flocking studies carried out from the statistical physics perspective. We classify them into two groups: with and without explicit alignment rules. This classification is analogous to the one used in this dissertation: models with explicit alignment rules use explicit information transfer, whereas models without alignment rules use implicit information transfer. Our goal here is only to contextualize our dissertation, and not to provide an exhaustive literature review of studies of flocking in statistical physics. A more comprehensive survey can be found in [Vicsek and Zafeiris \(2012\)](#).

Statistical physics models with explicit alignment rule

The first model developed to study flocking from a statistical physics perspective, the self-driven particle (SDP) model, was proposed by [Vicsek et al. \(1995\)](#). The SDP model considers point-particles that tend to align with each other while advancing at constant speed. At every time-step, each particle's orientation is set to the average orientation of its neighbors plus a random noise. The model includes no attraction nor repulsion. The authors found that a phase transition separates the ordered state (particles moving in a common direction) from the disordered state (particles with random orientation). The transition is reached by either varying the amount of noise or the particle density.

Grégoire et al. (2003) extended the SDP model by adding attraction and repulsion interactions. The orientation of each particle is computed as a weighted sum of an alignment term and an attraction/repulsion term. Noise is introduced by adding a small randomly oriented vector of length proportional to the noise level (an implementation often called *vectorial noise*) to the orientation. As in the Vicsek et al. (1995) case, the authors also observed here a transition between an ordered and a disordered state at a critical noise value. The details of this transition, however, appear to be different between Vicsek et al. (1995) and Grégoire et al. (2003), and this is still today a matter of debate.

Aldana and Huepe (2003) considered a different type of model, the vectorial network model (VNM), in which individuals are not represented by particles moving in space but by nodes of a network. Each node has an associated orientation that evolves following an SDP algorithm, and neighbors are defined as nodes directly connected through a link. The authors found that an order-disorder transition equivalent to that in the SDP model can be obtained, but only if links corresponding to long-range interactions are present. The VNM was extended by Turgut et al. (2008b) to analyze the self-organized flocking of robots that follow the algorithm proposed in Turgut et al. (2008a).

Finally, Cucker and Huepe (2008) studied theoretically the behavior of the swarm when some informed individuals are introduced. The authors considered a Laplacian model, where each particle tries to minimize the difference between its own velocity and that of its neighbors. They define the neighbors according to an arbitrary adjacency matrix that must satisfy certain properties. All particles follow the same basic interaction rules, but informed individuals then add a preferred heading direction. This model was originally developed to study a simple case of collective decision-making dynamics.

Statistical physics models without explicit alignment rule

Szabó et al. (2006) were the first to propose a model that displays ordered motion using locally interacting particles that do not follow explicit alignment rules. The goal was to model the dynamics of collectively migrating tissue cells, for which they had gathered experimental data. The authors found that the usual order-disorder phase transition is also observed here when the mean density of particles becomes smaller than a critical value. Their model has many similarities with the one we propose in Chapter 3 and whose results are presented in Chapter 6. For example, in both cases, particles interact through attraction-repulsion forces and move forward with a non-constant speed. An important difference, however, is that in their model each parti-

cle is not constrained to only move parallel to its orientation (forward or backward) but can also advance by sliding sideways. This difference makes their algorithm not directly applicable to standard non-holonomic mobile robots, which are constrained to move along the direction in which their wheels are pointing. In [Henkes et al. \(2011\)](#), the same model was studied in the context of active jamming, that is, they studied the behavior of particles at very high densities.

[Grossman et al. \(2008\)](#) considered a system composed of particles that interact only through inelastic collisions and move on a 2D plane with reflecting circular boundary conditions. They observed complex phenomena such as ordered motion, vortexes and chaos. They then studied the system in elliptic-shaped arenas, showing that more complex dynamical patterns, such as particles moving together in sub-groups, could be observed.

In [Romanczuk et al. \(2009\)](#), the authors considered particles with asymmetric non-isotropic interactions. Their model represents “escape-pursuit” dynamics. A particle is attracted to other particles in front of it that are moving away (pursuit), but is repelled from other particles approaching it from behind (escape). The authors find ordered phases displaying collective motion for high levels of the mean density. At low densities, collective motion is only achieved for specific combinations of the parameters that determine the escape-pursuit dynamics.

Finally, an experimental study with extremely simple robots was carried out by [Tarcai et al. \(2011\)](#). In this work, they placed a group of elongated boats that are constantly propelling themselves forward in an elliptic-shaped water tank. Even in such a simple system, the usual ordered phase of collective motion emerges, here given by all boats circling the tank in the same sense: clockwise or counterclockwise. They then introduced “informed” boats that had preference for going clockwise or counterclockwise, and showed that the global sense of collective motion is highly influenced by a few informed individuals. This system was then studied by performing numerical simulations on a realistic physical model of the boat dynamics. In [Menzel and Ohta \(2012\)](#), the authors achieved flocking using local pairwise repulsive interactions between spherical self-propelled particles that can be deformed.

2.4.3 Flocking in control theory

In control theory, many flocking algorithms have been proposed. Those algorithms typically assume individuals knowing the relative distance, bearing, and orientation of their neighbors. The main scientific objective of works in control theory is the analysis of the stability of the algorithms using analytical tools. In most of the cases, actuation and sensing noise are disregarded, with few exceptions ([Moshtagh et al.,](#)

2006).

Tanner et al. (2003a,b) proposed and proved the stability of a control algorithm for flocking based on relative distance, bearing and velocity information of neighboring robots. The algorithm includes attraction/repulsion rules, which depend on local distance measurement, and the alignment rule, which depends on local velocity measurement. The attraction/repulsion force is implemented using the virtual forces framework, also very common in control theory. The authors consider two cases in their study. In the first, the neighboring topology is fixed (Tanner et al., 2003a), while in the second case it is dynamic (Tanner et al., 2003b). However, in the second case, it is assumed that the neighboring graph remains connected, that is, that there is always a neighboring path connecting any individual to any other individual in the swarm. To prove stability, the authors used graph theory and Lyapunov's stability theorem for non-smooth systems in the two cases, respectively.

Jadbabaie et al. (2003) investigated under which condition the aligned motion of particles in a noiseless version of the SDP model (Vicsek et al., 1995) is stable. They showed that stability is ensured when the neighboring graph remains connected within a finite time interval, with time divided into infinitely many irregular intervals. The authors also proposed a more relaxed condition, that is none of the neighboring graphs is connected but the union of these graphs remains connected within a finite time interval.

Regmi et al. (2005) considered the control of two complex robots via a virtual leader in the fixed topology case. The robots are equipped with high-resolution encoders for displacement and velocity measurements, and an on-board computer with a high-speed communication facility which is used to broadcast their position and velocity to the other robot. Thus, each robot has the exact absolute position and velocity information of the other robot and of the virtual leader. Although this study considers an implementation of flocking on real robots, the authors consider only two robots and assume availability of global information and communication, which makes the algorithm not very scalable to larger groups of robots.

Lindhe et al. (2005) proposed a flocking algorithm based on Voronoi partitions. The algorithm uses relative distance and orientation information of each agent at each time-step. The agents move towards the centroid of this region, keeping their formation in a desired condition, if feasible. If these conditions are not realizable, agents do not move at the current time-step. This algorithm ensures stable and collision-free flocking in environments with complex obstacles.

Olfati-Saber (2006) considered the case of flocking without a leader in an environment with obstacles. The algorithm consists of a gradient-based attraction/re-

pulsion term and a velocity-matching term. The authors showed that the algorithm boils down to Reynolds' flocking algorithm (Reynolds, 1987) and that it is not stable for group sizes larger than 10. The addition of a group objective term is shown to be essential for stable flocking behavior in obstacle-free environments. Environments with obstacles are handled by introducing virtual agents to the algorithm, which are assumed to move in the periphery of obstacles.

Moshtagh et al. (2006) proposed an algorithm to align the orientations which does not resort to explicit measurement of the agent's orientation. Instead, it relies on computer vision techniques in which agents measure relative bearing, optical flow of neighboring agents and time-to-collision between the neighboring agents and use this information to deduce the orientation. Simulations revealed that the algorithm works successfully when the neighboring graph is connected, even in the case of noisy measurements. Additionally, as said above, this work represents one of the few examples where some actual experiments with real robots were also performed.

Hanada et al. (2007) introduced an algorithm for flocking in environments with obstacles. The agents are assumed to measure the range and bearing information of their neighbors and of the obstacles. It is also assumed that the goal direction is known by all of the agents. Each agent selects two neighbors and moves in such a way that an isosceles triangle is formed among the neighbors. At the same time, agents orient themselves toward the goal direction and avoid obstacles. Simulations have shown that the agents are able to split and rejoin in the presence of obstacles, and that they form equilateral triangles in the long-run as well.

Cao and Ren (2012) studied flocking and tracking of a dynamic virtual leader. The leader interacts with only a subset of a group, and all individuals use only local interaction and partial measurements of the states of the virtual leader and neighbors. The authors considered a kinematic case where velocity measurements are not available, under both fixed and switching network topologies. They also considered a dynamic case where the velocities are available but not the accelerations. To cope with changing virtual leader velocities, a distributed estimator was used. Finally, the authors also performed numerical simulation as a proof of concept.

2.4.4 Flocking in robotics

In robotics, flocking has received sustained attention for the last two decades. Here, we focus on studies that are either directly tested on real robots or that have the potential to be applied to real robots. As with Section 2.4.2, and to reflect the main categorization of this dissertation, we classify the literature according to two categories: works that use alignment (explicit information transfer) and works that do

not use alignment (implicit information transfer).

Robotics studies with alignment

We divide studies that use alignment into two further categories. In the first, we place studies that relied on external hardware beyond what robots have on-board. In the second, we place studies using on-board hardware only.

In works belonging to the first category, authors have either estimated relative orientation of the robots ([Hayes and Dormiani-Tabatabaei, 2002](#)) or emulated an orientation sensing device ([Holland et al., 2005](#)).

[Hayes and Dormiani-Tabatabaei \(2002\)](#) proposed a flocking behavior based on collision avoidance and velocity matching. The local range and bearing measurements are emulated and broadcast to the robots from an external computer. Using this information, robots compute the position and the velocity of the center-of-mass of their neighbors, in order to maintain cohesion and to align in the same direction. Furthermore, each robot is informed about the direction to a goal area.

[Holland et al. \(2005\)](#) proposed a flocking algorithm for unmanned aerial vehicles based on Reynolds' separation, cohesion and alignment behaviors. All the sensory information (relative distance, bearing and orientation of robots neighbors) is emulated externally and broadcasted to each robot individually.

In the works that fall within the second category, a local communication unit is always used to implement alignment. [Campo et al. \(2006\)](#) performed a study in which robots have to transport an object to a nest location. Although not strictly about flocking, this work is one of the pioneering studies of coordinated motion with purely on-board local communication. Robots are equipped with an LED ring and an omni-directional camera to communicate their estimates of the nest direction to their neighbors.

[Turgut et al. \(2008a\)](#) proposed an algorithm based on proximal control (separation and cohesion) and alignment control that achieves ordered flocking motion in a random direction. They used proximity sensors to implement proximal control and a virtual heading sensor (VHS) to implement alignment control. The VHS combines a digital compass and a communication unit used by each robot to measure its orientation and to broadcast it. In this way, the orientation of a robot is sensed "virtually" by its neighbors. In a follow-up study, [Gökçe and Şahin \(2010\)](#) introduced a goal-following behavior and studied the effect of having a noisy goal direction on the swarm motion. [Çelikkanat and Şahin \(2010\)](#), inspired by the work of [Couzin et al. \(2005\)](#) in biology, provided a goal direction to some of the robots and showed that a large swarm can be guided by just a few informed robots.

Robotics studies without alignment

In most of the studies where alignment is not used, ordered motion in a direction is induced by introducing a large majority of robots that are informed about a goal direction (Matarić, 1994) or light-source direction (Spears et al., 2004).

Matarić (1994) proposed a flocking algorithm based on a set of basic components: safe-wandering, aggregation, dispersion and goal-following. The robots sense obstacles in the environment, localize themselves with respect to a set of stationary beacons and broadcast their position. In this way, robots move cohesively in a goal direction known to all robots.

Kelly and Keating (1996) proposed a flocking method based on leader-following, where the leader is dynamically elected by the swarm and follows a random direction. They used a custom-made active infra-red sensing system to sense the range and bearing of other robots and a radio-frequency system for dynamic leader election. In their work, multiple informed robots can co-exist in the swarm and, in their presence, the swarm is able to split to overcome obstacles.

Baldassarre et al. (2003) used artificial evolution to evolve a flocking behavior with a group of four simulated robots. The robots are equipped with proximity and light sensors. They use the former to perceive each others' relative position and orientation, and the latter to perceive a common goal direction.

Nembrini et al. (2002) proposed a minimalist algorithm to achieve flocking using only a local communication device, an obstacle and a beacon sensor. Some robots are informed about a goal direction and signal their status using their beacon. The other robots perform U-turn maneuvers when they lose sight of the majority of neighbors or of the signaling robots. The authors achieved a swarming behavior where robots dynamically disconnect and reconnect to the swarm.

Spears et al. (2004) proposed a flocking algorithm based on attraction/repulsion and viscous forces. The robots form a regular lattice structure using the range and bearing of their neighbors and move in a goal direction given by a light source. Due to shadowing, some of the robots in their swarm cannot sense the light and have no information about the goal direction, whereas the rest of the swarm is informed.

Barnes et al. (2009) used artificial potentials for controlling the shape of a group of unmanned ground vehicles while they move along a desired trajectory. They performed experiments with four real robots where all robots received the precise GPS position of all other robots and also the desired coordinates of the center of mass.

Moslinger et al. (2009) proposed a minimalist flocking algorithm based on attraction and repulsion zones with different threshold levels. By adjusting these levels,

they achieved flocking with a small group of robots in a constrained environment. In this work, as in our study, robots use neither goal direction nor alignment information. However, the flocking behavior is limited since the group is not able to stay cohesive all the time.

[Monteiro and Bicho \(2010\)](#) developed a method based on leader-follower dynamics to move a swarm in formation towards a target. The location of the target is known to some informed robots, assumed to be identifiable within the swarm.

[Tarcai et al. \(2011\)](#) studied a system composed of very simple remote-controlled (RC) boats subject to inelastic collisions between each other. They put the boats in a toroidal pool and observed that the swarm organizes after a certain amount time. They also studied the effect of adding informed robots into the system. This work can be considered of hybrid nature, between robotics and statistical physics, and as such it has been also discussed in Section [2.4.2](#).

Chapter 3

Self-organizing flocking method

In this chapter, we describe the methods used in this dissertation for flocking and for information transfer. We introduce novel and state-of-the-art methods for flocking and we consider them, to the best of our knowledge, for the first time in the context of information transfer. We introduce methods for explicit information transfer, whereby robots explicitly share directional information using communication, and implicit information transfer, whereby robots do not communicate but can align to a common direction via the use of a special motion control method. Additionally, we also present the method used for swarm cohesion, adapted from some previous work (Spears et al., 2004; Turgut et al., 2008a).

The chapter is organized as follows. In Section 3.1, we give an high level introduction to the methodology by describing the general principles behind explicit and implicit information transfer. We then proceed to a more technical description of the method. In Section 3.2, we present the flocking control method in general and its components. In Section 3.3, we present proximal control, the method used to achieve swarm cohesion. In Section 3.4, we present alignment control and three explicit information transfer methods. In Section 3.5, we present two motion control methods, one of them capable of achieving implicit information transfer.

3.1 Explicit and implicit information transfer

A first important question in information transfer is who possesses the information within a swarm. As done in Couzin et al. (2005), in this dissertation we assume that information is scarce and possessed only by few *informed* robots, aware of a desired *goal direction*, while the rest of the swarm is composed of *non-informed* robots, not aware of the goal direction. We also do not require robots to identify who are the informed robots in the swarm, neither via sensing nor by communicating. In most

of the experiments, we consider information transfer from the informed towards the non-informed robots, although in a few cases we consider also information transfer among non-informed robots only.

Explicit information transfer is realized via alignment control, whereby robots explicitly exchange directional information through communication strategies. As a result, robots align either to a random direction or to a desired goal direction. We considered up to two goal directions, which we call goal direction *A* and goal direction *B*. Alignment control can be seen as equivalent to the alignment rule of Reynolds (1987). In this dissertation, we implemented three different strategies, two being novel, for explicit information transfer. These three strategies are described in Section 3.4. Experiments in flocking using the three strategies are presented in Chapter 4, whereas in Chapter 5 we present experiments in which one of these strategies has been used in another collective behavior: collective transport.

Implicit information transfer is realized through motion control. The main purpose of motion control is to convert sensing information into the forward and angular velocities of the robot. With the new method we propose (that we call magnitude dependent motion control or MDMC), robots align to a random direction or to the goal direction without the need of using alignment control. We explain motion control in Section 3.5, and we present the main results in Chapter 6. Motion control methods for flocking have not received much attention in the robotics community. Indeed, it is not intuitive to understand how motion control could lead to self-organizing flocking. For this reason, in Chapter 7, we thoroughly study its properties from the statistical physics perspective.

3.2 Flocking control method

The methodological framework we use in this dissertation is inspired by the work of Turgut et al. (2008a) and is based on virtual physics (see Section 2.2.2 for an overview). In this design methodology, a metaphor is used whereby robots are seen as particles subject to virtual forces. In this dissertation, we avoid the metaphor and merely describe these forces as control vectors that provide information that is processed.

Figure 3.1 depicts two schemata: in the first one, we show the structural relationship between the different components of the method (Figure 3.1a); in the second one, we show the dynamical input/output relationships among these components (Figure 3.1b). We describe different versions of the method. The choice of the version used depends on the focus of the experiment that we carried out, as explained

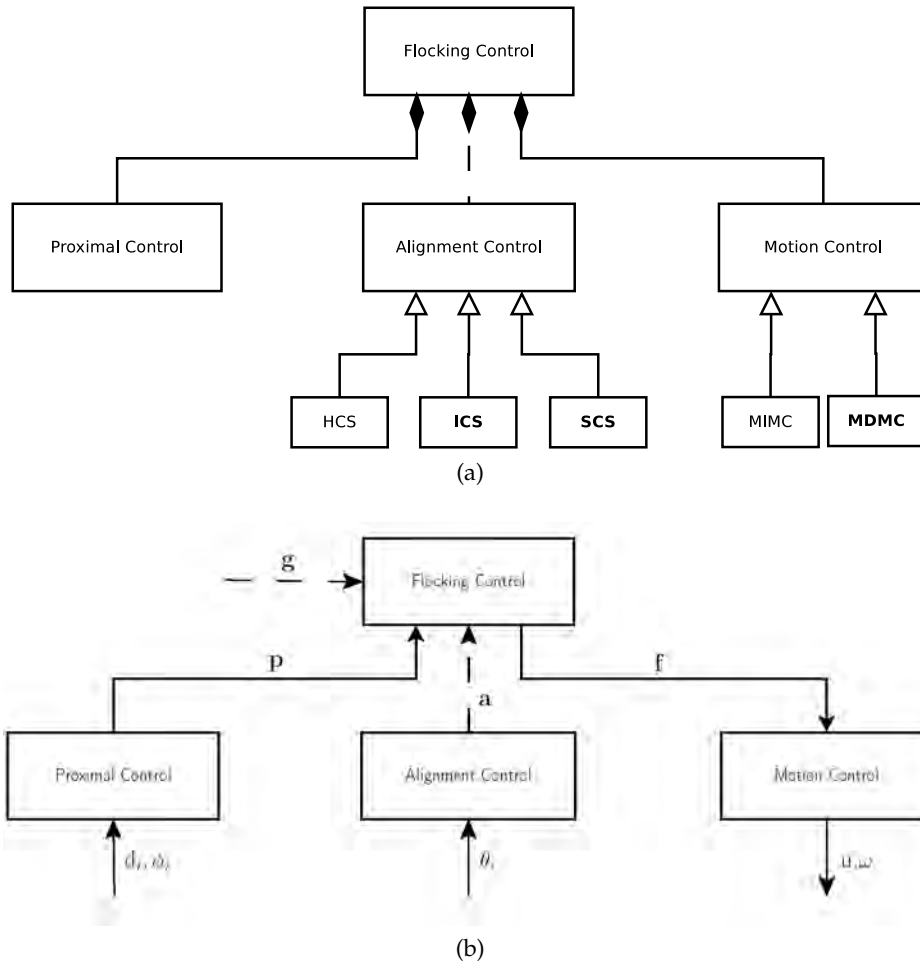


Figure 3.1: Static and input/output description of the flocking method. The are different versions of the method, depending on the set of components that are used. Figure (a) shows the different control components implemented and the relationship among them. Filled diamond symbol is used to indicate the decomposition of flocking control into the three control components: proximal, alignment and motion control. Relationships terminated by an empty triangle are *instantiations*, that indicate which versions of the component are possible. The components with the name in **bold** are the novel contribution of this dissertation. Figure (b) shows, through solid and dashed arrow lines, the input/output relationship among the components. In both graphs, solid arrow lines indicate that the component is present in all versions, whereas dashed lines indicate that the component (e.g. alignment control) can be absent in some of the versions. All the acronyms are explained in the text.

later in this section. The difference between the versions can be due to: a) alignment control being used or not (explicit versus implicit information transfer); b) the goal direction being used or not (with or without informed robots); c) different communication strategies for alignment control or different motion control methods.

We describe how the method works from the point of view of a single robot that

we call the *focal robot*. The focal robot computes a flocking control vector denoted by \mathbf{f} . The expression of the flocking control vector depends on the version of the method that is used. In the most complete version of flocking control, \mathbf{f} is computed using what we call Model 3.1:

$$\mathbf{f} = \alpha \mathbf{p} + \beta \mathbf{a} + \gamma \mathbf{g}_j. \quad (3.1)$$

Here, \mathbf{p} is the proximal control vector, which is used to achieve cohesion among robots and it is the output of proximal control described in Section 3.3; \mathbf{a} is the alignment control vector, that is used for explicit information transfer and it is the output of alignment control described in Section 3.4; \mathbf{g}_j is the goal direction vector, that encodes the goal following behavior used only by the informed robots to follow either goal direction A ($j = 1$) or goal direction B ($j = 2$). The parameters α , β and γ are used to control the relative importance of the three components.

We also consider other three versions of the algorithm (that we call Model 3.2, Model 3.3 and Model 3.4):

$$\mathbf{f} = \alpha \mathbf{p}, \quad (3.2)$$

$$\mathbf{f} = \alpha \mathbf{p} + \gamma \mathbf{g}_j, \quad (3.3)$$

$$\mathbf{f} = \alpha \mathbf{p} + \beta \mathbf{a}. \quad (3.4)$$

A focal robot using Model 3.2 is the most minimalistic robot, as it only uses proximal control. Intermediate cases are represented by Model 3.3 and Model 3.4. Model 3.1 and Model 3.4 include alignment control, therefore use explicit information transfer to share directional information. Conversely, Model 3.2 and Model 3.3 do not include alignment control, and as such can achieve flocking only through implicit information transfer. For this reason, to carry out explicit information transfer experiments in Chapter 4, we used Model 3.1 (for informed robots) and Model 3.4 (for non informed robots), and we considered and studied all the communication strategies for alignment control presented in Section 3.4. Conversely, to carry out implicit information transfer experiments in Chapter 6, we studied Model 3.2 and Model 3.3, as they do not include alignment control, and we compared the different motion control methods presented in Section 3.5.

3.3 Cohesion via proximal control

Proximal control is present in all the methods considered in this dissertation as we always require the swarm to stay cohesive. The main idea behind proximal control is that, in order to achieve cohesive flocking, the focal robot needs to keep a certain distance from its neighbors. The proximal control vector encodes the attraction and repulsion rules: the focal robot tends to move closer to its neighbors when the distance to them is higher than the desired distance and tends to move away from them when this distance is lower than the desired distance.

Proximal control assumes that the focal robot is able to sense the range and bearing of its neighboring robots within a maximum interaction distance D_p . The proximal control vector \mathbf{p} is computed as:

$$\mathbf{p} = \sum_{i \in \mathcal{N}_p} p(d_i) e^{j\phi_i}.$$

Here, $p(d_i) e^{j\phi_i}$ is a vector expressed in the complex plane, having magnitude $p(d_i)$ and angle ϕ_i ; \mathcal{N}_p denotes the set of neighboring robots perceived by the focal robot within range D_p ; and d_i and ϕ_i denote the relative range and bearing of the i^{th} neighboring robot, expressed in the body-fixed reference frame of the focal robot¹, respectively (Figure 3.2). The magnitude of the vector, $p(d_i)$, is computed as:

$$p(d_i) = -\frac{\partial P(d_i)}{\partial d_i} = -8\epsilon \left[2\frac{\sigma^4}{d_i^5} - \frac{\sigma^2}{d_i^3} \right].$$

This expression for $p(d_i)$ is defined as the negative derivative of the virtual potential function $P(d_i)$. A virtual potential function is often used in swarm robotics and statistical physics to model social interaction between robots (Spears et al., 2004) or particles (Grégoire et al., 2003). In our study, we use a modified version of the Lennard-Jones potential function:

$$P(d_i) = 4\epsilon \left[\left(\frac{\sigma}{d_i} \right)^4 - \left(\frac{\sigma}{d_i} \right)^2 \right].$$

The parameter ϵ determines the strength of the attraction/repulsion rule, whereas the desired distance d_{des} between the robots is linked to the parameter σ according to the formula $d_{des} = 2^{1/2}\sigma$.

A final note has to be made on the maximum interaction distance D_p . In general,

¹The body-fixed reference frame is fixed to the center of a robot, its x -axis points to the front of the robot and its y -axis is coincident with the rotation axis of the wheels.




Figure 3.2: The body-fixed reference frame ($x-y$) used on the robot, and the global reference frame ($X-Y$) common to all robots. The curvy arrow denotes the sense of rotation of positive angles.

D_p can be less than or equal to the maximum sensing range of the sensor. In studying implicit information transfer, in order for the proposed model to work in all cases, we found that the following precondition must hold: D_p has to be chosen so that, when the swarm reaches a stable formation, interactions among robots are limited to the first neighbors only in the Voronoi sense (Voronoi, 1908). In order to achieve this without having to compute the Voronoi tessellation, we set $D_p = 1.8 \cdot 2^{\frac{1}{2}} \sigma$, that is, the maximum interaction range is less than two times the desired distance $2^{\frac{1}{2}} \sigma$ between two robots. In this way, second neighbors (in the Voronoi sense) are outside the scope of proximal control.

3.4 Explicit information transfer via alignment control

Alignment control represents the method used in this dissertation to achieve direct information transfer. The main idea of alignment control is to have the focal robot computing the average of the directional information sent by its neighbors in order to achieve an agreement to a common direction with its neighbors.

Alignment control assumes that the focal robot can measure its own orientation θ_0 with respect to the reference frame common to all robots (see Figure 3.2). It also assumes the focal robot can send directional pieces of information, denoted as θ_{s_0} , using a communication device. The value of θ_{s_0} depends on the communication strategy that is being used, as described in Section 3.4.1, 3.4.2 and 3.4.3. The robot receives the information θ_{s_i} sent by one of its neighbors in the set \mathcal{N}_a within a given range D_a . The information represents directions expressed with respect to the com-

mon reference frame (Figure 3.2). Once received, each θ_{s_i} is converted into the body-fixed reference frame of the robot. In order to compute the average of the received directional information, all directions are converted into unit vectors with angles equal to θ_{s_i} , which are then summed up and normalized:

$$\mathbf{a} = \frac{\sum_{i=0}^{m_a} e^{j\theta_{s_i}}}{\left\| \sum_{i \in \mathcal{N}_a} e^{j\theta_{s_i}} \right\|},$$

We present three different communication strategies: heading communication strategy (HCS), information-aware communication strategy (ICS) and self-adaptive communication strategy (SCS).

3.4.1 Heading communication strategy (HCS)

In HCS, the piece of information θ_{s_0} sent by a robot to its neighbors is its own orientation $\theta_{s_0} = \theta_0$, measured with respect to the common reference frame. This strategy is used to reproduce the capability of a robot i to “sense” the orientation of a neighboring robot j , by making robot j communicate its own orientation to robot i .

3.4.2 Information-aware communication strategy (ICS)

ICS assumes that each robot is aware of whether it is *non-informed* or *informed*. If it is *non-informed*, it sends $\theta_{s_0} = \angle \mathbf{a}$ ($\angle \cdot$ denotes the angle of a vector) to its neighbors; otherwise, if it is *informed*, it sends $\theta_{s_0} = \angle \mathbf{g}_j$. The intuitive motivation behind this strategy is the following: in case the robot is *non-informed*, it helps the diffusion of the information originating from the *informed* robots; if instead it is *informed*, it directly propagates the information it possesses to its neighbors. Using this mechanism, the information eventually reaches the entire swarm. Note that only in HCS the communicated angle coincides with the robot’s current state (orientation), whereas this is not true for the other communication strategies.

3.4.3 Self-adaptive communication strategy (SCS)

SCS extends ICS by introducing a parameter w_t that represents the degree of confidence a robot has about the utility of its possessed information. The communicated directional information is computed in this way:

$$\theta_{s_0} = \angle [w_t \mathbf{g}_j + (1 - w_t) \mathbf{a}].$$

For *non-informed* robots, $w_t = 0$ (they do not possess information about \mathbf{g}_j). For *informed* robots, when $w_t = 1$, this strategy coincides with ICS. In SCS, however, we use the following rule to change w_t :

$$w_{t+1} = \begin{cases} w_t + \Delta w & \text{if } \|\mathbf{a}'\| \geq \mu; \\ w_t - \Delta w & \text{if } \|\mathbf{a}'\| < \mu, \end{cases}$$

where μ is a threshold and Δw is a step value, whose role is to control the dynamics of w_t as explained in the following. The quantity:

$$\|\mathbf{a}'\| = \left\| \frac{\sum_{i=0}^k e^{j\theta_{s_i}}}{k+1} \right\|$$

is the magnitude of the local consensus vector, or simply, the local consensus. We choose this quantity because inspired by the decision-making mechanism used by the Red Dwarf honeybee (*Apis florea*, the European honeybee): to perform nest selection, these bees wait to achieve locally a consensus to a given nest location before flying off (Makinson et al., 2011; Diwold et al., 2011).

The rationale behind SCS is the following. Informed robots communicate the goal direction when the local consensus $\|\mathbf{a}'\|$ is higher than the threshold μ . Local consensus measures how close the received pieces of information are to each other and to the information sent by the focal robot. When local consensus is close enough to 1, the angles received from the robot's neighbors are very similar to each other and to the angle sent by the focal robot. In this case, the focal robot increases (through steps of size Δw) its degree of confidence w_t on the possessed information \mathbf{g}_j . Conversely, when the local consensus is lower than the threshold μ , then the focal robot detects conflicting goal directions in the swarm. In this case, the focal robot decreases (through steps of size Δw) its level of confidence w_t on the goal direction \mathbf{g}_j , until $w_t = 0$, in which case the focal robot behaves as a non-informed robot. This mechanism facilitates the propagation of highest priority directional information that is available within a swarm. Note that the confidence w_t is a private piece of information that is never communicated by the focal robot. This makes our method applicable to a vast majority of robots, including not only robots with limited communication capabilities but also robots that communicate only using visual information (LEDs and cameras) that can only exchange directional information.

3.5 Implicit information transfer via motion control

Motion control is the method whereby implicit information transfer between robots is realized. Motion control is used to convert the flocking control vector into the forward speed u and the angular speed ω of the focal robot. Below, we describe the two motion control methods considered in this dissertation: magnitude dependent motion control (MDMC) and magnitude independent motion control (MIMC). Among these two methods, only through MDMC it is possible to achieve flocking via implicit information transfer, that is, without resorting to alignment control. The detailed explanation of why this is the case is very complex and is therefore not included in this but in a dedicated chapter: Chapter 7.

3.5.1 Magnitude dependent motion control (MDMC)

MDMC has been published and studied in [Stranieri et al. \(2011a\)](#) (with a different name: Variable forward-speed Motion Control, or VMC) and in [Ferrante et al. \(2012a\)](#). In MDMC, the forward and angular speed of the robot depend on both the magnitude and the direction of the flocking control vector. Therefore, MDMC uses all the information contained in this vector.

We define MDMC as follows. Let \mathbf{o} be the unit vector with direction equal to the orientation of the focal robot. Accordingly, $f_x = \mathbf{f} \cdot \mathbf{o}$ and $f_y = \mathbf{f} \cdot \mathbf{o}^\perp$ denote the projection of the flocking control vector (\mathbf{f}) on the x - and y -axis of the body-fixed reference frame, respectively. We let the forward speed u be directly proportional to the x component of the vector, pointing in the direction of motion of the robot. The angular speed ω , instead, is made directly proportional to the y component of the vector, pointing in the direction perpendicular to the motion of the robot. The corresponding equations are

$$u = K_1 f_x + U;$$

$$\omega = K_2 f_y,$$

where U is the forward biasing speed and K_1 and K_2 are the linear and angular gains, respectively.

The rationale of this method is the following: the larger the x -component of the flocking control vector, the faster we want the robot to move forward; the larger the y -component of the flocking control vector, the faster we want the robot to turn. Note that, with MDMC, robots can in principle move also backwards, but they have a propension to move forward due to the presence of the $+U$ term in the equation for the forward speed.

3.5.2 Magnitude independent motion control (MIMC)

MIMC was introduced for the first time in [Turgut et al. \(2008a\)](#) and later used in subsequent work by [Çelikkanat and Şahin \(2010\)](#). In MIMC, the forward and angular speed of the robot partially depend on the information contained in the flocking control vector. By “partially” we mean that, for calculating the speeds, only the direction of the flocking control vector is taken into account and the magnitude is discarded.

Let \mathbf{o} be the unit vector with direction equal to the orientation of the focal robot. The forward velocity and the angular velocities are computed as:

$$u = \begin{cases} \left(\mathbf{o} \cdot \frac{\mathbf{f}}{\|\mathbf{f}\|} \right) U, & \text{if } \mathbf{o} \cdot \frac{\mathbf{f}}{\|\mathbf{f}\|} \geq 0; \\ 0, & \text{otherwise;} \end{cases}$$

$$\omega = K_3 (\angle \mathbf{o} - \angle \mathbf{f}),$$

where K_3 is the angular gain and \cdot denotes the dot product of two vectors.

The rationale of the method is the following. The robot modulates its forward speed according to the angular difference between the flocking control vector and its direction of motion. When the two directions are aligned, the robot moves at maximum speed U . When the angular difference is equal to or larger than 90 degrees, the robot does not move forward but rotates on spot. In between these two cases, the forward speed is interpolated between 0 and U according to the dot product between the flocking control vector and the direction of motion. The angular speed is instead proportional to the angular difference between the flocking control vector and the direction of motion. Note that, with MIMC, the robot are not able to move backwards.

3.5.3 Wheel speed computation

To compute the speed of the left and right wheel, we first limit the forward speed to the interval $[-U, U]$, and the angular speed to the interval $[-\Omega, \Omega]$. We then use the differential drive model as in [Turgut et al. \(2008a\)](#) to convert the forward speed u and the angular speed ω into the linear speeds of the left (N_L) and right (N_R) wheel:

$$N_L = u + \frac{\omega}{2}l,$$

$$N_R = u - \frac{\omega}{2}l,$$

where l is the distance between the wheels.

Chapter 4

Explicit information transfer via communication

In this chapter, we present the experiments performed with explicit information transfer. In general, explicit information transfer requires robots capable of sharing directional information via a communication device such as a local radio (as we do here) or visual communication (e.g. via LED signaling and camera detection, as done in previous studies performed within the swarm-bots project¹). After presenting the results on flocking in this chapter, in Chapter 5 we present further experiments of explicit information transfer on another collective behavior: collective transport.

This chapter is organized as follows. After the problem description and motivation (Section 4.1), we present the metrics and the experimental setup (Section 4.2 and Section 4.3). We then consider experiments with a single stationary goal direction (Section 4.4) and a single non-stationary goal direction (Section 4.5) followed by experiments with two stationary goal directions (Section 4.6) and two non-stationary goal directions (Section 4.7). We conclude the chapter with a summary and discussion (Section 4.8).

4.1 Problem description

Explicit information transfer is being studied in biology as one of the main mechanisms underlying flocking. In some situations, animals achieve flocking in presence of multiple, static or dynamic, sources of information with different priorities. An example is represented by the dynamics of some animals that are subject to attacks by predators. The escape direction from a predator is a very fast changing piece of

¹Swarm-bots project, <http://http://www.swarm-bots.org/>, March 2013.

information and it might possibly conflict with the direction to a food source. To deal with these situations, animals developed communication mechanisms to transfer information effectively and efficiently throughout the group (Franks et al., 2007; François et al., 2006).

By analogy, similar situations can be faced also by robotic swarms in some example scenarios, such as: a swarm that, while following a curvy trajectory, needs to change direction often and rapidly; a swarm going in one direction while avoiding an obstacle; a swarm that has to avoid a dangerous location while going to a target location. Within large swarms, the goal location and the dangerous locations might be perceived by a small proportion of informed robots. We can imagine this happening in at least two possible ways. In the first scenario, we might have only a few robots equipped with advanced sensors required for getting directional information from the environment. Here, the robots with these advanced sensors can be assumed to be randomly distributed in the swarm. In the second scenario, all the robots would be equipped with the same sensors, but only some robots might have access the relevant directional information due to their privileged position within the swarm. For example, only the robots in the front might be able to sense the goal direction as they can directly sense it through a camera, while the others are shadowed by other robots. In this case, there is a spatial correlation between the relative location of the robots in the swarm and the information they possess. In all these situations, a typical objective would be to get all robots to a goal area without losing any, that is, by keeping the swarm cohesive, even when there is a dangerous area to be avoided on the way.

Motivated by the above examples, in this chapter we study explicit information transfer of one and of two distinct goal directions in a flocking robot swarm. We consider the case of stationary and non-stationary goal directions possessed by a small minority of robots, the informed robots, that need to influence the movement of the entire swarm while the latter needs to keep cohesion. We perform experiments using the communication strategies explained in Section 3.4. After introducing the experiment with one goal direction (Section 4.4 and Section 4.5), we present the main experiments that focus on the two goal directions case. Here, we require the swarm to remain cohesive, and we tackle two different macroscopic objectives: a) move to the average direction among the two goal directions (for example to avoid an obstacle) (Section 4.6); b) select the most important of the two goal directions (for example the direction that allow the robots to avoid a danger) and follow it (Section 4.7). We denote the two goal directions as goal direction A and goal direction B . In the experiments presented in Section 4.6, the two goal directions have the same priority,

they are present for the entire duration of the experiment, and the goal of the swarm is to follow the average direction between the two. In the experiments presented in Section 4.7, goal direction B has a higher priority with respect to goal direction A and is present during only a limited period of time.

4.2 Metrics

Here, we describe analytically the metrics used to measure the performances. The first two metrics are also utilized in Chapter 6.

Order: The order metric ψ measures the degree of agreement of the orientations of the robots in the swarm, and has been widely adopted both in physics (Vicsek et al., 1995) and in robotics (Turgut et al., 2008a). To define it, we first need to compute the vectorial sum of the orientations of all N robots:

$$\mathbf{b} = \sum_{i=1}^N e^{j\theta_i}.$$

The order is then defined as:

$$\psi = \frac{1}{N} \|\mathbf{b}\|,$$

such that $\psi \approx 1$ when the robots have a common orientation and $\psi \approx 0$ when they point in different directions.

Average direction: For experiments with two stationary goal directions, we measure the average swarm direction. It is obtained by computing the average direction of the orientations of all robots:

$$\angle \mathbf{b} = \angle \sum_{i=1}^N e^{j\theta_i}.$$

We compare the average group direction against the theoretical average direction, which takes into account the number of robots informed about each goal direction:

$$\hat{\theta} = \angle(N_A \mathbf{g}_1 + N_B \mathbf{g}_2),$$

where N_A (N_B) is the number of robots informed about goal direction A (goal direction B).

Accuracy: The accuracy metric δ measures how accurately robots follow the goal direction. It has been used both in a seminal biology paper by Couzin et al. (2005) but

also later in robotics by [Çelikkanat and Şahin \(2010\)](#). Consistently with [Çelikkanat and Şahin \(2010\)](#), we define the accuracy metric in the $[0, 1]$ interval for convenience:

$$\delta = 1 - \frac{1 - \psi \cos(\angle \mathbf{b} - \angle \mathbf{g}_i)}{2}.$$

Here, $\angle \mathbf{b}$ is the average group direction and $\angle \mathbf{g}_i$ is goal direction A or B . Robots having a common orientation and moving in the goal direction will have $\delta \approx 1$ and $\psi \approx 1$. Conversely, if robots are moving in a direction which is opposite to the goal direction (even if they are well aligned), then $\delta \approx 0$. Note that, when the swarm is not aligned, the value of the accuracy is $\delta \approx 0.5$, since $\psi \approx 0$ nullifies the contribution of the cosine term.

Number of groups: We use the number of groups present at the end of an experiment to detect whether the swarm has split into different groups or has kept cohesion and stayed as a single group. We use the following criteria to define what a *group* is and how to calculate the number of groups. We first find the distance between all pairs of the robots. If the distance between the robots in a pair is smaller than the maximum sensing range of the range and bearing sensor (2 meters), we set it as an equivalence pair and append it to the list containing the other equivalence pairs. We then use the equivalence class method on the list to determine the equivalence class of each pair. The total number of equivalence classes calculated is equal to the number of groups. For the details of the equivalence class method please refer to [Press et al. \(1992\)](#).

4.3 Experimental setup

In this section, we describe the experimental setups used in simulation and with the real robots.

4.3.1 Simulation experimental setup

We execute experiments in simulation using the ARGoS simulator, which is described in more details in [Appendix A.2](#). ARGoS allows the cross-compilation of the controllers both in simulation and on real-robots. This allowed us to seamlessly port the same controller studied in simulation to real robot.

In the experiments, N mobile simulated robots are placed at random positions in a circle of variable radius and with random orientations uniformly distributed in

Variable	Description	Value
N	Number of robots	{100, 300}
R	Number of runs per setting	100
ρ_{I1}	Proportion of robots informed about $\angle \mathbf{g}_1$	{0.01, 0.1}
ρ_{I2}	Proportion of robots informed about $\angle \mathbf{g}_2$	0.1
$T_{\Delta p}$	Duration of two-goal phase	600 s
T_s	Duration of experiments in stationary environments	300 s
T_n	Duration of experiments in one-goal non-stationary environments	$4T_s$ s
T_p	Duration of experiments in two-goal non-stationary environments	$T_s + T_{\Delta p} + 2T_s$ s
α	Weight of proximal control	1
β	Weight of alignment control	4
γ	Weight of goal direction	1
μ	Threshold value used in SCS	0.999
Δw	Step value used in SCS	0.1
U	Maximum forward speed	20 cm/s
Ω_{max}	VMC max angular speed	$\pi/2$ rad/s
K_1	VMC linear gain	0.5 cm/s
K_2	VMC angular gain	0.06 rad/s
l	Inter-wheel distance	0.1 m
ϵ	Strength of attraction-repulsion	1.5
σ	Distance-related proximal control parameter	0.4 m
d_{des}	Desired inter-robot distance	0.56 m
D_p	Maximum perception range of proximal control	1.0 m
D_a	Maximum perception range of alignment control	2.0 m
ξ	Amount of noise	0.1
Δt	ARGoS integration time-step and real robot control step	0.1 s

Table 4.1: Experimental values or range of values for all constants and variables used in simulation

Variable	Description	Value
N	Number of robots	8
R	Number of runs per setting	10
ρ_{I1}	Proportion of robots informed about $\angle \mathbf{g}_1$	0.125
ρ_{I2}	Proportion of robots informed about $\angle \mathbf{g}_2$	0.125
$T_{\Delta p}$	Duration of two-goal phase	100 s
T_s	Duration of experiments in stationary environments	100 s
T_n	Duration of experiments in one-goal non-stationary environments	$2T_s$ s
T_p	Duration of experiments in two-goal non-stationary environments	$50 + T_{\Delta p} + 50$ s
	All the other control parameters	See Table 4.1

Table 4.2: Experimental values or range of values for all constants and variables used with the real robots. Note that all the parameters related to the controllers are the same as in simulation, that is, the controller used on the real robot is exactly the same as in simulation.

the $[-\pi, \pi]$ interval. A light source is also placed at a fixed position in the arena, far away from the robot but with a very high intensity.

We conducted four sets of experiments:

1. **Stationary environment:** A stationary environment is an environment with only one goal direction, which is fixed at the beginning and does not change over time. In stationary environments, we randomly select a proportion ρ_{I1} of robots and we inform them about goal direction A . All the other robots remain uninformed during the entire experiment. Goal direction A is selected at random in each experiment. The duration of one run is T_s simulated seconds.
2. **One-goal non-stationary environment:** A one-goal non-stationary environment or, in short, non-stationary environment, is an environment where there is only one goal direction that does not change for an amount of time and then changes as a step function. This process repeats itself four times during a single experiment. Thus, a non-stationary environment consist of four stationary phases of equal duration. The proportion of informed robots ρ_{I1} is kept fixed during the entire run. However, goal direction A and the informed robots are randomly re-selected at the beginning of each stationary phase. The duration of the entire run is T_n simulated seconds.
3. **Two-goal stationary environment:** A two-goal stationary environment is an environment in which two goal directions are present during the entire experiment. A proportion of ρ_{I1}, ρ_{I2} robots are informed about goal directions A and B , whose value are here denoted by $\angle g_1$ and $\angle g_2$, respectively. Thus here, $N_A = N\rho_{I1}$ is the number of robots informed of goal direction θ_1 and $N_B = N\rho_{I2}$ is the number of robots informed of goal direction θ_2 .
4. **Two-goal non-stationary environment:** A two-goal non-stationary environment is an environment where goal direction A is present for the entire duration of the experiment and goal direction B is present only within a time window that lasts $T_{\Delta p}$. In two-goal non-stationary environments, we first randomly select a proportion ρ_{I1} of robots that are informed about goal direction A . At a certain time T_s , we randomly select a proportion ρ_{I2} of robots that are informed about goal direction B . To capture the most difficult case, which corresponds to the case with maximal conflict (angular difference) between the two goal directions, we let goal direction B always point to the opposite direction with respect to goal direction A . At time $T_s + T_{\Delta p}$, we reset all informed robots and we re-sample a proportion ρ_{I1} of robots and we make them informed about goal direction A for additional $2T_s$ simulated seconds. We call the phase between




Figure 4.1: Two pictures that explain the two selection mechanisms. (a) Non-spatial selection in stationary environment: grey circles, that represent robots informed about goal direction A , are selected at random locations in the swarm (the white circles represent non-informed robots). (b) Spatial selection during the two goal phase in two-goals non-stationary environment: informed robots (grey and black circles) are selected at the periphery of the swarm. Grey circles represent robots informed about goal direction A (which in this case points left), whereas black circles represent robots informed about goal direction B (which in this case points right).

time T_s and time $T_s + T_{\Delta p}$ the two-goal phase. The total duration of one run is $T_p = 3T_s + T_{\Delta p}$ simulated seconds. The proportion ρ_{I2} is always set to 0.1. Note that robots informed about goal direction B use SCS with fixed $w_t = 1$ as they possess the information with the highest priority and as such they do not need to change their confidence into their goal direction.

Experiments in Environment 1, 2 and 4 are further classified according to how informed robots are selected. This selection mechanism is either *non-spatially* or *spatially* correlated. Figure 4.1 depicts the difference between the two selection mechanisms.

Non-spatial selection: With this selection mechanism, the informed robots are selected at random at the beginning of each stationary phase (see Figure 4.1a).

Spatial selection: With this selection mechanism, informed robots are selected in a way such that they are always adjacent to each other. Furthermore, the selected robots are at the periphery of the swarm and their relative position is correlated to the goal direction (see Figure 4.1b).

In experiments in Environment 1, 2 and 4, we compare all the communication strategies presented in Section 3.4 (HCS, ICS and SCS) and we study the impact of the proportion of informed robots ρ_{I1} and of the swarm size N . The format of the plots

is always the same. On the same row we report results with the same number of robots (N), whereas on the same column we report results with the same proportion of informed robots (either 1% or 10%). In plots where the x -axis is time, we report the median values (50% percentile), the first and the third quartile (25% and the 75% percentiles).

In the experiments in Environment 3, we only considered HCS as it was sufficient to reach our objectives. We categorize these experiments in a different way compared to the ones in the other three environments. In particular, here we are interested in determining i) the impact of the difference between N_A and N_B and ii) the impact of the difference between two goal directions $\angle \mathbf{g}_2 - \angle \mathbf{g}_1$ and iii) the impact of the total proportion of informed robots. To study i), we classify the experiments in three sets: no difference between N_A and N_B ($N_A = N_B$), small difference ($N_A - N_B = 2$) and large difference ($N_A - N_B > 2$). To study ii) and to reduce the parameter space, we fix the $\theta_1 = 0$, and we only vary $\theta_2 \in \{10, 20, \dots, 170, 179, 180\}$. To study iii), we consider the following values of ρ_{I1} and ρ_{I2} (proportion of informed robots): $\{(0.01, 0.01), (0.01, 0.05), (0.02, 0.04), (0.1, 0.1), (0.09, 0.11), (0.01, 0.19)\}$.

In all of the experiments, we add noise to several components of our system: to the orientation measurement θ_0 , to the proximal control vector \mathbf{p} and to the goal directions $\mathbf{g}_1, \mathbf{g}_2$. We consider noise only in angle, as commonly done in flocking studies (Vicsek et al., 1995; Turgut et al., 2008a), and we model it as a variable uniformly distributed in the $[-\xi 2\pi, +\xi 2\pi]$ range. The parameter ξ is used to control the magnitude of the noise.

4.3.2 Real robot experimental setup

In order to validate the experiments performed in simulation in Environment 1, 2 and 4, we also performed experiments on real robots. More details on the used robots and on the implementation can be found in Appendix A. At the beginning of each run, eight robots are placed in the arena depicted in Figure 4.2, each with a random orientation. At the left of the arena, a light source area is present. To measure order and accuracy over time, we built a custom-made tracking system. We place carton hats, having a directional marker, on top of each robot². This marker is detected by an overhead camera placed on the back side of the arena, at an height of about 3 meters and pointing on the ground towards the arena (Figure 4.2 has been obtained by this camera). We recorded a movie for each experiment and we then analyzed each video off-line using the Halcon software³. The analysis of a video

²Note that such hats are used for tracking purposes only and are not detectable by the robot themselves.

³<http://www.halcon.de/>

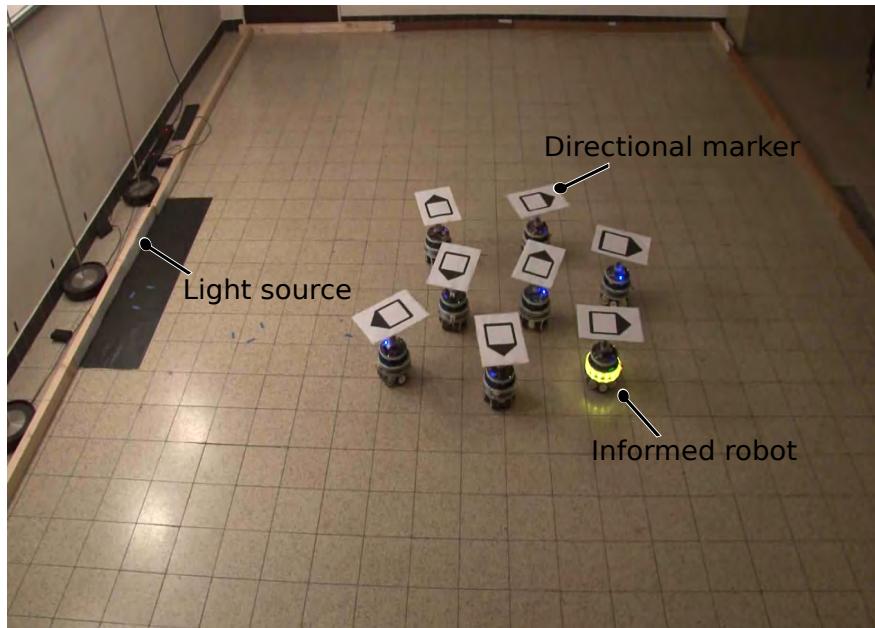


Figure 4.2: The arena as seen by the overhead camera used for tracking: on the left side we created a light source by placing four lamps; a carton hat with a directional marker is placed on each robot, in order to detect its orientation to compute the required metrics; the glowing robots are informed about the desired goal direction (this example is taken from experiments with informed robots). Note that LEDs and carton hats are used for debugging, visualization and measurement purposes only, but are never utilized in the controller algorithm.

produced a file containing, for each frame, the orientation of every robot detected.

The setting used on the robots is the same as in simulation (Section 6.3.1), with only two exceptions: the one-goal non-stationary environment consists of 2 stationary phases instead of 4 and the duration of the phases in all three settings are different and summarized in Table 4.2. We decided to reduce the duration of each experiment due to the limited size of the arena, which does not allow very long experiments involving robots that keep on going in one direction during the entire experiment. Furthermore, since experiments in simulation showed almost no difference in results between non-spatial and spatial selection, and due also to the limited size of the real robot swarm, on the real robots we consider only the non-spatial selection case.

For each experimental setting and for each of the three strategies, we execute 10 runs and we report the median values, the first and the third quartile. Since we are considering few runs, we also perform the Wilcoxon rank sum test to validate the statistical significance of our claims. The statistical test is performed by comparing vectors containing the time-average performance of a given method during a given




Figure 4.3: Results in simulation. HCS, ICS and SCS in the stationary environment using the non-spatial selection mechanism: effect on the accuracy. Ticker (central) lines represent the medians of the distributions, whereas thinner lines represent the 25% and the 75% percentiles.

phase of the experiment.

Table 4.2 summarizes all the parameters of the setup, while the parameters of the controllers are reported in Table 4.1 as they are the same as those used in simulation. We manually tune all parameters to the values reported in the table. In particular,

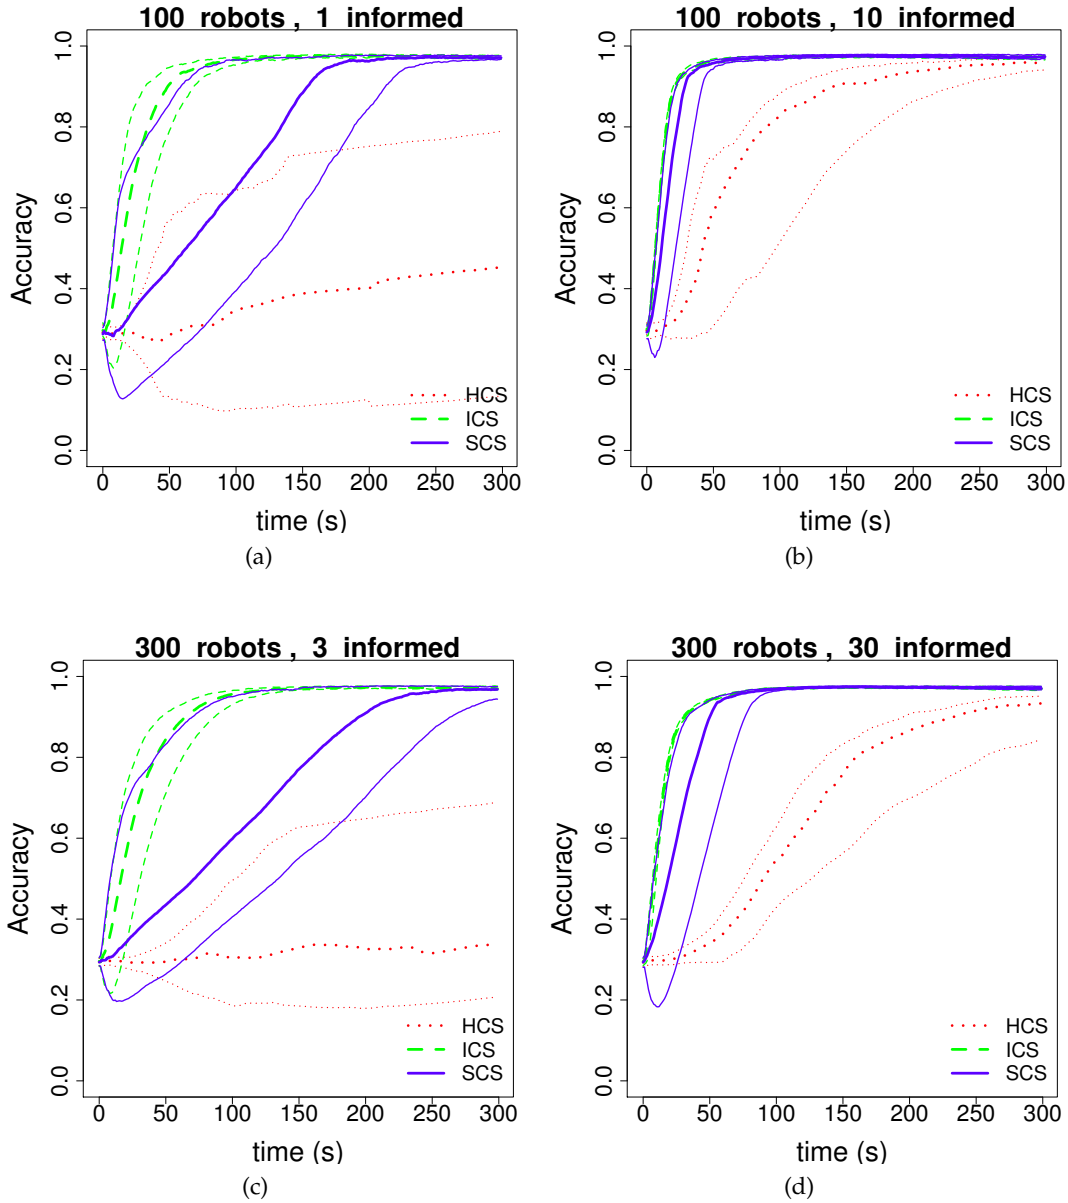


Figure 4.4: Results in simulation. HCS, ICS and SCS in the stationary environment using the spatial selection mechanism: effect on the accuracy. Ticker (central) lines represent the medians of the distributions, whereas thinner lines represent the 25% and the 75% percentiles.

the choice of values for the parameters of SCS (Section 3.4) is justifiable in this way: Δw is set to 0.1 as larger values would produce large fluctuations of w while smaller values would correspond to a slower convergence time, and μ is set to 0.999 as it is enough to detect low local consensus with a very good precision. This is in turn possible due to fact that the range and bearing communication device is noise-free.

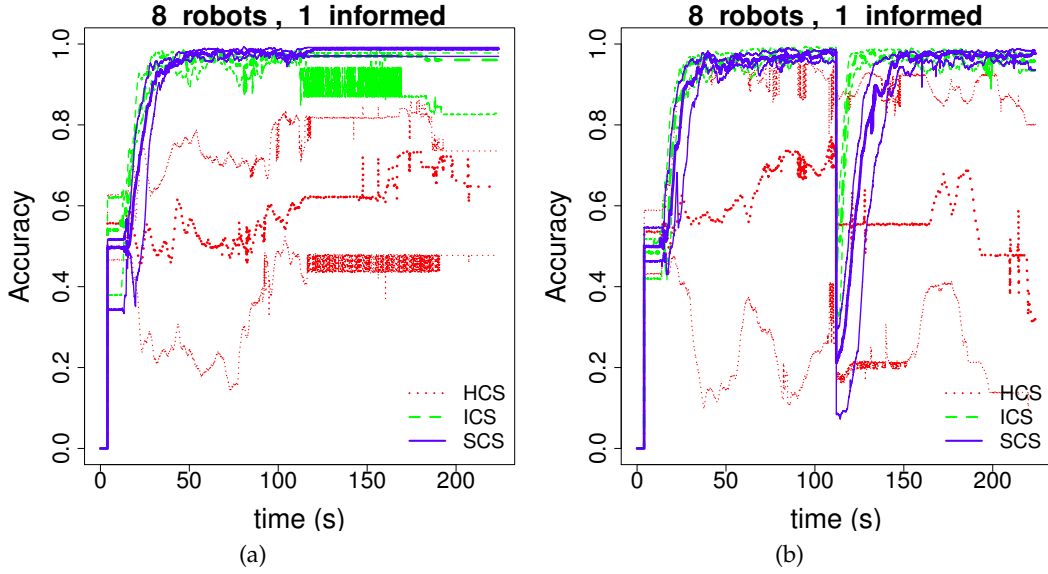


Figure 4.5: Results with real robots. HCS, ICS and SCS in one-goal stationary and non-stationary environments. Figures (a), (b) plot the distribution of the accuracy metric over time for stationary, one-goal non-stationary, respectively. Ticker (central) lines represent the medians of the distributions, whereas thinner lines represent the 25% and the 75% percentiles.

4.4 One stationary goal direction

4.4.1 Simulation

Figure 4.3 shows the results obtained in stationary environments when using a non-spatial selection mechanism. Figure 4.3a and Figure 4.3c show that ICS outperforms the other two strategies when only 1% of the robots are informed. When we consider the median values, SCS reaches the same level of accuracy as ICS in a slightly larger amount of time. In the best runs (above the 75% percentile), performance of SCS is very close to those obtained with ICS, whereas in the worst runs (below the 25% percentile), results are slightly worse. We also observe that results with HCS have larger fluctuations than the one obtained with the other two strategies. Additionally, the novel strategy SCS shows a reasonable level of accuracy compared to ICS and performs much better than HCS.

Figure 4.4 shows the results obtained in stationary environments when using a spatial selection mechanism. Figure 4.4a and Figure 4.4c show that, when only 1% of the robots are informed, the median values of SCS is slightly worse with respect to the non-spatial selection case (Figure 4.3c). This can be explained by the fact that,

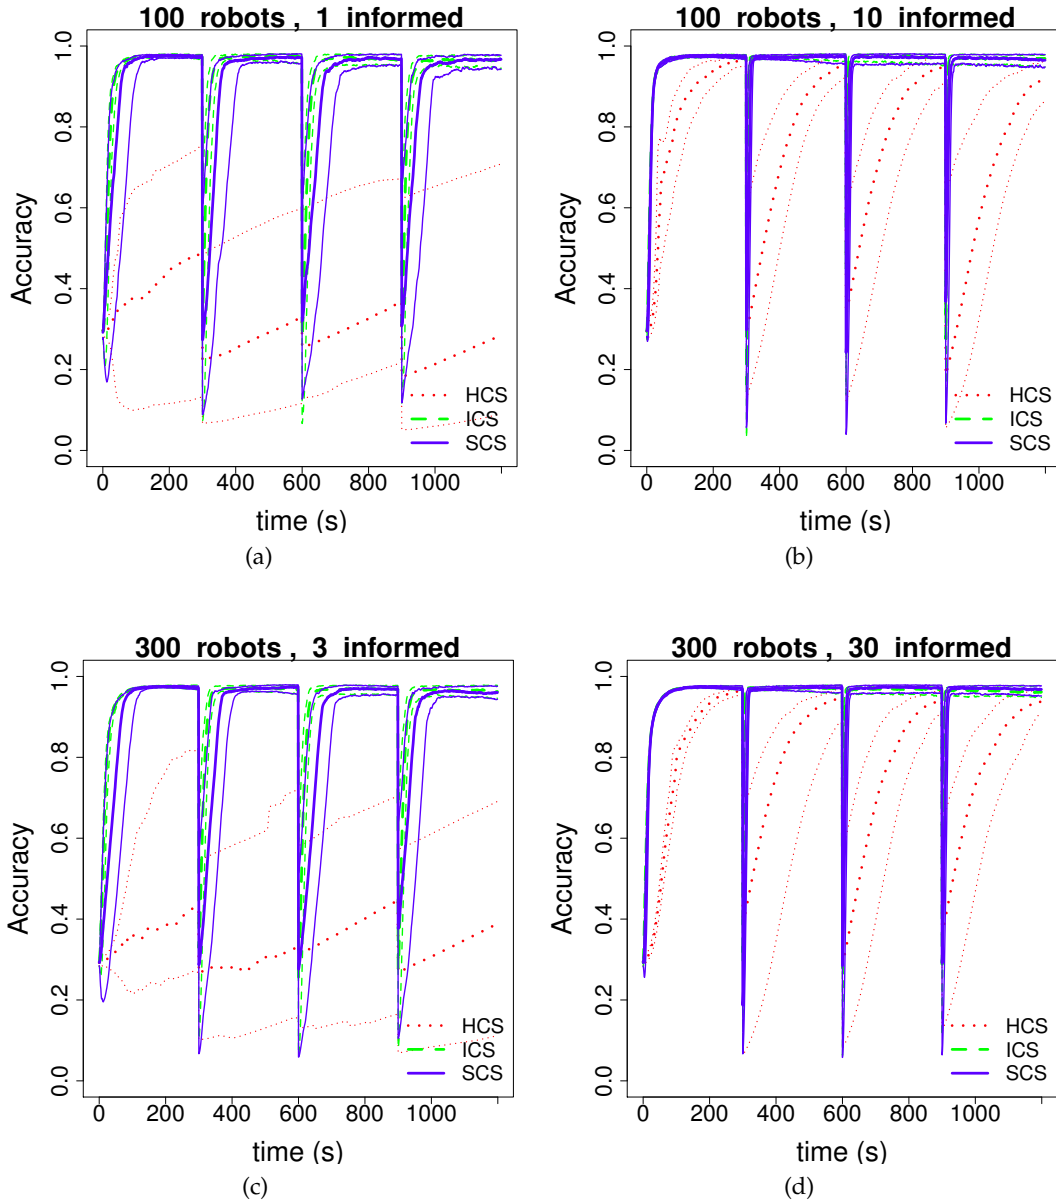


Figure 4.6: Results in simulation. HCS, ICS and SCS in the one-goal non-stationary environment using the non-spatial selection mechanism: effect on the accuracy. Ticker (central) lines represent the medians of the distributions, whereas thinner lines represent the 25% and the 75% percentiles.

in this case, informed robots are at the boundaries instead of being at random positions. Hence, the propagation of the goal direction in the swarm takes a bit longer. When 10% of the robots are informed (Figure 4.3b versus Figure 4.4b and Figure 4.3d against Figure 4.4d), results with the spatial selection mechanism show a minor dif-

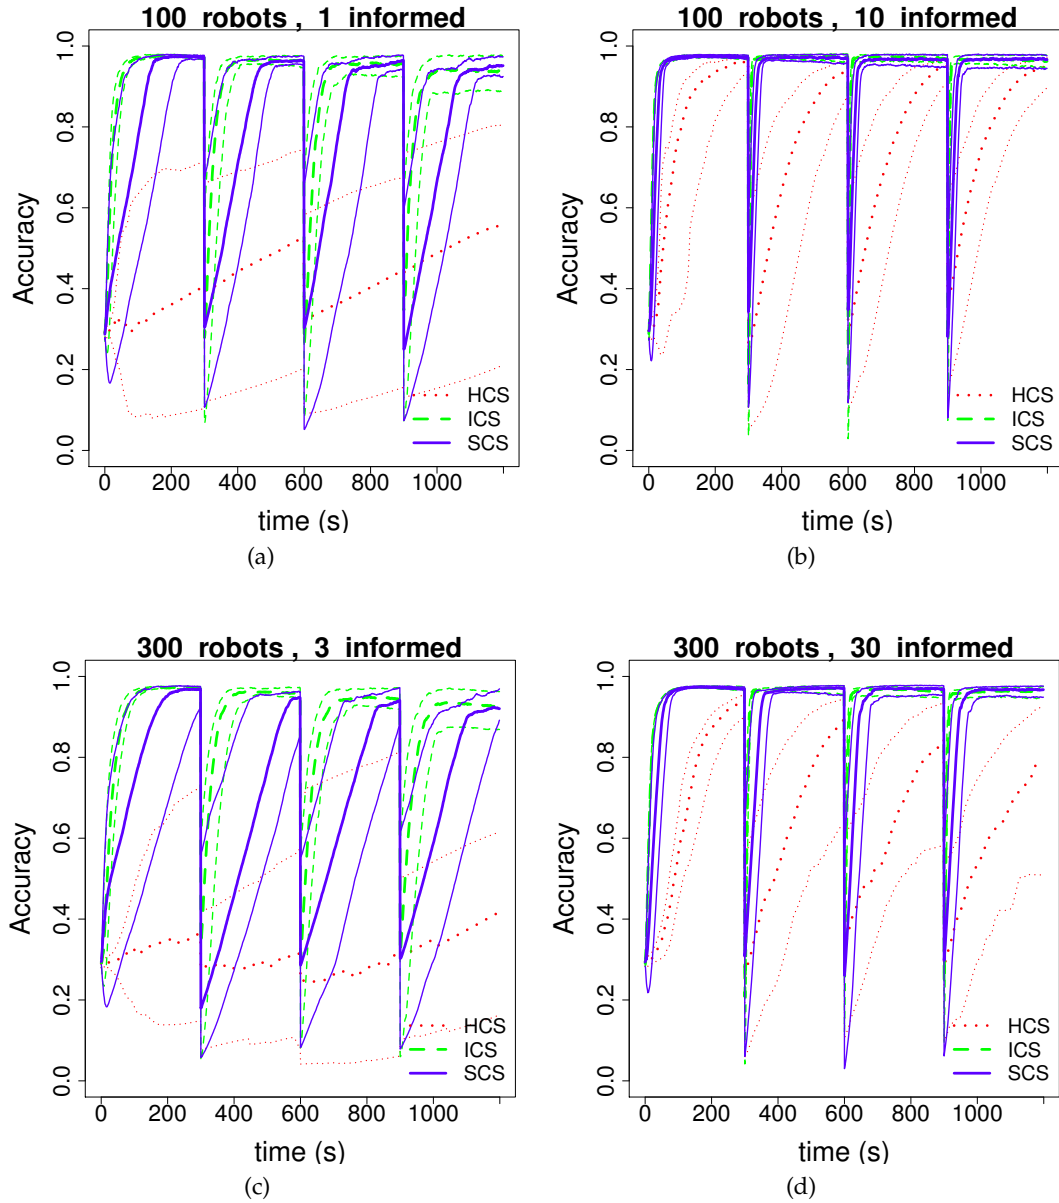


Figure 4.7: Results in simulation. HCS, ICS and SCS in the one-goal non-stationary environment using the spatial selection mechanism: effect on the accuracy. Ticker (central) lines represent the medians of the distributions, whereas thinner lines represent the 25% and the 75% percentiles.

ference in performance for the two selection mechanisms. In the supplementary material page (Ferrante et al., 2011), we report also the time evolution of the order metric and the distribution of the number of groups at the end of the experiment, showing that the swarm remains always cohesive.

When 10% of the robots are informed, ICS and SCS have very similar performance. In all cases, HCS is outperformed by the two strategies, that is still consistent with the results in (Ferrante et al., 2010b).

4.4.2 Real robots

Figure 4.5 reports all the results obtained in the real robot experiments. Figure 4.5a shows that results obtained in the stationary environment are similar to those obtained in simulation (Figure 4.3 and Figure 4.4). Both ICS and SCS perform very well (with a maximum average difference < 0.075 with a confidence level of 95%), whereas HCS is not able to reach reasonable levels of accuracy in the same amount of time, that is, 100 seconds (p -value < 0.01).

4.5 One non-stationary goal direction

4.5.1 Simulation

Figure 4.6 and Figure 4.7 show the results obtained in non-stationary environments when using a non-spatial selection and spatial selection mechanisms, respectively. These results show two points. First, within each stationary phase, the results are all consistent with the results obtained in the stationary environment case. Second, we find that all strategies exhibit, to some extent, some degree of adaptation to the changes in the goal direction. In all the cases, the ranking of the three strategies is the same. SCS performance is always comparable to ICS performance, although slightly lower. On the other hand, SCS is either better than HCS when the proportion of robots is 10% (Figure 4.6b, Figure 4.7b, Figure 4.6d and Figure 4.7d) or much better when there is only 1% informed robots (Figure 4.6a, Figure 4.7a, Figure 4.6c and Figure 4.7c). In the supplementary material page (Ferrante et al., 2011) we report also the time evolution of the order metric and the distribution of the number of groups at the end of the experiment, showing that, also in this case, the swarm remains always cohesive.

4.5.2 Real robots

Results of experiments in one-goal stationary environment (Figure 4.5b) confirm the trend observed in simulation: during both phases, ICS and SCS perform considerably well whereas, with HCS, the informed robots (in this case one) are not able to lead the swarm along the desired direction (p -value < 0.01).

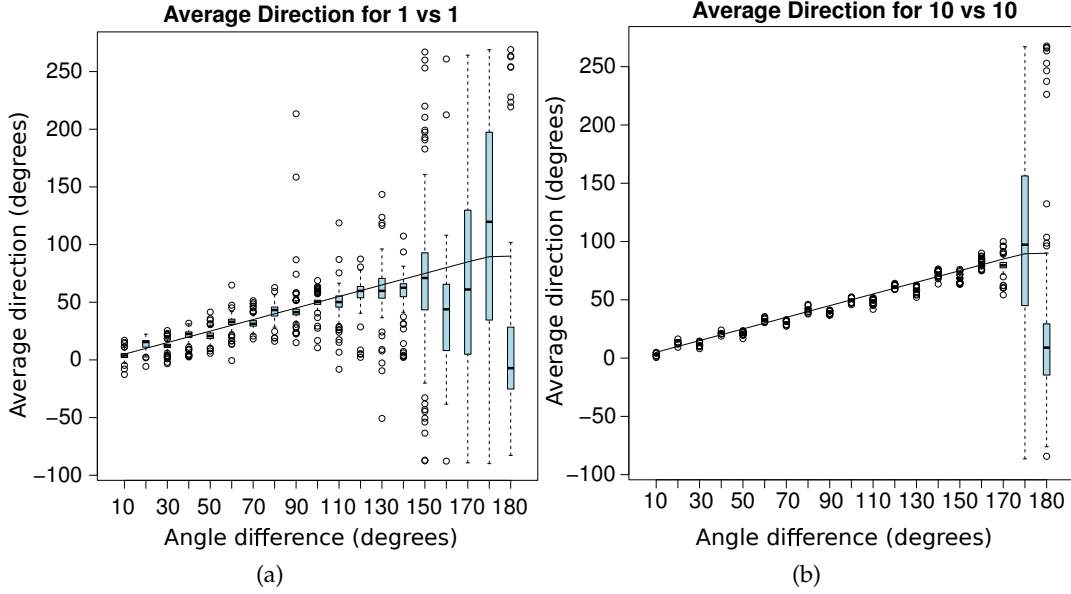


Figure 4.8: Distribution of average direction for $N_A = N_B$. Solid black line represents the theoretical average direction in total informed robots. The above and nether edges of the box indicate first and third quartiles. The black center line indicates the median for each dataset.

4.6 Two stationary goal directions

In the experiments with two stationary goal directions, we only studied HCS and we always observed a cohesive swarm, no matter the value of N_A , N_B and $|\theta_2 - \theta_1|$ (results not shown).

We now report and discuss the average swarm direction in the three cases: $N_A = N_B$ (no difference), $N_B - N_A = 2$ (small difference) and $N_B - N_A > 2$ (large difference).

No difference ($N_A = N_B$) Figure 4.8 shows that, in most of the cases, the average direction strictly follows the theoretical average direction $\hat{\theta}$. The most noticeable exception is the $\rho_{I1} = \rho_{I2} = 0.01$ case (Figure 4.8a). In this case, the robots are not able to follow the theoretical average direction when $\angle g_2 - \angle g_1$ is too high, and the distribution of the average swarm direction has high standard deviation. This can be explained by the fact that the total proportion of informed robots is not high enough to drive the swarm in the desired direction, as argued in [Çelikkanat and Şahin \(2010\)](#). In fact, Figure 4.8b shows that, with higher total proportion of informed robots, the swarm follows the theoretical average direction more precisely for most configurations of $\angle g_2 - \angle g_1$.

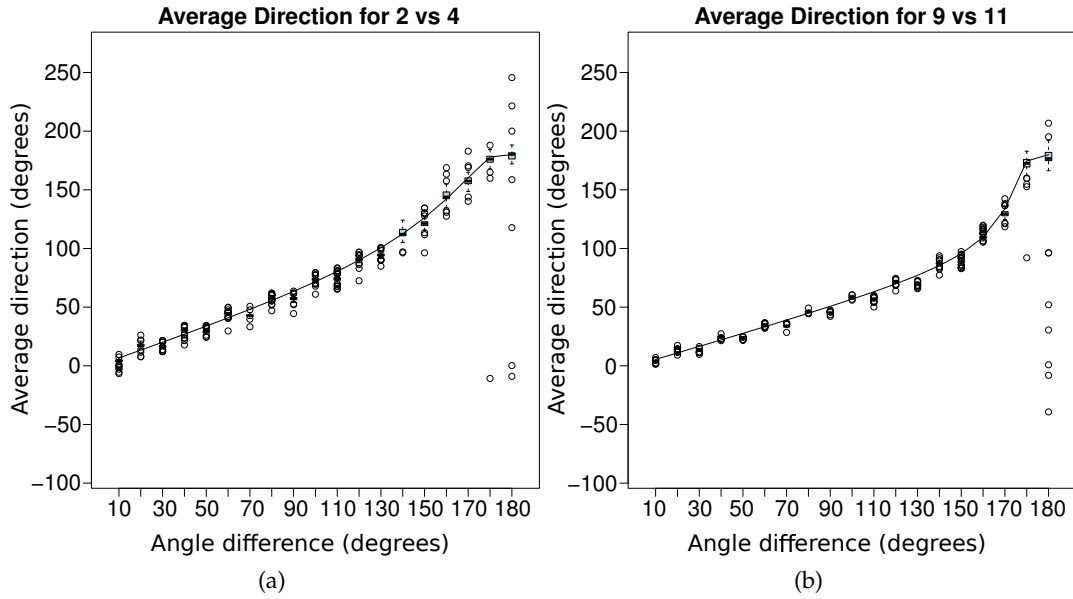


Figure 4.9: Distribution of average direction $N_B - N_A = 2$. Solid black line represents the theoretical average direction in total informed robots. The above and nether edges of the box indicate first and third quartiles. The black center line indicates the median for each dataset.

Small Difference ($N_B - N_A = 2$) Figure 4.9 shows that, no matter $\angle g_2 - \angle g_1$, the group follows the theoretical average direction. When $\angle g_2 - \angle g_1$ is high, larger total proportion of informed robots (Figure 4.9b) correspond to lower spread in the distribution. $\angle g_2 - \angle g_1 = 180$ degrees is a special case as it presents many outliers. Otherwise, the swarm is able to follow the theoretical average even for $\angle g_2 - \angle g_1 = 179$ degrees.

Large Difference ($N_B - N_A > 2$) Figure 4.10 shows that the theoretical average direction represents the goal direction that is known by the majority. Here, a precise following of the theoretical average direction always takes place, even when $\angle g_2 - \angle g_1 = 180$. Additionally, when the total proportion is small, there are a few outliers (Figure 4.10a). These outliers are otherwise not present for higher total proportion of informed (Figure 4.10b).




Figure 4.10: Distribution of average direction for $N_B - N_A > 2$. Solid black line represents the theoretical average direction in total informed robots. The above and nether edges of the box indicate first and third quartiles. The black center line indicates the median for each dataset.

4.7 Two non-stationary goal directions

4.7.1 Simulation

In this setting, we report not only the accuracy over time for the non-spatial (Figure 4.11) and spatial (Figure 4.12) selection mechanisms, but also the data regarding the number of groups present at the end of the experiment (Figure 4.13 and Figure 4.14). Figure 4.11 shows the results obtained in two-goal non-stationary environments when using a non-spatial selection mechanism. We first focus on the results for the 1% informed robots case (Figure 4.11a and Figure 4.11c). In the first phase, between time 0 and T_s , we observe similar results as those observed in stationary environments. Subsequently, during the two-goal phase, all strategies are able to track goal direction B . This can be explained by our choice of parameters, $\rho_{I2} = 0.1 > \rho_{I1} = 0.01$, which makes robots informed about goal direction B easily able to lead the swarm in that direction as only one robot is opposing them. After time $T_s + T_{\Delta p}$, we observe that HCS continues tracking goal direction B , whereas ICS and SCS are able to follow again goal direction A . In Figure 4.13a and Figure 4.13c, we observe that the swarm splits only when using ICS. These results show that both ICS and SCS are preferable to HCS in terms of accuracy, because they are both able

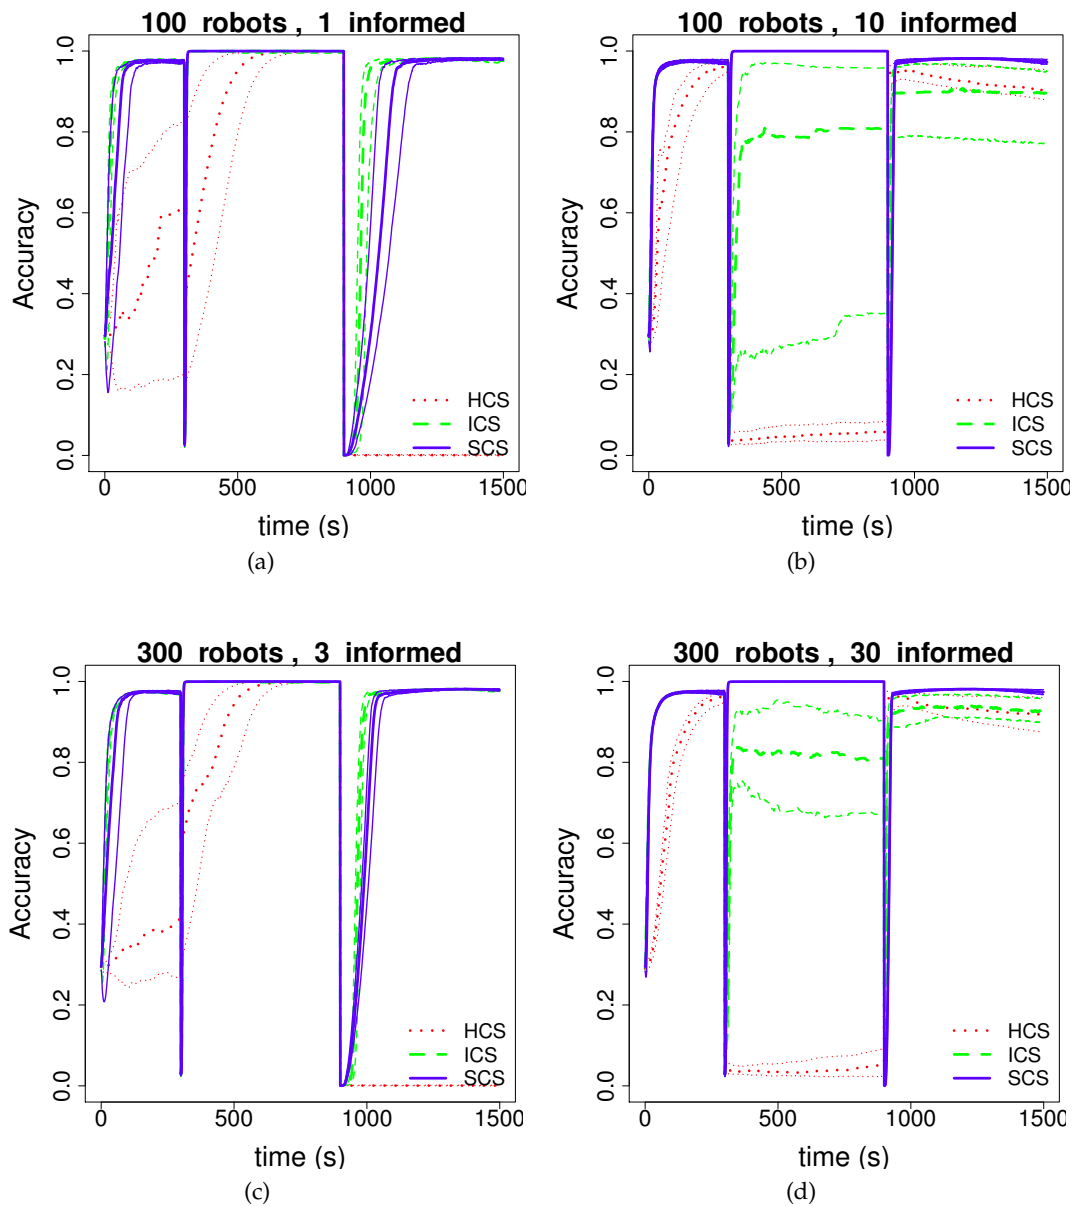


Figure 4.11: Results in simulation. HCS, ICS and SCS in the two-goal non-stationary environment using the non-spatial selection mechanism: effect on the accuracy. Ticker (central) lines represent the medians of the distributions, whereas thinner lines represent the 25% and the 75% percentiles.

to track the goal directions (first A , then B , then A again). However, SCS is better than ICS because it keeps swarm cohesion all the times whereas ICS does not.

When the proportion of informed robots is set to 10%, results are slightly different. In fact, HCS is not able to track goal direction B . This is due to the fact that,

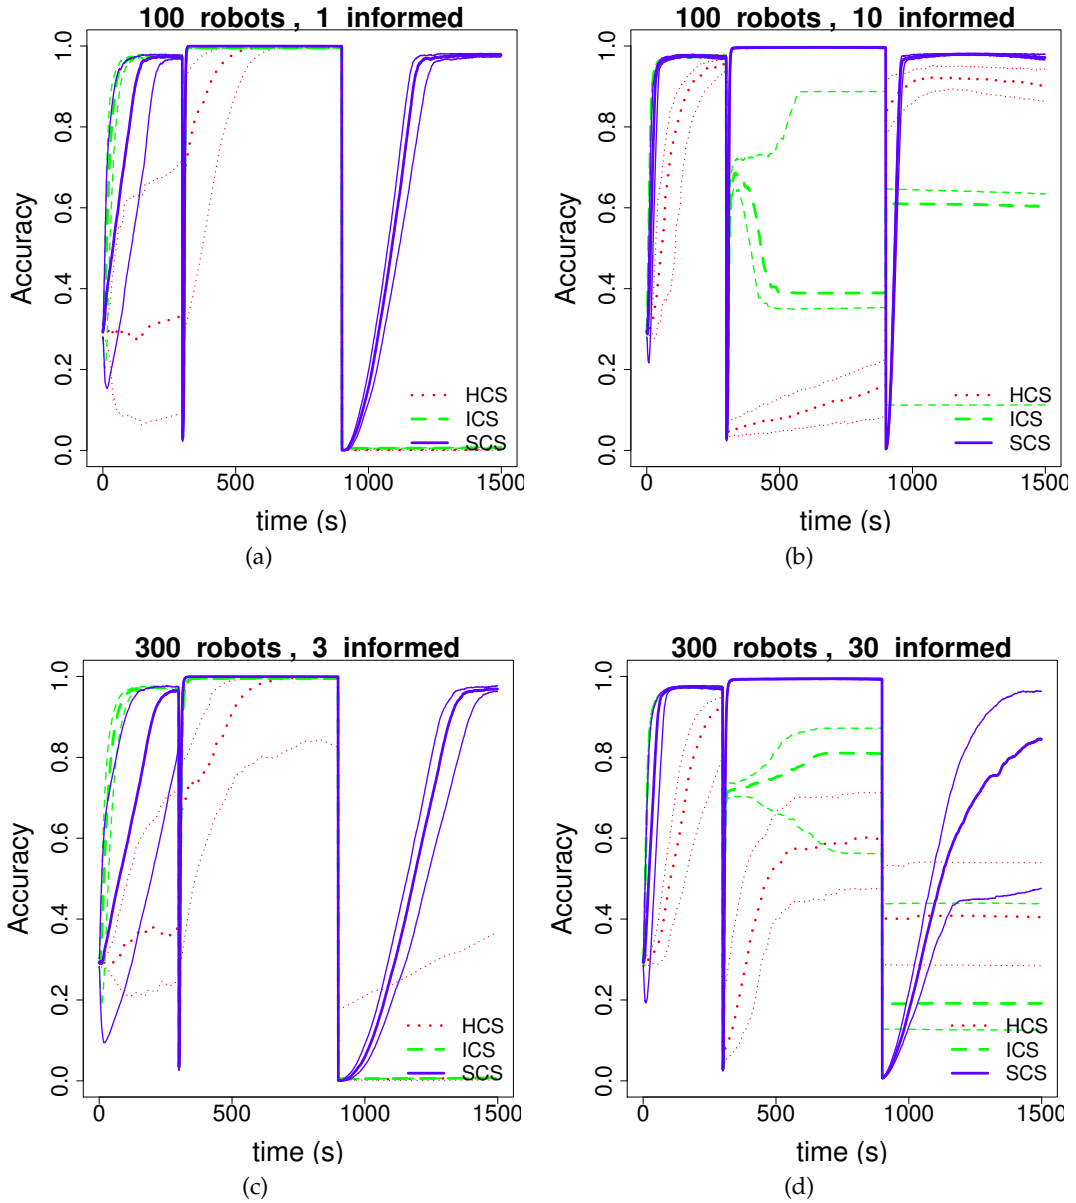


Figure 4.12: Results in simulation. HCS, ICS and SCS in the two-goal non-stationary environment using the spatial selection mechanism: effect on the accuracy. Ticker (central) lines represent the medians of the distributions, whereas thinner lines represent the 25% and the 75% percentiles.

when $\rho_{I1} = \rho_{I2}$ and the swarm already achieved a consensus decision on goal direction A , the number of robots informed about goal direction B is not large enough to make the swarm change this consensus decision. However, the swarm almost never splits, as shown in Figure 4.13b and Figure 4.13d. Figure 4.13b and Figure 4.13d

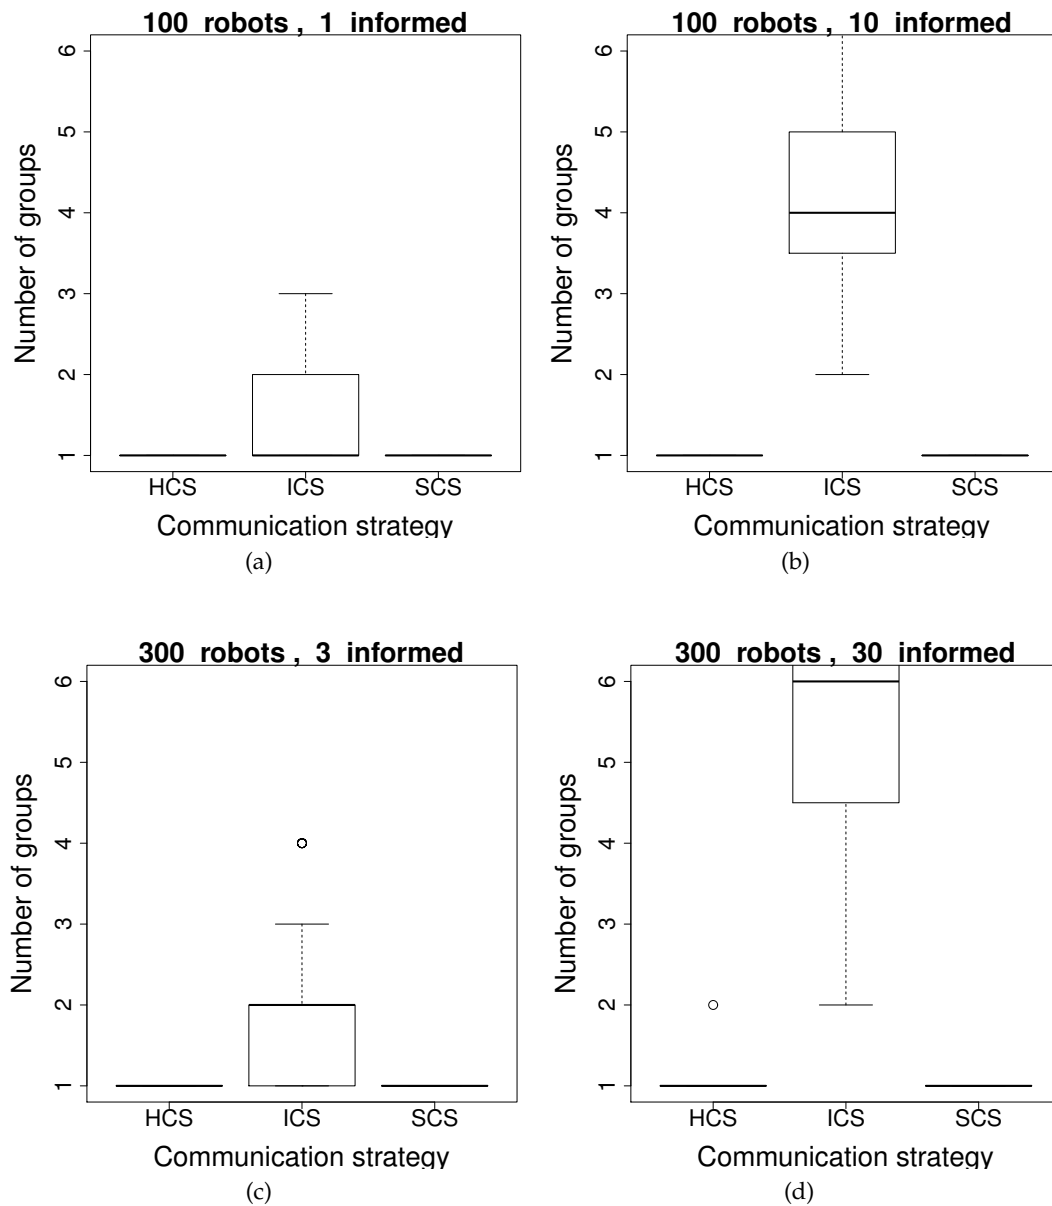


Figure 4.13: Results in simulation. HCS, ICS and SCS in the two-goal non-stationary environment using the non-spatial selection mechanism: number of groups at the end of the experiment.

show instead that the swarm does not keep cohesion when the strategy used is ICS. This translates into an intermediate level of accuracy during the two-goal phase (Figure 4.11b and Figure 4.11d), due to the fact that when the swarm splits, part of it tracks goal direction A and the other part tracks goal direction B . The relative sizes of these groups change from experiment to experiment, which is directly linked to

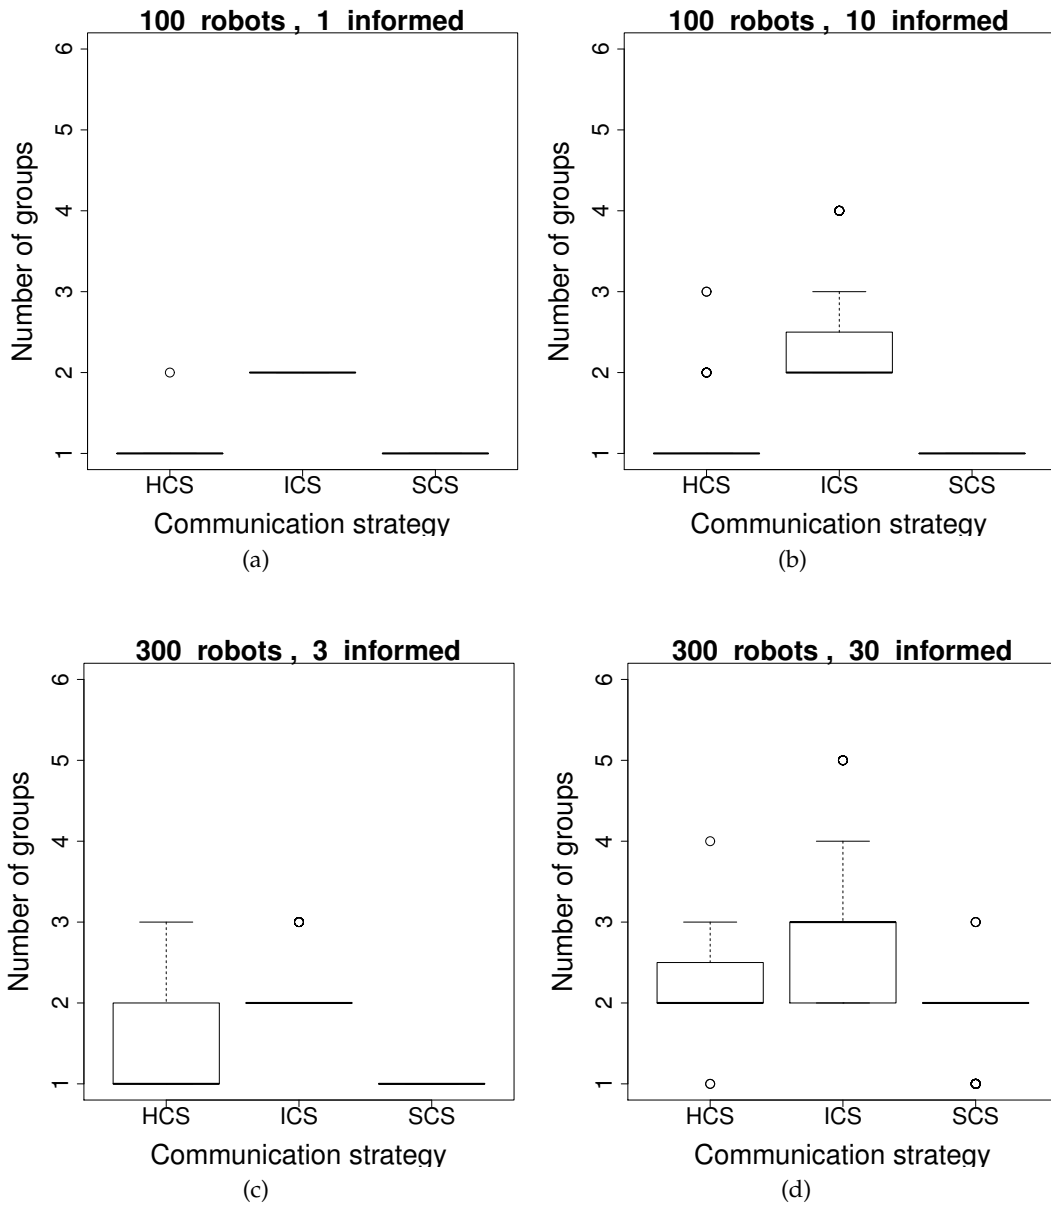


Figure 4.14: Results in simulation. HCS, ICS and SCS in the two-goal non-stationary environment using the spatial selection mechanism: number of groups at the end of the experiment.

the observed fluctuations around the median value during the two-goal phase of ICS. The best results in these experiments are produced by using SCS. In fact, the swarm is able to first track goal direction A , then track goal direction B and then again goal direction A and the swarm cohesion is always guaranteed, even in large swarms.

Figure 4.12 shows the results obtained in two-goal non-stationary environments

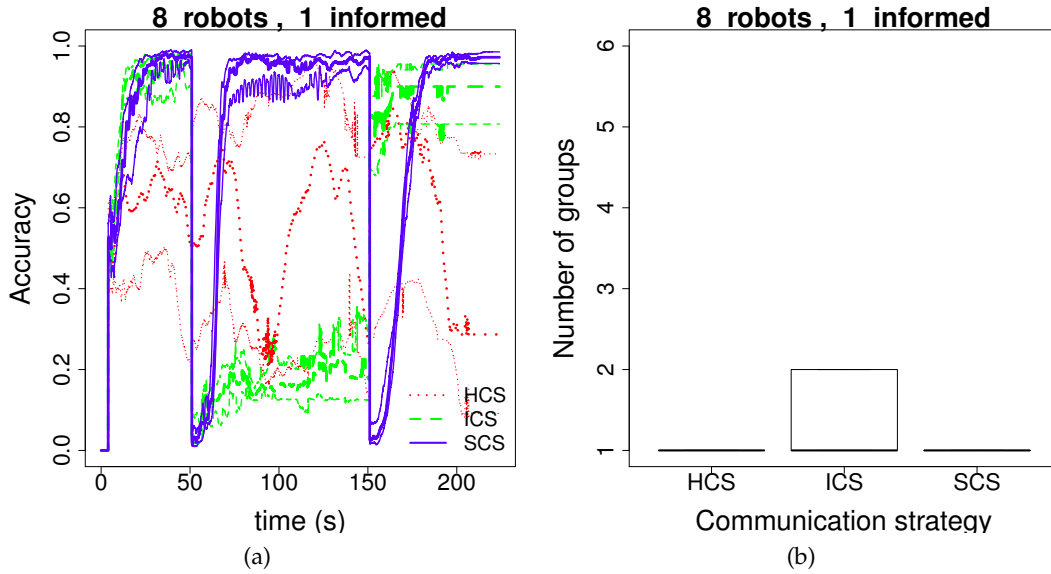


Figure 4.15: Results with real robots. HCS, ICS and SCS in two-goal non-stationary environment. (a) plots the distribution of the accuracy metric over time. Ticker (central) lines represent the medians of the distributions, whereas thinner lines represent the 25% and the 75% percentiles. Figure (b) plots the distribution of the number of groups at the end of the experiments in two-goal non-stationary environment.

when using a spatial selection mechanism. When we first focus on the experiments with only 1% of informed robots (Figure 4.12a and Figure 4.12c), results show that SCS outperforms the other two strategies, as it is the only strategy able to track changes in goal direction (A to B and back to A). HCS behaves as in the non-spatial selection mechanism. Conversely, ICS performs dramatically worse in this case, as the swarm always splits during the two-goal phase (Figure 4.14a and Figure 4.14c), which is due to the fact that informed robots are always selected along the periphery of the swarm. In this case, the swarm can no longer track the goal direction A , as robots informed about goal direction A disconnected from the rest of the swarm during the two-goal phase. Results with 100 robots and 10% informed (Figure 4.14b) are similar to the ones reported, in the analogous case, for the non-spatial selection mechanism. However, with 300 robots, we observe that swarm cohesion is not guaranteed anymore, even when using SCS (Figure 4.14d). This case is in fact the most challenging one, and we included it to show the limits of our method. A large number of robots placed along the periphery is stretching the swarm in two different directions, eventually causing it to split. As a result, the accuracy metric is also affected (Figure 4.12d). This case is unlikely in practice, as in a real application information would be either randomly distributed in the swarm (with robots having

heterogeneous sensors) or possessed by the robots that are leading the swarm in front.

Figure 4.13 and Figure 4.14 show that the number of groups obtained when using ICS differs between the spatial and the non-spatial selection cases. In the non-spatial selection cases more subgroups are formed compared to the spatial selection case. This can be explained by the following argument: when using the non-spatial selection mechanism, several subgroups emerge and split from the main group at different moments of the experiment due to the presence of non-uniform “cluster of informed robots”; when using the spatial selection mechanism, instead, informed robots are spatially distributed in one unique cluster, so that the number of emerging subgroups is smaller and closer to two. For the time evolution of the order metric and for the distribution of group sizes for the first two environment refer to the supplementary materials page (Ferrante et al., 2011).

4.7.2 Real robots

Figure 4.15a shows the results obtained in the two-goal non-stationary environment. HCS performs poorly during all the duration of the experiment since informed robots are never able to stabilize the swarm along one direction. This might be due to the limited time available for real robot experiments, or to the different nature of noise which prevents the control of the direction of the swarm without an effective communication strategy. However, the swarm is aligned along the same direction as the order metric is high (see the supplementary materials page Ferrante et al. (2011)). Using ICS and SCS instead introduces a degree of control on the direction of the swarm. During the first phase (between time 0 and T_s), the results are consistent with the results in the stationary environment case: ICS and SCS have both good performance, that is, they both track goal direction A , compared to HCS (p -value < 0.01).

Figure 4.15a also shows that SCS has very good results, comparable to the ones obtained in simulation, also during the subsequent phases, as it first tracks goal direction A , then goal direction B and finally goal direction A . When using ICS, instead, the swarm continues tracking goal direction A during the two-goal phases in 70% of the runs (7 out of 10), in which the swarm does not split (Figure 4.15b). However, in the remaining runs (3 out of 10), the swarm splits in two or more groups: one group follows goal direction B , whereas the other group continues following goal direction A . This causes the accuracy metric to have the distribution depicted in Figure 4.15a, that shows median values close to 0.8 and an high spread.

We performed the Wilcoxon rank sum test to compare the medians of the time-

average performance of SCS against HCS and of SCS against ICS during all three phases. The test suggested that SCS consistently outperforms HCS during all three phases (p -value < 0.01), outperforms ICS during the second phase (p -value < 0.01) and performs comparatively as ICS in the first and third phase, as described above.

4.8 Summary and Discussion

In this chapter, we performed experiments with explicit information transfer. In doing so, we employed the communication strategies presented in Section 3.4. The new communication strategies are loosely inspired by the one introduced by [Campo et al. \(2006\)](#) in the context of collective transport. His strategy can be considered a predecessor of ICS: all robots were informed about a noisy goal direction at the beginning of the simulation and kept an estimate of the actual goal direction, which was updated by averaging the goal direction estimated by their neighbors. The main difference between our and their work is twofold: first, they consider a different problem (estimating the actual goal direction by averaging out noise); second, in their work every robot estimates the goal direction and signals its own estimate, while in ICS uninformed robots do not compute an explicit estimate of the goal direction, whereas informed robots do not update the information they send but always send the goal direction.

In Table 4.3 we report a complete overview of the results as a table of experiments. In stationary and non-stationary environments, we showed that SCS and ICS, the two explicit information transfer methods presented in this dissertation, can tackle the problem of following one stationary and non-stationary goal direction successfully. This means that, using ICS or SCS, flocking along a goal direction is possible with a high level of accuracy even if only a few robots are informed about the goal direction and when the desired goal direction changes over time.

In the two-goal stationary environment, we showed that HCS alone is capable of making the swarm follow the average direction between the two without splitting. This result presents some differences when compared to the results of [Couzin et al. \(2005\)](#), where they also studied the conflicting goal direction case but showed that the resulting average direction strongly depends on the difference between the two goal directions. This lack of agreement might be due either to the different methodology or to the different level of detail in the simulations.

Results obtained in two-goal non-stationary environments reveal the benefits of using the most advanced of the explicit information transfer strategies studied in this dissertation: SCS. In fact, SCS provides swarm cohesion in almost all cases without

N	ρ_{I_1}	Environment	Selection	Best strategies	Split
100	0.01	1-goal stat.	non-spatial	ICS, SCS	none
100	0.1	1-goal stat.	non-spatial	ICS, SCS	none
300	0.01	1-goal stat.	non-spatial	ICS, SCS	none
300	0.1	1-goal stat.	non-spatial	ICS, SCS	none
100	0.01	1-goal stat.	spatial	ICS	none
100	0.1	1-goal stat.	spatial	ICS, SCS	none
300	0.01	1-goal stat.	spatial	ICS	none
300	0.1	1-goal stat.	spatial	ICS, SCS	none
8	0.125	RR 1-goal stat.	random	ICS, SCS	none
100	0.01	1-goal non-stat.	non-spatial	ICS, SCS	none
100	0.1	1-goal non-stat.	non-spatial	ICS, SCS	none
300	0.01	1-goal non-stat.	non-spatial	ICS, SCS	none
300	0.1	1-goal non-stat.	non-spatial	ICS, SCS	none
100	0.01	1-goal non-stat.	spatial	ICS	none
100	0.1	1-goal non-stat.	spatial	ICS, SCS	none
300	0.01	1-goal non-stat.	spatial	ICS	none
300	0.1	1-goal non-stat.	spatial	ICS, SCS	none
8	0.125	RR 1-goal stat.	random	ICS, SCS	none
Difference	$\rho_{I_1} + \rho_{I_2}$	Environment	Selection	Tracked average	Split
no	0.02	2-goal stat.	random	yes	none
no	0.2	2-goal stat.	random	yes	none
small	0.06	2-goal stat.	random	yes	none
small	0.2	2-goal stat.	random	yes	none
large	0.06	2-goal stat.	random	yes	none
large	0.2	2-goal stat.	random	yes	none
N	ρ_{I_1}	Environment	Selection	Best strategies	Split
100	0.01	2-goal non-stat.	non-spatial	SCS	ICS
100	0.1	2-goal non-stat.	non-spatial	SCS	ICS
300	0.01	2-goal non-stat.	non-spatial	SCS	ICS
300	0.1	2-goal non-stat.	non-spatial	SCS	HCS, ICS
100	0.01	2-goal non-stat.	spatial	SCS	HCS, ICS
100	0.1	2-goal non-stat.	spatial	SCS	HCS, ICS
300	0.01	2-goal non-stat.	spatial	SCS	HCS, ICS
300	0.1	2-goal non-stat.	spatial	none	HCS, ICS, SCS
8	0.125	RR 2-goal non-stat.	random	SCS	ICS

Table 4.3: Table of experiments summarizing all the results of this chapter. The order of the experiments is the same as they appear in the text. For experiments in one-goal (stationary and non-stationary) environments and in two-goal non-stationary environments, we report what are the best performing strategies and with which strategies the swarm splits (last two columns) for each setting (first four column, RR in some rows of the third column stands for “real robots”). In experiments in two-goal stationary environments, we report whether the swarm tracked the average direction and whether it has split, and we recall that here we used always HCS with a swarm of 100 robots.

sacrificing accuracy. On the other hand, ICS is very strong in providing high level of accuracy but performs dramatically worse in maintaining swarm cohesion. This general message holds for both experiments executed in simulation and with the real robots.

Chapter 5

Explicit information transfer in collective transport

In this chapter, we present an application of explicit information transfer to collective transport. The goal is to control a group of robots that are attached to an object and that need to move in a coordinated fashion while carrying the object to a required location. The environment is cluttered by obstacles: the robots use explicit information transfer to share directional information and to integrate the direction to the goal location with the obstacle avoidance direction.

The chapter is organized as follows. In Section 5.1, we introduce the collective transport problem, some related literature and our approach. In Section 5.2, we describe the collective transport method. In Section 5.3, we describe the experimental setup. In Section 5.4, we present experimental results, and in Section 5.5 we conclude with a summary and a discussion.

5.1 Problem description

In collective transport, a group of robots cooperates in order to transport an object that cannot be transported by a single robot due to several reasons such as the weight or the size of the object. Collective transport within swarm robotics has received substantial attention during the last decade. [Campo et al. \(2006\)](#) investigated the use of a goal negotiation strategy similar to HCS presented in this dissertation (see Section 3.4 and Section 4.8 for a description and discussion). The strategy was used to achieve collective transport in a goal direction without obstacles. Other authors used artificial evolution to tackle the collective transport problem ([Groß and M.Dorigo, 2009](#); [Trianni and Dorigo, 2006](#); [Baldassarre et al., 2006](#)). The latter two articles studied a collective transport problem similar to the one considered in this

chapter. In [Trianni and Dorigo \(2006\)](#), robots are physically connected to each other and must navigate in an environment with holes to be avoided and without a goal direction. Similarly, in [Baldassarre et al. \(2006\)](#) physically connected robots must navigate in an environment with obstacles, furrows and holes towards a light source (goal direction). Through artificial evolution, both authors synthesized a behavior that heavily exploits the traction sensor, a specialized sensor that is able to detect forces exerted among the connected robots. The use of the traction sensor can be considered an implicit information transfer mechanism, since it allows robots to implicitly share information on their orientation in a similar way as done in this dissertation for flocking (see Chapter 6).

In this chapter, we propose a control strategy that can be considered an application of ICS (see Section 3.4.2 and Section 3.4.3) to collectively transport an object to a goal location and in presence of obstacles. This chapter's contribution to this dissertation is two fold. First, it strengthens the value of the proposed explicit information transfer mechanism as we show that it can also be applied to another collective behavior. Second, together with existing work in collective transport ([Baldassarre et al., 2006](#)), it contributes to the information transfer framework used in this dissertation: we use explicit information transfer via local communication while [Baldassarre et al. \(2006\)](#) used implicit information transfer via traction sensor.

5.2 Explicit information transfer and collective transport

As we did in Chapter 4, we consider two conflicting goal directions: the first, denoted by $\angle g_1$, is the direction to the goal location; the second, denoted by $\angle g_2$, is the direction to an obstacle to be avoided. As in Chapter 3, we explain the method from the point of view of a single robot which we call the focal robot. At a given moment, the focal robot might have access to both $\angle g_1$ and $\angle g_2$, to only one of them or to none of them. This is because obstacles might not be always in the range of sight of the focal robot ($\angle g_2$ not available) or because the goal direction is not perceived by the focal robot due to occlusions by obstacles ($\angle g_1$ not available).

We designed a hierarchically organized control method, in which the *Collective transport* component is decomposed into three components that we call *goal detector*, *obstacle detector* and *social mediation*.

Social mediation, explained in Section 5.2.1, is the explicit information transfer method used to negotiate the navigation direction used in *collective transport*, in turns explained in Section 5.2.2. The two low level control components *goal detector* and *obstacle detector* are used as follows. *Goal detector* is used to obtain the goal direction

Notation	Meaning	Component
$\angle \mathbf{g}_1$	Goal direction	Collective transport
$\angle \mathbf{g}_2$	Obstacle direction	Collective transport
$\angle \mathbf{g}_P$	Preferred direction that takes into account $\angle \mathbf{g}_1$ and $\angle \mathbf{g}_2$ (if present)	Social mediation
$\angle \mathbf{a}$	Socially mediated angle $\angle \mathbf{a} = \angle \sum_{i \in \mathcal{N}_{CT}} e^{j\theta_{s_i}}$	Social mediation, collective transport
$\angle \bar{\mathbf{a}}$	Weighted time average of $\angle \mathbf{a}$	Collective transport
θ_0	Robot orientation	Collective transport
θ_{s_0}	Direction sent by social mediation: $\angle \mathbf{a}$ when <i>amSocial</i> is <i>true</i> or $\angle \mathbf{g}_P$ when it is <i>false</i>	Social mediation
$\theta_{s_1} \dots \theta_{s_{ \mathcal{N}_{CT} }}$	Direction received from the $ \mathcal{N}_{CT} $ neighbors	Social mediation

Table 5.1: Explanation of the notation used to describe the control method.

$\angle \mathbf{g}_1$; *obstacle detector* is used to detect the presence of obstacles and to obtain the angle $\angle \mathbf{g}_2$ of the closest obstacle when detected.

In the following, we use a notation that is as much as possible consistent with the one used in the other chapters. The notation is explained and summarized in Table 5.1.

5.2.1 Social mediation

Social mediation is an explicit information transfer strategy that consists in the adaptation of ICS (see Section 3.4.2) to collective transport. It needs two pieces of directional information: $\angle \mathbf{a}$ and $\angle \mathbf{g}_P$. $\angle \mathbf{a}$ is the socially mediated direction, which is computed in a similar way as the angle of the alignment control vector (see Section 3.4). $\angle \mathbf{g}_P$ is the robot's preferred direction, and needs to take into account both $\angle \mathbf{g}_1$ and $\angle \mathbf{g}_2$ when present, as explained in Section 5.2.2. The main idea behind the algorithm is the following: When the focal robot is non informed (i.e., it does not have any information on the goal direction \mathbf{g}_1 or on the obstacles direction \mathbf{g}_2), it has an internal flag *amSocial* set to *true*. In this case, the focal robot computes $\angle \mathbf{a}$, the average of the directional information available to its neighbors, and sends this value to the other robots. However, when the focal robot is informed, the *amSocial* flag is set to *false*. In this case, it sends its own preferred direction $\angle \mathbf{g}_P$, which takes into account the goal directions. The *amSocial* flag of this component, as well as the preferred direction $\angle \mathbf{g}_P$, are updated by collective transport, as explained in Section 5.2.2. Thus, when all other robots are non informed and sending $\angle \mathbf{a}$, the opinion of the focal robot diffuses in the group, that is $\angle \mathbf{a}$ through the group converges to $\angle \mathbf{g}_P$.

Algorithm 1 explains the steps executed at every control step: The focal robot receives the directional piece of information $\theta_{s_1} \dots \theta_{s_{|\mathcal{N}_{CT}|}}$ of its neighbors, where \mathcal{N}_{CT} is the set containing the neighbors of the focal robot. The socially mediated direction $\angle \mathbf{a}$ is computed by averaging the directional information received by the neighbors (line 2), with the robot's

own information θ_{s_0} . The directional information θ_{s_0} is then determined (lines 3-7) and sent (line 8).

Algorithm 1 Social mediation control loop

```

1: Receive( $\theta_{s_1} \dots \theta_{s_{|N_{CT}|}}$ )
2:  $\angle \mathbf{a} \leftarrow \angle \sum_{i \in N_{CT}} e^{j\theta_{s_i}}$ 
3: if amSocial then
4:    $\theta_{s_0} \leftarrow \angle \mathbf{a}$ 
5: else
6:    $\theta_{s_0} \leftarrow \angle \mathbf{g}_P$ 
7: end if
8: Send( $\theta_{s_0}$ )

```

We now describe the main control method, that uses social mediation and the other two components to achieve collective transport with obstacle avoidance.

5.2.2 Collective transport

In this section we present the main control method for collective transport with obstacle avoidance. This main component uses the other three components explained at the beginning of this section. We explain the method in Algorithm 2, where *Component* :: denotes an interaction with sub-component *Component* :: (*SocialMediation* ::, *GoalDetector* :: or *ObstacleDetector* ::). Additionally, here directional information is always considered as relative to the focal robot orientation, denoted with θ_0 , obtained via the common light source.

At the beginning of Algorithm 2, sensors are used to verify whether the goal and/or obstacles are perceived (lines 1-2). The corresponding directions $\angle \mathbf{g}_1$, corresponding to the goal direction, and $\angle \mathbf{g}_2$, corresponding to the angle of the closest obstacle, are also obtained.

Depending to which information is available to the focal robot, the value of the *amSocial* flag is determined (lines 3-7): If the focal robot perceives an obstacle or the goal direction, *amSocial* is set to *false*; otherwise, if neither the goal direction nor an obstacle are perceived, *amSocial* is set to *true*.

When only the goal direction is perceived, the focal robot simply informs the other robots about the goal direction by setting its preferred direction $\angle \mathbf{g}_P$ to the goal direction $\angle \mathbf{g}_1$ (line 9). When only the obstacle is perceived, the focal robot avoids the obstacle and computes the preferred direction as $\angle \mathbf{g}_P = \angle \mathbf{g}_2'$ (opposite direction of $\angle \mathbf{g}_2$) by setting a weight factor $w_{CT} = 1$ (line 15), used to compute a weighted average (see below).

When, however, both the obstacle and the goal direction are perceived, the focal robot needs to integrate these two pieces of information and to compute the preferred direction $\angle \mathbf{g}_P$ accordingly: $\angle \mathbf{g}_1$ and $\angle \mathbf{g}_2$ are thus averaged using a weighted average and the result is assigned to $\angle \mathbf{g}_P$ (lines 17). The weighted average uses a weight $w_{CT} \in [0, 1]$, which is a function of the distance between the robot and the obstacle (line 13) and represents

Algorithm 2 Collective transport control loop

```

1: [ $\angle \mathbf{g}_1, goalPerceived$ ] = GoalDetector :: PerceiveGoal()
2: [ $\angle \mathbf{g}_2, d_o, obstaclePerceived$ ] = ObstacleDetector :: PerceiveObstacle()
3: if goalPerceived or obstaclePerceived then
4:   SocialMediation :: amSocial  $\leftarrow$  false
5: else
6:   SocialMediation :: amSocial  $\leftarrow$  true
7: end if
8: if goalPerceived then
9:   SocialMediation ::  $\angle \mathbf{g}_P \leftarrow \angle \mathbf{g}_1$ 
10: end if
11: if obstaclePerceived then
12:   if goalPerceived then
13:      $w_{CT} \leftarrow -\frac{d_o}{\min(d_o, D_{CT})} + 1$ 
14:   else
15:      $w_{CT} \leftarrow 1$ 
16:   end if
17:   SocialMediation ::  $\angle \mathbf{g}_P \leftarrow \angle(w_{CT}\mathbf{g}'_2 + (1 - w_{CT})\mathbf{g}_1)$ 
18: end if
19: SocialMediation :: ControlStep()
20:  $\angle \bar{\mathbf{a}} \leftarrow \angle((1 - \varpi)\bar{\mathbf{a}} + \varpi\mathbf{a})$ 
21: MotionControl( $\bar{\theta}_S$ )

```

how urgent it is to avoid obstacles: $w_{CT} = 1$ when the obstacle is very close ($d_o = 0$) and $w_{CT} = 0$ when it is far away ($d_o = D_{CT}$, the maximal perception range of obstacles). We set $D_{CT} = 0.75$ meters, half of the maximal range of the distance scanner, and we use the min operator to avoid negative values for w_{CT} .

Once $\angle \mathbf{g}_P$ is computed, the control step of social mediation is executed (line 19). As a result, the angle $\angle \mathbf{a}$ is computed by social mediation. This angle is then filtered by computing a weighted time average (line 20) to filter out the effect of noise, with a factor $\varpi = 0.25$.

Finally, the motion control logic uses the filtered socially mediated direction $\bar{\mathbf{a}}$ as a reference direction to be followed (line 21). The focal robot first converts the socially mediated direction to its local frame of reference using its own orientation θ_0 . It then computes the left and right wheel speeds:

$$\begin{aligned}
 N_L &= u + \frac{\omega}{2}l, \\
 N_R &= u - \frac{\omega}{2}l, \\
 \omega &= K_{CT}\angle \bar{\mathbf{a}},
 \end{aligned}$$

where N_L, N_R are the wheel rotation speeds of the left and right wheel speed respectively, $l = 0.1$ m is the distance between the two wheels, u and ω are the forward and angular velocities respectively. The forward velocity $u = 7$ cm/s is kept constant, whereas we vary



Figure 5.1: Picture of robot setup: three mobile robots are attached to an irregularly shaped object, another robot, that can only be grasped only in certain regions.

the angular velocity ω proportionally to the socially mediated direction $\angle \bar{a}$ to be followed, where $K_{CT} = 3$ is a proportional factor (we assume a clockwise convention for the angles). This motion control method is equivalent to MIMC (presented in Section 3.5). Here, we use MIMC and not MDMC as it is the only one of the two that does not require full vectorial information but only directional information.

5.3 Experimental setup

A group of three simulated robots are attached to an irregularly shaped object (Figure 5.1). The task is to collectively transport the object from an initial to a goal location. The object is irregular in the sense that it cannot be grasped through its entire perimeter but only in certain regions.

We performed three sets of experiments. The first two sets consider a simple environment, where we position only one obstacle at the center of the arena with varying angle ϑ (see Figure 5.2a). For each set of experiments, we executed 100 runs. Our hypothesis is that the more ϑ tends to 0, the longer it takes to avoid the obstacle in collective transport. We also expect (and confirm) that the proposed control method is robust enough to always accomplish the task (move from an initial to a goal location, see Figure 5.2b) in this simple environment. We hence report the completion times as a function of ϑ . The difference between the first and the second set of experiments is that in the first set we just analyze the impact of the angle ϑ by keeping the projected size of the obstacle l_{obs_x} fixed (Figure 5.2a),

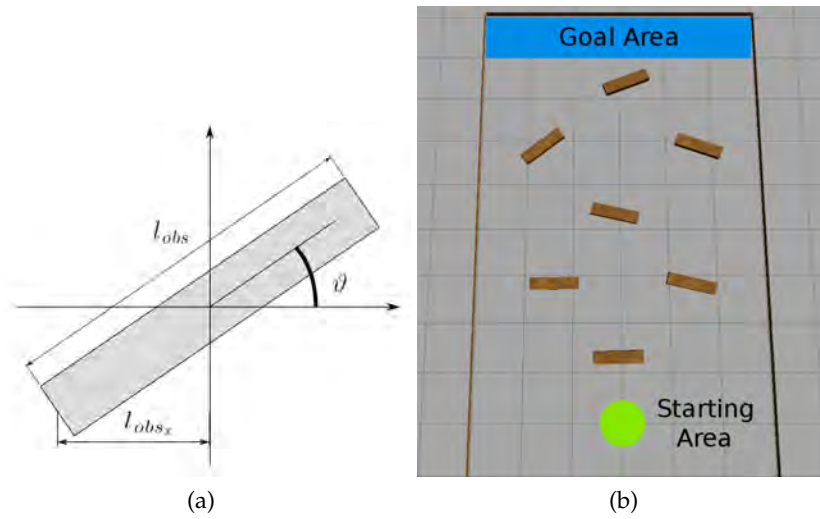


Figure 5.2: (a) The controlled obstacle parameters in the first two sets of experiments. (b) An example of complex environment.

whereas in the second set we also analyze the impact of the varying projected size, keeping l_{obs} fixed. Execution times are reported in time-steps. Each simulated second corresponds to 10 time-steps.

In the third set of experiments, we randomly generate complex environments, of the type depicted in Figure 5.2b, in which a number of cuboid-shaped obstacles are present, each with an arbitrary position and orientation. Since here we are not interested in any comparison, we only report the success rate of the method. We executed a total of 1000 runs, where in each run the angle and an offset of the position of each obstacle is generated at random.

As done in Chapter 4 and 6 for flocking, a light is placed in the environment to establish a common frame of reference for all the robots (more details on the implementation and on the used sensors and actuators can be found in Appendix A.1.3). For the sake of simplicity, the robots use the direction of the light source as the goal direction, that is, they perform photo-taxis.

5.4 Results

Figure 5.3 shows the results for the first two sets of experiments performed in the simple environments. The success rate is 100%, and the initial hypothesis can be accepted, as the execution times solely depends on ϑ and not on the projected length l_{obs_x} of the obstacle. In fact, execution times increase with increasing values for ϑ , that is, the more the obstacle is perpendicular to the direction of motion, the longer it takes for the robots to perform obstacle avoidance. However, execution times are almost identical in the two cases (Figure 5.3a and

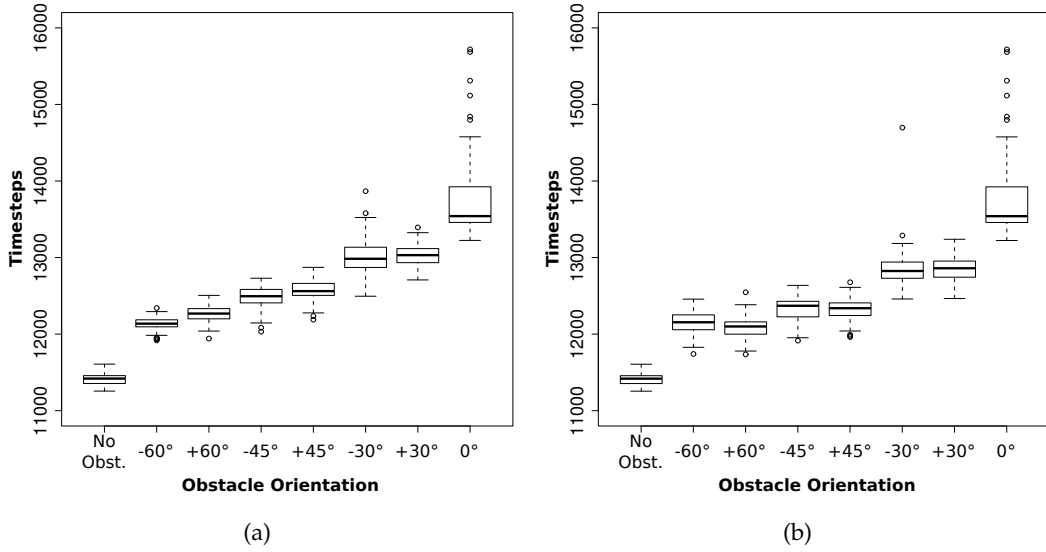


Figure 5.3: Box plot of completion time for the experiment set with fixed l_{obs_x} (a) and for fixed l_{obs} (b).

Figure 5.3b) for the same value of ϑ .

The case $\vartheta = 0$ is particularly problematic. Average times are much higher, and many more outliers are present (not fully shown due to scale differences). This is explained by the fact that, when the obstacle is perpendicular to the direction of motion, i.e. $\vartheta = 0$, the avoidance direction takes some time to converge to one of the two possible obstacle avoidance sides. Despite this, all the runs were successful and no collision was registered.

In the third set of experiments, results showed a success rate of 96%. In the remaining 4% of the cases, robots hit an obstacle and hence the corresponding run was terminated. After analyzing failure cases separately, we found out that they were all due to the slow turning rate of the compound robot structure in the goal direction after avoiding an obstacle. This slow turning rate made the robot hit the next obstacle when too close. A video showing one typical run for this set of experiments can be found in a web page ([Ferrante et al., 2010a](#)).

5.5 Summary and Discussion

In this chapter, we presented an application of explicit information transfer to tackle obstacle avoidance in collective transport. The task involves collective transport of an object by a group of three robots. In this task, robots assembled to an object need to navigate to a given goal location while avoiding obstacles.

The proposed method consists of two interacting components. The first component is called social mediation and is used to perform negotiation of a common direction through

explicit information transfer by taking into account possibly conflicting perceptions of the members of the group. The second component achieves collective transport, using this mediated direction.

Experiments were performed in a simple environment with one obstacle placed at different angles and in a more complex environment with several obstacles. Results in the simple environment show that the efficiency (inversely linked to execution times) of the method solely depends on the angle at which obstacles are placed, and that the more the obstacle is placed perpendicularly to the direction of motion the more time it takes to avoid it. In a more complex environment, we measured the success rate of the proposed approach, obtaining 96% of successful runs.

This work has been used in a practical situation. In the context of the Swarmanoid project, the method presented in this chapter was used in the final demonstration scenario: two foot-bots transported an hand-bot from a source location to a destination and back to the source, while going through doors and avoiding obstacles and other groups of robots performing the same task¹. The work has also attracted the attention of the research community working on formal methods. In [Gjondrekaj et al. \(2012\)](#), a theoretical model of the system was developed and used to prove some properties of the algorithm, such as that no situations involving a cyclic dependency (i.e., no “deadlocks”) can arise. Finally, the work can be extended in a number of directions. For example, it would be interesting to perform a direct comparison between explicit and implicit information transfer in collective transport. This would imply re-synthesizing or re-engineering the implicit information transfer strategy developed in [Baldassarre et al. \(2006\)](#) to directly compare it to the one presented in this chapter. A second and more ambitious objective would be to control a group of an arbitrary number of robots connected to each other and/or to an irregular object at different positions: We speculate that the social mediation methodology can be extended to tackle dynamic negotiation of heading direction with an arbitrary number of robots.

¹Swarmanoid: the movie, <http://www.youtube.com/watch?v=M2nn1X9X1ps>, January 2013

Chapter 6

Implicit information transfer via motion Control

In this chapter, we present the experiments performed on implicit information transfer. We show that it is possible to achieve flocking without the need of explicit communication. We achieve this via a novel motion control method, that is, a novel method to convert sensing information into robot's wheel speed. Despite the simplicity of the method, the mechanism underlying implicit information transfer is complex and hard to understand. For this reason, after analyzing the method empirically in this chapter, in Chapter 7, we use statistical physics to better analyze its dynamics and to understand what are the driving forces behind it.

The chapter is organized as follows. We first present the problem description (Section 6.1). We then define the metrics (Section 6.2) and the experimental setup (Section 6.3). In Section 6.4, we analyze an hybrid explicit/implicit information transfer setting. We then perform flocking experiments with implicit information transfer, both without (Section 6.5) and with (Section 6.6) informed robots. We conclude the chapter with a summary and discussion (Section 6.7).

6.1 Problem description

In this chapter, we propose and study *magnitude dependent motion control* (MDMC), a mechanism that can be used to achieve flocking via implicit information transfer.

In the robotics literature, to the best of our knowledge, flocking has been achieved mainly relying on explicit information transfer (Turgut et al., 2008a; Çelikkanat and Şahin, 2010). This means that alignment control, or an equivalent method, is used. In other work, flocking has been achieved either relying on simulated sensors (Hayes and Dormiani-Tabatabaei, 2002; Holland et al., 2005) or by making all or the majority of the robots aware of the desired goal direction (Matarić, 1994; Spears et al., 2004), thus not needing directional information

transfer. However, when the knowledge of the goal direction is widely available in the swarm, it is not clear whether the alignment of all robots in the same direction results from this global information or rather from a self-organization process.

MDMC has been described in Section 3.5. MDMC is a method that, together with proximal control, achieves flocking with minimal components. By minimal we mean that, differently from Turgut et al. (2008a), alignment control is not needed; differently from (Hayes and Dormiani-Tabatabaei, 2002; Holland et al., 2005), external hardware is not needed; differently from (Matarić, 1994; Spears et al., 2004), information about the goal direction is not needed. Our method achieves ordered and cohesive flocking using very simple robot hardware: robots only need to be non-holonomic and to sense their neighbors' range and bearing. Furthermore, we believe that our model is simple but powerful enough to be used as a reference model for biology, for studying the movement of very simple animals (Simpson et al., 2006) or cells (Szabó et al., 2006), in which information exchange appears to be implicit. Additionally, as explained in Chapter 7, the model is useful also in the context of statistical physics, for studying the principles underlying swarms of active particles (Vicsek et al., 1995; Romanczuk et al., 2009).

In this chapter, we directly compare MDMC with *magnitude independent motion control* (MIMC), which was used in Turgut et al. (2008a) and in Çelikkanat and Şahin (2010), showing that only MDMC is able to achieve flocking without alignment control and without informed robots. We systematically study self-organized flocking by performing extensive simulations and validating our results on real robots. First, we analyze the case with hybrid explicit/implicit information transfer: here, only a fraction of the robots, which we call the *aligning* robots, use alignment control (Section 6.4). We then conduct different experiments in which either all robots or none of the robots use alignment control and in which informed robots are not present (Section 6.5). We then repeat the latter experiments by also introducing informed robots (Section 6.6).

6.2 Metrics

In this section we describe the metrics used to measure performance. In addition to order (Vicsek et al., 1995) and accuracy (Couzin et al., 2005), presented in Section 4.2, in this chapter we introduce and use the following new metrics:

Group cohesion: This metric is used only in experiments with hybrid explicit/implicit information transfer. To compute it, we determine the number of groups g present at the end of each experiment (Couzin et al., 2005), as described in Section 4.2. Group cohesion is then computed as:

$$\varphi = 2 - \min(2, g).$$

and takes therefore values in $[0, 1]$.

Rescaled group speed: This metric is also only used in experiment with hybrid explicit/implicit information transfer. To compute it, we first calculate the average group speed as:

$$s = \left\| \frac{\mathbf{c}_T - \mathbf{c}_0}{T} \right\|,$$

where T is the total duration of one experiment and \mathbf{c}_T and \mathbf{c}_0 are the positions of the center of mass of the swarm at the end and at the beginning of the experiment, respectively. We then rescale the average group speed:

$$s_r = \frac{s}{U},$$

where U is the maximum forward speed of motion control. s_r takes therefore values in the $[0, 1]$ interval.

Steady-state values: To analyze the asymptotic dynamics of the system, we assume that the system always reaches a steady state, and we measure the steady-state value of the order ($\bar{\psi}$) and of the accuracy ($\bar{\delta}$) metrics by computing the asymptotic values reached by the order and the accuracy metric at the end of an experiment. They are computed as the average value of the metric in the last 100 simulated seconds:

$$\bar{\psi} = \frac{\sum_{(t=T-100)}^T \psi_t}{100}, \quad \bar{\delta} = \frac{\sum_{(t=T-100)}^T \delta_t}{100}.$$

Here, ψ_t denotes the order and δ_t the accuracy at time t .

Settling times: The order and accuracy settling time are the time needed to reach the steady-state value of the order and of the accuracy metric, respectively. More precisely, the order (accuracy) settling time t_ψ^* (t_δ^*) is defined as the time for which $\forall t \geq t_\psi^*$ ($\forall t \geq t_\delta^*$) we have $\psi_t \geq 0.95\bar{\psi}$ ($\delta_t \geq 0.95\bar{\delta}$). In other words, t_ψ^* (t_δ^*) is the time after which the value of the order (ψ_t) (accuracy δ_t) metric remains consistently above the 95% of the steady-state value $\bar{\psi}$ ($\bar{\delta}$) of the order (accuracy) metric.

Effective traveled distance: The effective traveled distance is the distance, in meters, traveled following the desired goal direction. It is computed as the total distance traveled projected onto the goal direction \mathbf{g}_1 :

$$D_e = (\mathbf{c}_T - \mathbf{c}_0) \cdot \mathbf{g}_1.$$

Variable	Description	Value(s) / Range
N	Number of robots	{25, 100}
R	Number of runs per setting	50
ρ_A	Prop. of aligning robots	{0.4, 0.8}
β_1/α_1	Aligning robots parameters	{1, 2, 4, 6, 8, 10}
α_2	Non-aligning robots parameter	{1, 2, 4, 6, 8, 10}
U	Maximum forward speed	1.5 cm/s
K_1	MDMC linear gain	0.25 cm/s
K_2	MDMC angular gain	0.1 rad/s
K_3	MIMC angular gain	0.5 rad/s
l	Inter-wheel distance	0.1 m
U	max forward speed	20 cm/s
Ω	max angular speed	$\pi/2$ rad/s
ϵ	Strength of potential function	0.5
d_{des}	Desired inter-robot distance	0.6 m
ξ	Amount of noise	0.1
T	Duration of experiments	600 s
Δt	ARGoS integration time-step and real robot control step	0.1 s

Table 6.1: Experimental values or range of values for all constants and variables used in Experimental setting 1

Here, \mathbf{c}_T and \mathbf{c}_0 are the position vectors of the center of mass of the swarm at time $t = T$ and $t = 0$, respectively; \mathbf{g}_1 is the goal direction (unique in this case), and \cdot denotes the dot (scalar) product of two vectors.

6.3 Experimental setup

In this section, we describe the experimental setup used in simulation and with the real robots.

6.3.1 Simulation experimental setup

We compared MDMC, the proposed motion control method, with MIMC, the motion control method used in [Turgut et al. \(2008a\)](#), and [Çelikkanat and Şahin \(2010\)](#). We conducted three sets of experiments.

1. Hybrid explicit-implicit information transfer Here, there is no robot informed about the desired goal direction, and the swarm is composed of a proportion $\rho_A \in [0, 1]$ of *aligning robots*, that use alignment control. The rest of the swarm is composed of *non-aligning robots* that do not use alignment control.

Variable	Description	Value
N	Number of robots	{10, 50, 100 , 500, 1000}
R	Number of runs per setting	100
ρ_I	Proportion of informed robots	{ 0.01 , 0.05, 0.1, 0.15, 0.2}
α	Weight of proximal control	1
β	Weight of alignment control	1
γ	Weight of goal direction	1
U	Maximum forward speed	0.005 m/s
Ω	Maximum angular speed	$\pi/2$ rad/s
K_1	MDMC linear gain	0.5 cm/s
K_2	MDMC angular gain	0.06 rad/s
K_3	MIMC angular gain	0.25 rad/s
l	Inter-wheel distance	0.14 m
ϵ	Strength of potential function	1.5
σ	Distance-related proximal control parameter	0.4 m
d_{des}	Desired inter-robot distance	0.56 m
D_p	Maximum interaction range of proximal control	$1.8d_{des} = 1.008$ m
D_a	Maximum interaction range of alignment control	2.0 m
ξ	Amount of noise	0.1
T	Duration of experiments	2500 s
Δt	ARGoS integration time-step and real robot control step	0.1 s

Table 6.2: Parameter values or range of values used in the controller and in simulations in Experimental setting 2 and 3. The values in **bold** are the nominal values, that is, the values used in experiments when the effect of another parameter is explored.

2. Implicit information transfer without informed robots Here, there is no robot informed about the goal direction. We perform experiments with all robots and with no robot using alignment control, and we analyze the impact of different swarm sizes on the transient and on the steady-state dynamics of the system.

3. Implicit information transfer with informed robots Here, we select a proportion ρ_I of informed robots and we inform them about the goal direction. The informed robots and the goal direction are randomly selected at the beginning of each run. All the other robots remain uninformed during the experiment. Similarly to Experimental setting 2, we perform experiments with all robots and with no robot that uses alignment control.

We now explain the experimental setup by making the distinction between each set of experiments.

Variable	Description	Value
N	Number of robots	8
R	Number of runs per setting	10
ρ_I	Proportion of informed robots	0.2
T	Duration of experiments	300 s
	All the other control parameters	See Table 6.2

Table 6.3: Parameter values or range of values used for the real robot experiments in Experimental setting 2 and 3. All the parameters related to the controllers were omitted in this table, since they are the same as in simulation (refer to Table 6.2).

Experimental setting 1

In Experimental setting 1, aligning robots compute the flocking control vector using both proximal control and alignment control (see Section 3.2):

$$\mathbf{f} = \alpha_1 \mathbf{p} + \beta_1 \mathbf{h},$$

whereas non-aligning robots only use proximal control:

$$\mathbf{f} = \alpha_2 \mathbf{p}.$$

We study the impact of the relative contribution of proximal and alignment control (α_1 , α_2 and β_1). We further classify Experimental setting 1 in two sets:

MIMC-MIMC In this case, all robots use MIMC. Here, we study the effect of the ratio $\frac{\beta_1}{\alpha_1}$, and we do not change α_1 and β_1 independently, since MIMC does not utilize the magnitude of \mathbf{f} , but only its angular component. As such, multiplying both α_1 and β_1 with the same constant value will produce no difference in the robot motion. For the same reason, α_2 does not affect the robot motion.

MIMC-MDMC In this case, aligning robots use MIMC whereas non-aligning robots use MDMC. For the non-aligning robots, the magnitude of \mathbf{f} plays a role in their motion. Thus, additionally to studying the impact of $\frac{\beta_1}{\alpha_1}$ to the dynamics of the aligning robots, we study the impact of α_2 to the dynamics of the non-aligning robots.

We analyze flocking with medium ($N = 25$) and large ($N = 100$) swarm sizes and with low ($\rho_A = 0.4$) and high ($\rho_A = 0.8$) proportions of aligning robots. We study the effect of changing the ratio $\frac{\beta_1}{\alpha_1} \in \{1, 2, 4, 6, 8, 10\}$ and, for the MIMC-MDMC case, we also study the effect of changing $\alpha_2 \in \{1, 2, 4, 6, 8, 10\}$, but we report only the results obtained with the best case, that is, $\alpha_2 = 10$ (refer to Stranieri et al. (2011b) for the complete set of results). In a

supplementary web page (Stranieri et al., 2011b), we also report the flocking performance as a function of $\rho_A \in \{0.2, 0.4, 0.6, 0.8, 1.0\}$.

We study how the heterogeneous flocking performance is influenced by: i) the way robots implement their motion (MIMC-MIMC motion versus MIMC-MDMC motion), ii) the parameters that affect the strength of the proximal control vector and of the alignment control vector, that is, $\frac{\beta_1}{\alpha_1}$ and α_2 , and iii) the ratio of aligning robots ρ_A .

Experimental setting 2 and 3

In Experimental setting 2 and 3, we utilize the models described in Chapter 3 in the following way: In experiments without informed robots and without alignment control, all robots use Model 3.2 (Equation 3.2 - only proximal control). In experiments without informed robots and with alignment control, all the robots use Model 3.4 (Equation 3.4 - proximal control and alignment control). In experiments with informed robots and without alignment control, informed robots use Model 3.3 (Equation 3.3 - proximal control and goal direction), whereas the rest of the swarm uses Model 3.2 (Equation 3.2 - only proximal control). Finally, in experiments with informed robots and with alignment control, informed robots use Model 3.1 (Equation 3.1 - proximal control, alignment control and goal direction) and the rest of the swarm uses Model 3.4 (Equation 3.4 - proximal control and alignment control).

In these experiments, we study both the transient and steady state behavior of the system. When analyzing the transient behavior, we report the time evolution of the median value of the order or of the accuracy metrics, with the corresponding error bars showing the values at the 25% and the 75% percentile scores (see Figure 6.3 for an example). When analyzing the steady-state behavior, we report the box-plot of the distribution of the steady-state values and of the settling times. In experiments with informed robots, we also report the box-plot of the distribution of the effective traveled distances.

Experimental setting 1, 2 and 3

In all experiments, we add noise to several components of our system: to the proximal control vector \mathbf{p} , to the alignment control vector \mathbf{a} , and to the goal direction vector \mathbf{g}_1 . We consider only angular noise, as commonly done in flocking studies (Vicsek et al., 1995; Turgut et al., 2008a), and we model it as a random variable uniformly distributed in the range $(-\xi 2\pi, +\xi 2\pi)$. For each experimental setting, we execute R independent runs.

At the beginning of each run, N mobile robots are placed within a circle at random positions and orientations. The radius of the circle is chosen so that the mean density of the initial configuration is kept fixed at 5 robots per square meter, independently of the number of robots N . In the experiments with informed robots or with aligning robots, a light source is also placed at a fixed position in the arena, in such a way to be seen by all robots that use it as a global reference frame.

6.3.2 Real robot experimental setup

In order to validate MDMC, we perform experiments on real robots for Experimental setting 2 and 3. The details of the experimental setup are the same as the ones reported in Section 4.3.2 and will not be repeated here.

We conduct two sets of experiments as we did in simulation: without and with informed robots. All parameters are kept the same as in the simulations. We only change the number of robots and the duration of the experiments in order to comply with arena-size constraints. We consider a proportion of 25% informed robots. The parameters used with real robots are summarized in Table 6.3. All other parameters are the same as used in simulation (Table 6.2).

For each of the two sets of experiments, and for each of the two motion control methods, we execute 10 runs and report the median values of the metrics together with their first and third quartiles.

6.4 Hybrid explicit-implicit information transfer

We present here the results of the hybrid explicit-implicit information transfer case.

6.4.1 Results in the MIMC-MIMC case

The experimental results for MIMC-MIMC case are depicted in Figure 6.1. We first focus on the $\rho_A = 0.8$ case, for both $N = 25$ (Figure 6.1a) and $N = 100$ (Figure 6.1b). Results show that the swarm is cohesive in most runs. However, order and rescaled group speed are high only when $\frac{\beta_1}{\alpha_1} \geq 2$. Furthermore, while order is high at different values of the ratio $\frac{\beta_1}{\alpha_1}$, rescaled group speed increases with increasing values of $\frac{\beta_1}{\alpha_1}$, until it saturates at around $\frac{\beta_1}{\alpha_1} = 6$. This shows that, when the contribution of alignment control is higher, robots tend to move faster. This is explained by the fact that the alignment control vector has less fluctuations, over time, compared to the proximal control vector. Thus, the higher the contribution of alignment control, the more the robots tends to move forward rather than turning. This allows the swarm to move faster, until speed saturates at the maximum forward speed U .

When the proportion of aligning robots is $\rho_A = 0.4$, performance gets sensibly worse (Figures 6.1c and 6.1d). In both cases ($N = 25$ and $N = 100$), we observe two possible outcomes: for small values of the ratio $\frac{\beta_1}{\alpha_1}$ the swarm remains cohesive but does not move. This happens because the contribution of alignment control is not high enough to have the aligning robots pulling the swarm in a direction. For larger values of the ratio $\frac{\beta_1}{\alpha_1}$, rescaled group speed and order are higher. However, in this case, the swarm splits in at least 25% of the runs. This happens because, in those runs, clusters of non-aligning robots are formed. Since the motion of non-aligning robots is governed only by proximal control, they tend to turn on spot more frequently than moving forward. Thus, non-aligning robots are not able

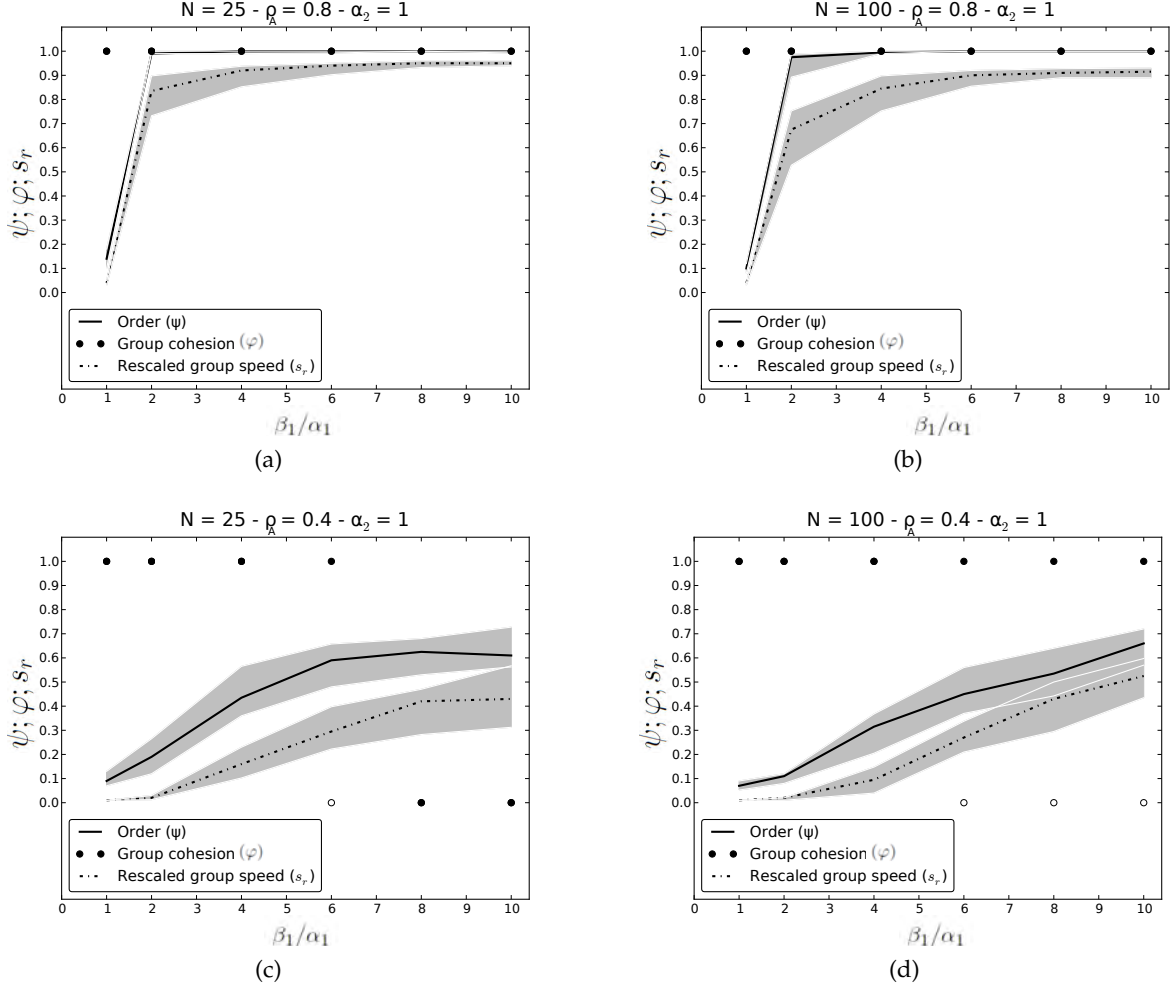


Figure 6.1: MIMC-MIMC case experiments for varying swarm size ($N \in \{25, 100\}$) and ratio of aligning robots ($\rho_A \in \{0.4, 0.8\}$). Thick lines show the median values, whereas the gray areas show the 25% and the 75% inter-quartile range of the data. For group cohesion, filled circles correspond to median values and empty circles to the 25% percentile score of the data.

to match the higher speed of the aligning robots and remain disconnected from the group.

In [Stranieri et al. \(2011b\)](#), we also report the performance as a function of ρ_A . We consider the case $\frac{\beta_1}{\alpha_1} = 10$, as it generally provides the best overall results. As shown in [Stranieri et al. \(2011b\)](#), only $\rho_A \geq 0.6$ corresponds to acceptable flocking performance, both with $N = 25$ and $N = 100$.

6.4.2 Results in the MIMC-MDMC case

In the MIMC-MDMC case, results with $\rho_A = 0.8$ (Figures [6.2a](#) and [6.2b](#)), are similar to the results obtained, with the same ratio, in the MIMC-MIMC case. Instead, the results with

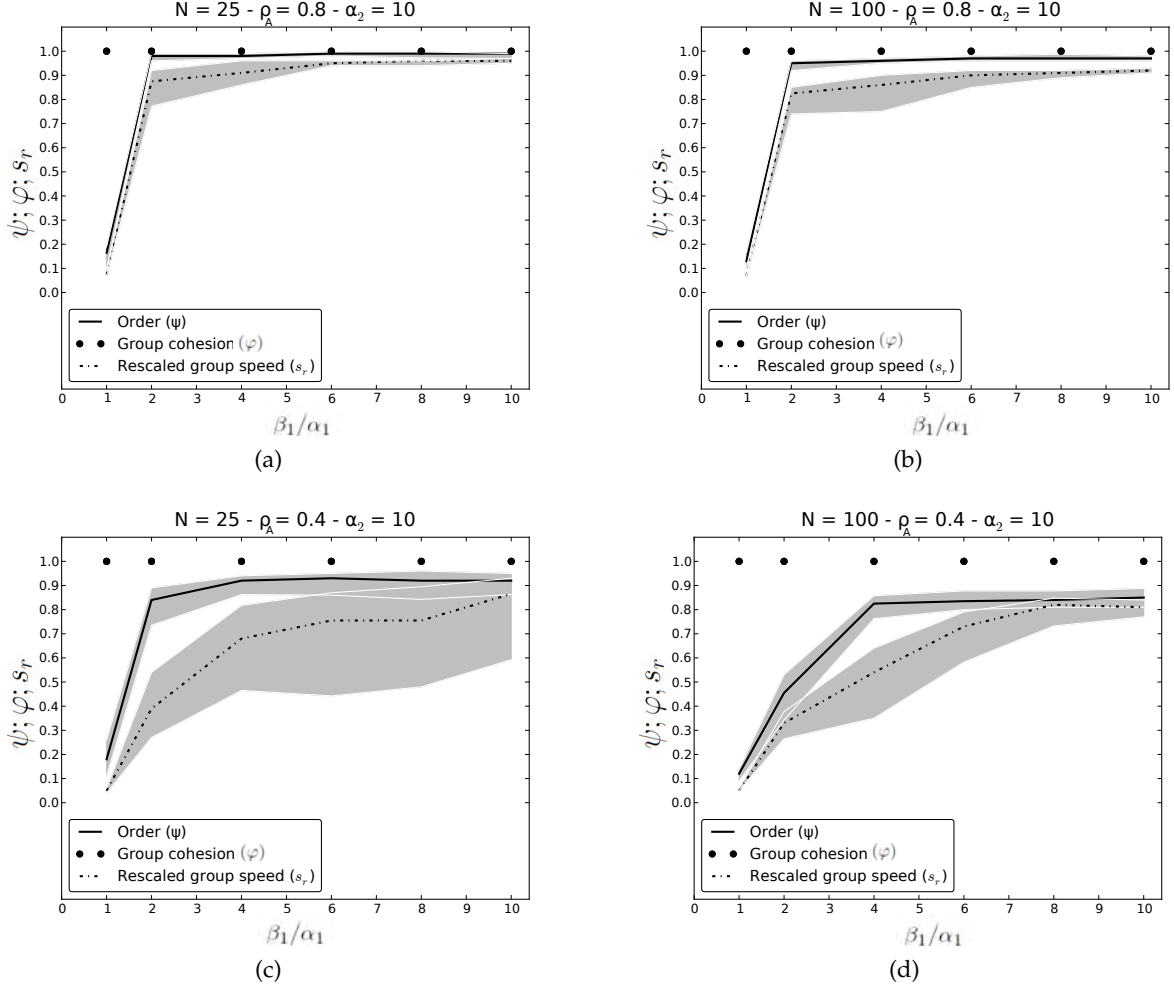


Figure 6.2: MIMC-MDMC case experiments for varying swarm size ($N \in \{25, 100\}$) and ratio of aligning robots ($\rho_A \in \{0.4, 0.8\}$). Thick lines show the median values, whereas the gray areas show the 25% and the 75% inter-quartile range of the data. For group cohesion, filled circles correspond to median values and empty circles to the 25% percentile score of the data.

$\rho_A = 0.4$ are much better in the MIMC-MDMC case (Figures 6.2c and 6.2d) with respect to the MIMC-MIMC case (Figures 6.1c and 6.1d). With both swarm sizes we have that, when $\frac{\beta_1}{\alpha_1} > 2$, the swarm is able to effectively flock together at the cost of a reduced speed.

In Stranieri et al. (2011b), we also report the flocking performance as a function of ρ_A for $\frac{\beta_1}{\alpha_1} = 10$ and $\alpha_2 = 10$. Differently from the MIMC-MIMC case, in the MIMC-MDMC case the performance of flocking degrades more gracefully as the proportion of non-aligning robots decreases.

The improved capability of the swarm to stay together is due to the advantage of using MDMC in the non-aligning robots. In fact, non-aligning robots are able to respond to the

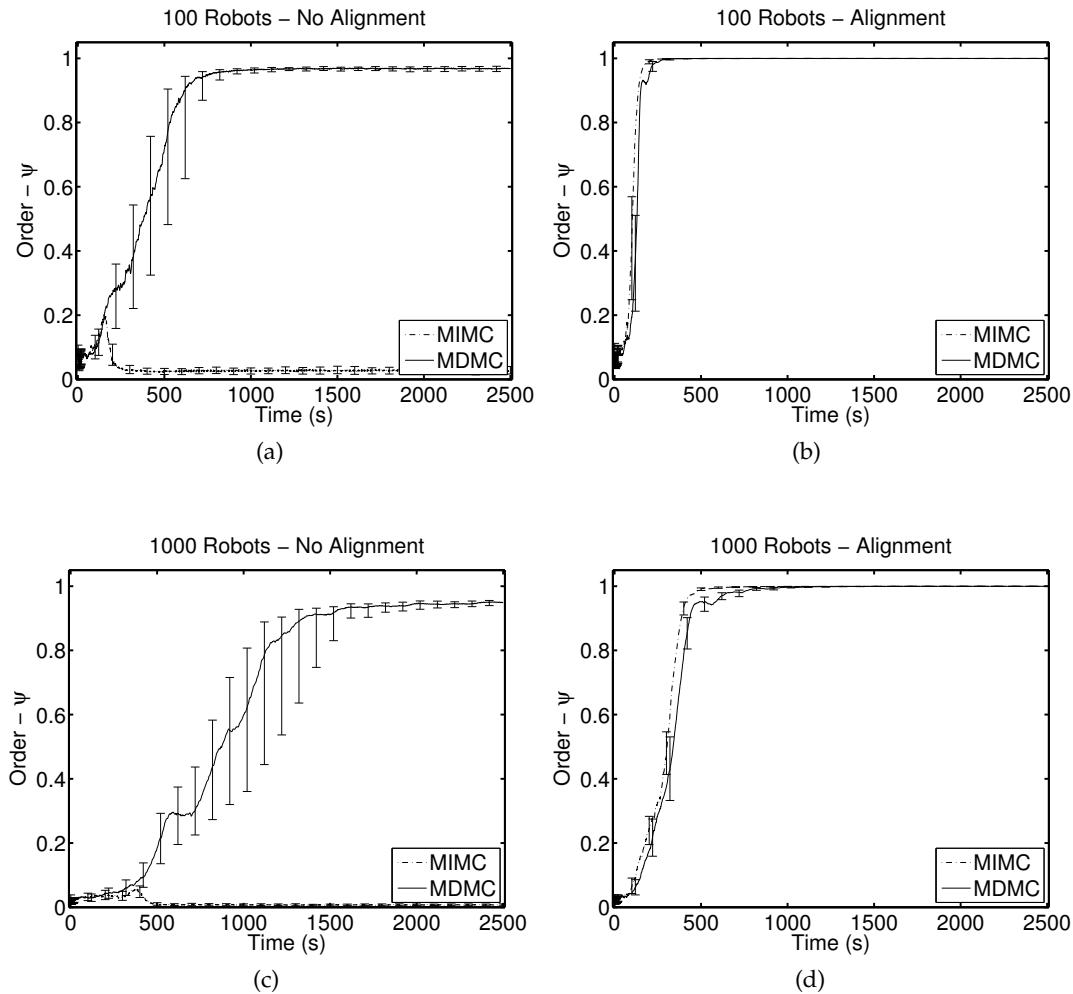


Figure 6.3: MDMC vs MIMC in simulation experiments without informed robots. Time evolution of the order metric for (a) 100 robots - without alignment control, (b) 100 robots - with alignment control, (c) 1000 robots - without alignment control and (d) 1000 robots - with alignment control. The line denotes the median of the distribution obtained over $R = 100$ runs whereas the error bars represent the interval between the 25% and the 75% quartiles of the distribution.

high variations in the proximal control vector much better if they are also able to modify their forward speed. This, in turns, is possible with MDMC but not with MIMC. As such, non-aligning robots are also able to remain together with the aligning robots more successfully. Finally, the reduced speed and the high variation of speed among runs is due to the following fact. In presence of a low proportion of aligning robots, we observed that the direction of the swarm is stable over short periods of time but changes over long periods of time due to the disturbances caused by the non-aligning robots. This results in a non-linear swarm trajectory, which is different at each run. Since the rescaled group speed is computed




Figure 6.4: MDMC vs MIMC in simulations with no informed robots and without alignment control. Box-plots of the distribution of (a) steady-state values and (b) settling times of the order metric for different swarm sizes.

assuming a linear trajectory, this metric has large variation from run to run.

6.5 Implicit information transfer without informed robots

We now present results with implicit information transfer without informed robots.

6.5.1 Simulations

Transient behavior: Figure 6.3 shows the transient behavior of the system in a medium-sized ($N = 100$) and in a large ($N = 1000$) swarm. Figures 6.3a and 6.3c show the results obtained without alignment control (Equation 3.2). As we can see, the swarm reaches an ordered state only when using MDMC. The ordered state is reached within 700 simulated seconds in the medium-sized swarm and within 1500 simulated seconds in large swarms. When MIMC is used, the swarm remains disordered during the entire experiment. Figures 6.3b and 6.3d show the results obtained with alignment control. In this case, the alignment state is reached with both motion control methods, and the response of the system is much faster than without alignment control. The performance of MDMC and MIMC are comparable: the system self-organizes in less than 200 seconds in medium-sized swarms and in less than 400 seconds in large swarms.

Steady state and settling time without using alignment control: Figure 6.4 shows the steady-state behavior of the system for different swarm sizes. Figure 6.4a shows box-plots




Figure 6.5: Results with no informed robots using MDMC in real-robots experiments. Time evolution of the order metric for 8 robots (a) without alignment control and (b) with alignment control. The line denotes the median of the distribution obtained over $R = 10$ runs whereas the error bars span the interval between the 25% and the 75% quartiles of the distribution.

of the distributions of the asymptotic values of the order metric reached without alignment control. Figure 6.4b shows box-plots of the distributions of the settling times without alignment control.

An important result of this chapter is shown in Figure 6.4a: by using MDMC, the system achieves an ordered state without the need of alignment control for all swarm sizes. The use of MDMC is critical for achieving this result. In fact, as shown in Figure 6.4a, the system never reaches the ordered state when using MIMC.

Figure 6.4a also displays points indicating results that were considered as outliers. These are of two different types. Some correspond to cases where the system is trying to self-organize but is too slow to reach its stationary state within the maximum run time of $T = 2500$ s, and can be easily identified by looking at the time series. The others correspond to cases in which the swarm, instead of achieving aligned motion, converged to a rotating state with fixed center of mass, low values of order metric ($\psi < 0.5$) and high angular rotation. A video of a rotating state can be found in our supplementary material page (Ferrante et al., 2012b). For the purpose of this chapter, we consider these as outliers and focus on the convergence to parallel states. Convergence to a rotating state is rare, becomes less probable if more robots are used, and can be easily suppressed by imposing minor constraints to the robot dynamics. Its analysis goes beyond the scope of this chapter and is presented together with the statistical physics analysis in Section 7.6.1.

In Figure 6.4b we report the settling time needed to reach the ordered state. The settling

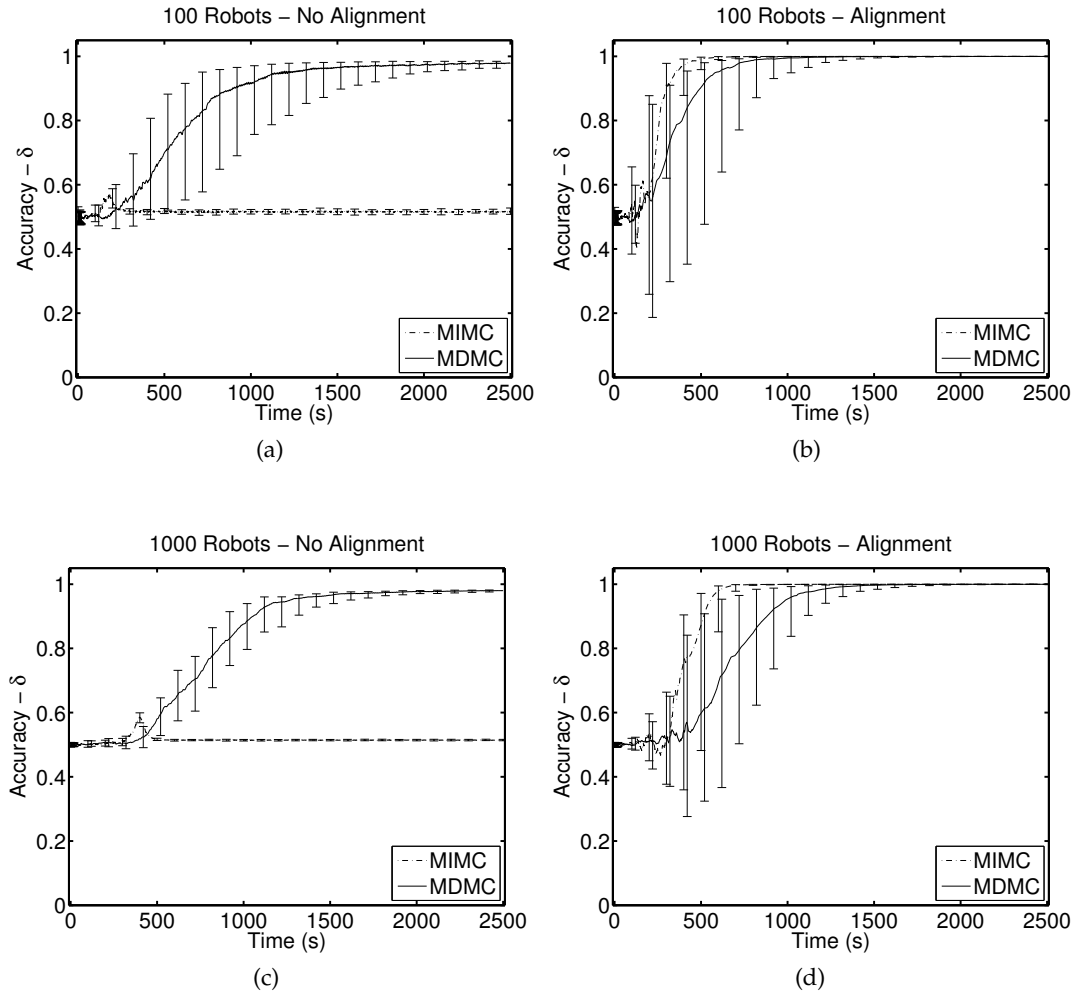


Figure 6.6: MDMC vs MIMC in simulations with a proportion of $\rho_I = 0.01$ informed robots. Time evolution of the accuracy metric for (a) 100 robots without alignment control, (b) 100 robots with alignment control, (c) 1000 robots without alignment control and (d) 1000 robots with alignment control. The line denotes the median of the distribution obtained over $R = 100$ runs whereas the error bars represent the interval spanned between the 25% and the 75% quartiles of the distribution.

time obtained with MDMC increases with increasing swarm sizes, as expected. We also noticed some outliers in the distribution of settling times for MDMC. These correspond to runs where the systems first self-organizes in the rotating solution, and only later, due to noise, reorganizes in the ordered solution. Settling times for MIMC are not reported as the system does not converge to a steady-state in the allotted time (2500 s).

Steady state and settling time when using alignment control: Additional results reported in the supplementary material page (Ferrante et al., 2012b) show that the system reaches the

ordered state for any swarm size and any motion control method when alignment control is used. This is consistent with previous studies in flocking (Turgut et al., 2008a) and, additionally, it shows that MDMC can also be implemented together with alignment control. Additionally, for both methods the medians of the settling times increase with increasing swarm sizes.

6.5.2 Real robots

Figure 6.5 shows the results obtained with MDMC using real robots. Results in Figure 6.5a show that, when alignment control is not used, the system reaches reasonable levels of order also with the real robots. The difference in performance when compared to the simulations can be attributed to the different nature and amount of noise on the real robots, which is very difficult to replicate in simulation. Figure 6.5b shows that, when alignment control is used, results are less noisy. This is because alignment control uses a communication device, the RAB, that is less noisy. Its only source of noise comes from orientation measurements via the light sensor, which appears to be precise enough. Results with no alignment control also exhibit much larger fluctuations of the order metric than those with alignment control. Finally, note that both in Figure 6.5a and in Figure 6.5b the value order metric becomes constant at $t \approx 280$ s and at $t \approx 180$ s, respectively. This is because by time $t \approx 280$ s (respectively $t \approx 180$ s) all experiments had been stopped due to the robots having reached the borders of the arena.

6.6 Implicit information transfer with informed robots

We now present results with implicit information transfer with informed robots.

6.6.1 Simulations

Accuracy vs time: Figure 6.6 shows the transient behavior of the system for $N = 100$, $N = 1000$ and a proportion of informed robots $\rho_I = 0.01$. Figures 6.6a and 6.6c show results obtained without alignment control. For this case, the swarm reaches a value of accuracy close to 1 only when using MDMC. This value is reached within 1000 s and 1500 s for swarms with $N = 100$ and $N = 1000$ robots, respectively. When using MIMC, accuracy of the system remains $\delta = 0.5$. As explained in Section 6.2, this is the value obtained when the order metric is close to 0. In other words, when using MIMC without alignment control, the system is not able to reach an ordered state when the proportion of informed robots is small ($\rho_I = 0.01$). Complete results that include also the value of the order metric are reported in our supplementary material page (Ferrante et al., 2012b). However, as we will show later, when the proportion of informed robots is increased, the system self-organizes also with

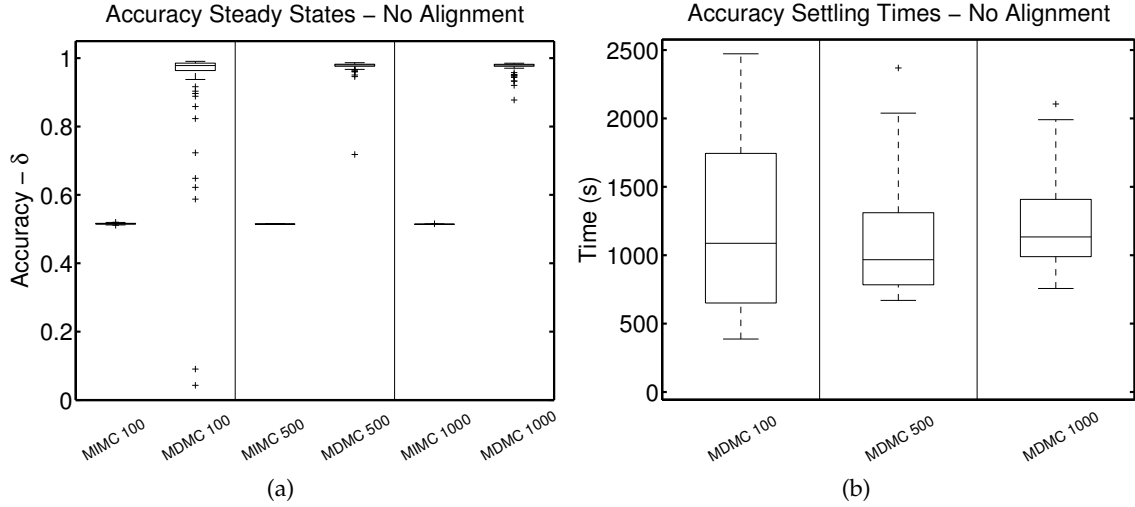


Figure 6.7: MDMC vs MIMC in simulations with a proportion of $\rho_I = 0.01$ informed robots without alignment control. Box-plots of the distribution of (a) steady-state values and (b) settling times of the accuracy metric for different swarm sizes.

MIMC. Figures 6.6b and 6.6d show the results obtained with alignment control. In this case, a proportion $\rho_I = 0.01$ of informed robots is enough to reach high levels of accuracy for both MIMC and MDMC. Furthermore, MIMC seems to have better performance in terms of convergence time.

Steady state and settling time of accuracy without using alignment control: Figure 6.7 shows the steady-state behavior of the system for different swarm sizes, with $\rho_I = 0.01$ and without alignment control. Figure 6.7a shows the distribution of the asymptotic value of the accuracy metric. As we can see, MDMC dominates MIMC in all three cases. In fact, the system does not reach an ordered state when MIMC is used. This makes the accuracy metric settle approximately at $\delta \approx 0.5$. In Figure 6.7b, we see that the distribution of settling times for MDMC has 1000 s as median value.

Steady state and settling time of accuracy when using alignment control: Additional results reported in the supplementary material page (Ferrante et al., 2012b) show that, when alignment control is used, the steady-state values of the accuracy are almost always close to 1. The few outliers that are observed are due to the fact that, in rare cases, the system splits into two groups. The median settling time, as well as its spread, increases with increasing swarm sizes, and increases more with MDMC than with MIMC.




Figure 6.8: MDMC vs MIMC in simulations with informed robots, no alignment control, a swarm size of $N = 100$ robots and $\rho_I = \{0.01, 0.05, 0.1, 0.15, 0.2\}$. Box-plots of the distribution of (a) steady state values and (b) settling times of the accuracy metric for different proportions of informed robots.

Impact of the proportion of informed robots: Figure 6.8 reports the results of the study conducted with fixed swarm size $N = 100$ and various proportions of informed robots. Figure 6.8a shows that, when alignment control is not used, MDMC dominates MIMC in terms of accuracy only for $\rho_I = 0.01$. In the other cases, the two methods achieve very similar performance level. However, as shown in Figure 6.9 and explained below, this apparent equivalence between MDMC and MIMC in accuracy does not result in an equivalence in the effective traveled distance. Finally, Figure 6.8b shows that the system converges faster with MDMC than with MIMC in the $\rho_I = 0.01$ case, and that the opposite is true for the other cases.

Additional experiments performed with alignment control (Ferrante et al., 2012b) show that the system achieves very good performance levels in all cases. However, MDMC converges slower than MIMC when the proportion of informed robots is low.

Effective traveled distance: Figure 6.9 shows the box-plots of the distribution of the effective traveled distance for different swarm sizes (fixing $\rho_I = 0.01$) and of the proportion of informed robots (fixing $N = 100$). When alignment control is not used, MDMC consistently outperforms MIMC for any swarm size (Figure 6.9a). This was already known since, unless the swarm is small, the system is not able to self-organize or to achieve high accuracy values when using MIMC. Figure 6.9b shows that MDMC also consistently outperforms MIMC for any proportion of informed robots. Although the system can self-organize and achieve a

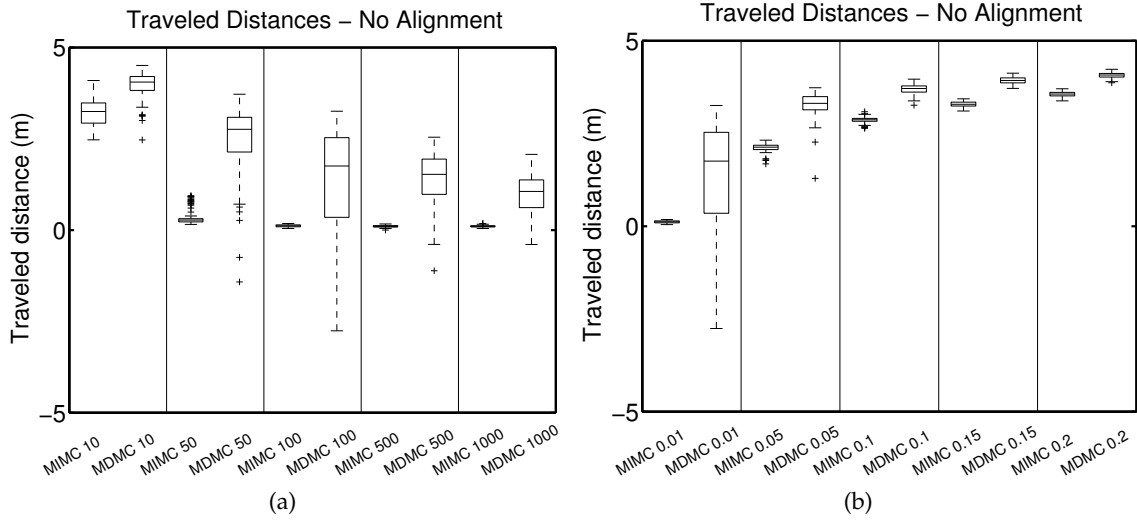


Figure 6.9: MDMC vs MIMC in simulations with informed robots and no alignment control. Box-plots of the distribution of effective traveled distance for (a) different swarm sizes ($\rho_I = 0.01$) and (b) different proportions of informed robots ($N = 100$).

high level of accuracy for a larger proportion of informed robots, MDMC still has the advantage of being able to let the swarm travel farther in the desired direction of motion.

Additional experiments performed using alignment control (Ferrante et al., 2012b) show that the swarm is able to travel in the desired goal direction for all swarm sizes. We note that MIMC performs slightly better than MDMC with alignment control, although the performance of the latter are still acceptable. Finally, we also observe that the two methods perform almost equally for a large proportion of informed robots.

6.6.2 Real robots

Figure 6.10 shows the results obtained with MDMC using real robots. The results in Figure 6.10a show that, without alignment control, the system reaches reasonable levels of accuracy. The difference in performance when compared to simulation could be due to the different way in which noise affects the system, since our simulation does not attempt to model realistically the noise. Figure 6.10b shows that, surprisingly, MDMC performs worse in real robots with alignment control. In fact, in a group of 8 real robots with 2 informed robots, the range and bearing communication device made the swarm agree on a common, random direction very fast. When this happened, the interactions between informed and non-informed robots were not strong enough to let the swarm change its direction in the desired one. Furthermore, the swarm, due to noise, continuously changed its direction, which may explain the increase in the spread and the decrease of the median of the distribution




Figure 6.10: Results with informed robots using MDMC in real-robots experiments. Time evolution of the accuracy metric for a group of 8 robots and 2 informed robots (a) without alignment control and (b) with alignment control. The line denotes the median of the distribution obtained over $R = 10$ runs whereas the error bars represent the interval spanned between the 25% and the 75% quartiles of the distribution.

after $t = 100$ s. By contrast, when alignment control is not used, the swarm never agrees on a common direction. Hence, the movement of the swarm is almost completely driven by the interaction between the informed and the non-informed robots, that in this case is enough to make the swarm move in the desired direction of motion.

6.7 Summary and discussion

In this chapter, we performed several experiments with hybrid explicit-implicit information transfer and with implicit information transfer. In doing so, we explored the capabilities of Magnitude Dependent Motion Control (MDMC), the novel motion control method presented in Section 3.5. In Table 6.4 we report a complete overview of the results of this chapter as a table of experiments.

In experiments with hybrid explicit-implicit information transfer, we have shown that flocking is possible when some individuals in the swarm lack the capability to explicitly transfer information. In particular, we showed that: i) a higher proportion of aligning robots always corresponds to better performance; ii) performance is affected by the relative contribution (weight) of alignment and proximal control, and iii) for smaller proportions of aligning robots, flocking is possible only when the non-aligning robots also change their forward speeds.

In experiments with implicit information transfer without informed robots, we have

N	Informed	ρ_I	Alignment	ρ_A	Best strategies
25	no	$\rho_I = 0.0$	hybrid	$\rho_A = 0.8$	MIMC-MIMC, MIMC-MDMC
100	no	$\rho_I = 0.0$	hybrid	$\rho_A = 0.8$	MIMC-MIMC, MIMC-MDMC
25	no	$\rho_I = 0.0$	hybrid	$\rho_A = 0.4$	MIMC-MDMC
100	no	$\rho_I = 0.0$	hybrid	$\rho_A = 0.4$	MIMC-MDMC
10	no	$\rho_I = 0.0$	no	$\rho_A = 0.0$	MDMC
50	no	$\rho_I = 0.0$	no	$\rho_A = 0.0$	MDMC
100	no	$\rho_I = 0.0$	no	$\rho_A = 0.0$	MDMC
500	no	$\rho_I = 0.0$	no	$\rho_A = 0.0$	MDMC
1000	no	$\rho_I = 0.0$	no	$\rho_A = 0.0$	MDMC
10	yes	$\rho_I = 0.01$	no	$\rho_A = 0.0$	MDMC
50	yes	$\rho_I = 0.01$	no	$\rho_A = 0.0$	MDMC
100	yes	$\rho_I = 0.01$	no	$\rho_A = 0.0$	MDMC
500	yes	$\rho_I = 0.01$	no	$\rho_A = 0.0$	MDMC
1000	yes	$\rho_I = 0.01$	no	$\rho_A = 0.0$	MDMC
100	yes	$\rho_I = 0.01$	no	$\rho_A = 0.0$	MDMC
100	yes	$\rho_I = 0.05$	no	$\rho_A = 0.0$	MDMC
100	yes	$\rho_I = 0.1$	no	$\rho_A = 0.0$	MDMC
100	yes	$\rho_I = 0.15$	no	$\rho_A = 0.0$	MDMC
100	yes	$\rho_I = 0.2$	no	$\rho_A = 0.0$	MDMC

Table 6.4: Table of experiments summarizing all the results of this chapter. We report which strategy performed best for each setting. The third and fifth column report the proportion of aligning robots and the proportion of informed robots, respectively. In experiments with informed robots, we used the effective traveled distance to determine the winning strategy.

shown that MDMC is the only method capable of achieving self-organized flocking without alignment control. We also showed that the method performs as good as Magnitude independent motion control (MIMC) when alignment control is used.

In experiments with implicit information transfer with informed robots, we have shown that MDMC always outperforms MIMC when alignment control is not used. MDMC requires less informed robots than MIMC to make the swarm align to the desired goal direction. Furthermore, the swarm is able to travel on average around 20% more in the desired direction when using MDMC compared to MIMC.

All experiments with implicit information transfer and that used MDMC were also validated on real robots, obtaining results comparable to those obtained in simulation.

Chapter 7

Statistical physics analysis of implicit information transfer

In this chapter, we use statistical physics to analyze how MDMC, the motion control method presented in Section 3.5, can achieve implicit information transfer. The analysis is performed using both numerical simulations and analytical methods.

To carry out the analysis from the point of view of statistical physics, we use a slightly different set of assumptions and terminology compared to what we have been using in the previous chapters about robotics. The main differences between this chapter and the chapters dedicated to robotics are the following:

Particles We consider a *robot* as a *particle*. Particles can be seen as simplified robot models in which only the relevant sensing and actuation mechanisms are included. In particular, particles sense the relative position of their neighbor and can move using kinematic equations and are subject to non-holonomic constraints.

Model formulation We formulate a model by rewriting MDMC using the global frame of reference rather than the local one, as it is a standard practice in statistical physics (Equation 7.1 and Equation 7.2). We kept the same variable naming where applicable, in order to clearly show the equivalence between the model analyzed here and MDMC.

Topology Initial experiments in statistical physics were conducted considering the same type of interactions as the one used with the robots: particles sensing the relative position of their neighbors within a given range of sight. This means that the neighbors of a particle can potentially dynamically change during one experiment if particle positions reshuffle. In this chapter and in the related publications (Ferrante et al., 2013b,c), we place our work in the context of active solid and active crystals. By active solid we mean a structure in which self-propelled particles are bound to each other via a fixed set of connections which depend on the molecular structure of the solid. An ac-

tive crystal is a special case of active solid in which particle-to-particle connections are homogeneous within the entire structure. In order to implement this, we consider a lattice topology in which local interactions are decided at the beginning of the simulation and kept fixed for the entire duration of the experiment. The two types of interactions, the one used in robotics and the one used in this chapter, present no qualitative difference in the results, as particles rarely change their positions with respect to each other during one experiment due to the attraction-repulsion rule used.

Potential function The proximal control method described in Section 3.3 is based on the Lennard-Jones potential function (Jones, 1924). Differently from this, here we use a simpler linear potential function (Equation 7.3) in order to simplify the analytical calculations. Our initial experiments performed with the original Lennard Jones potential showed no substantial qualitative differences with the linear potential function, provided that particles are initially close to the equilibrium positions (all particles placed close to the desired distance).

Noise model To compare our work with previous studies in statistical physics (Vicsek et al., 1995; Grégoire and Chaté, 2004), we consider two different noise models (explained in Section 7.2), both of which are different from the uniformly distributed noise considered in previous chapters.

Tools To analyze the dynamics of the system, we used standard and non-standard statistical physics tools such as phase-transition analysis (Vicsek et al., 1995), non-linear elasticity theory (Fetter and Walecka, 2003), and stability analysis (Ogata, 2001). For numerical simulations, differently from the previous chapters, we did not use the ARGoS simulator, but we use and extended a simple multi-particle simulator¹. Also, in this chapter we measure quantities, such as the potential energy, the kinetic energy and the binder cumulant, that were not considered in the robotics chapter.

The chapter is organized as follows. In Section 7.2, we introduce our statistical physics model. In Section 7.3, we characterize its typical dynamics and stationary solutions, focusing on the self-organization process that leads to flocking and on the order-disorder transition observed at the critical noise. In Section 7.4, we describe the energy cascading mechanism responsible for collective motion in the model and characterize its convergence dynamics. Section 7.5 presents an analytical linear stability calculation that provides a necessary condition for a model to sustain flocking. Section 7.6 explores three examples of possible dynamics achievable with the model: an elastic ring-shaped configuration with oscillating radius, the propagation of an angular perturbation on a group of aligned particles, and a variation of

¹Cristián Huepe - Computer Programs, <http://people.esam.northwestern.edu/~cristian/>, March 2013

the model where each particle has a different self-propulsion speed. Finally, in Section 7.7 we conclude with a summary and a discussion.

7.1 Problem description

From the perspective of statistical physics, a flocking process is seen as a non-equilibrium system. In flocking, as in many other processes typical of living systems, energy is injected at the smallest scale, at the level of each individual. Due to this very nature, and despite intense recent research activity, there is still no comprehensive understanding of the underlying mechanisms that can lead groups of self-propelled agents to self-organize and move in a common direction.

The prevailing paradigm in the theory of flocking in statistical physics has been strongly influenced by the seminal work of Vicsek et al. (1995), who introduced a minimal model for flocking, the *Vicsek model*, that has become a referent in the field (Deutsch et al., 2012; Vicsek and Zafeiris, 2012). The model describes a group of particles advancing at a fixed common speed that interact with each other via an alignment rule equivalent to alignment control: each particle is steered towards the average direction of motion of all particles within a given radius (Vicsek et al., 1995; Czirók et al., 1997). As the amount of noise is decreased, the system undergoes a dynamical phase transition at a critical noise level, below which particles self-organize and start moving in a common direction. Within Vicsek's framework, a swarm can be viewed as a fluid of self-propelled spins with aligning interactions, which can be described as an extension of the XY-model (Binney et al., 1992) where spins advance in their pointing direction rather than remaining affixed to a lattice. In the same spirit of the Vicsek model, but using a continuous description rather than particle-based simulations, Toner and Tu introduced a hydrodynamic theory of active fluids (Toner and Tu, 1995, 1998).

While most of the work on flocking has focused on systems with explicit alignment (explicit information transfer), some studies have considered particles that do not explicitly exchange information on their orientation (implicit information transfer). Examples, already reviewed in Section 2.4.2, include Romanczuk et al. (2009), where flocking is driven by escape-pursuit interactions only; Grossman et al. (2008), where it is driven by inelastic collisions between isotropic particles; Szabó et al. (2006) and Henkes et al. (2011), where it is driven by short-range radial forces that are coupled to the particles' turning dynamics; Menzel and Ohta (2012), where it is achieved due to local pairwise repulsive interactions between soft deformable self-propelled spherical particles. Given that Vicsek-like algorithms rely on explicit alignment rules to achieve flocking (Vicsek et al., 1995; Couzin et al., 2002; Grégoire et al., 2003; Grégoire and Chaté, 2004), it was initially surprising that other models could self-organize without such rules. One can argue that these models must include an indirect effect that produces an effective aligning interaction between individual particles in

order to achieve flocking, thus reducing the dynamics again to explicit information transfer. An example of these aligning dynamics is found in [Grossman et al. \(2008\)](#), where inelastic collisions tend to align post-collision trajectories due to the conservation of momentum. Despite these studies, it remains unclear if the same underlying mechanism is responsible for the emergence of flocking in all these cases and whether or not the mechanism leading to flocking must involve explicit or implicit aligning interactions.

In this chapter, we analyze the mechanism for flocking that is at the basis of MDMC. We explain here the mechanism using a novel paradigm: the emergence and growth of regions of coherent motion due to standard elasticity processes ([Fetter and Walecka, 2003](#)). We explore this mechanism by introducing a simple two-dimensional Active Elastic Sheet (AES) model where the individual particle motion is determined by attraction-repulsion rules only (implicit information transfer) and no orientation information is exchanged between particles (no explicit information transfer). Rather than considering an *active fluid* where particles can flow with respect to each other ([Toner and Tu, 1995, 1998](#)), we consider here a two-dimensional *active solid* or *active crystal* by describing a configuration of self-propelled particles that act as an elastic lattice, where interacting neighbours remain coupled by linear elastic forces throughout the dynamics, regardless of the amount of strain in the system.

7.2 Active Elastic Sheet model

We consider a system of N particles moving on a two-dimensional plane. The position \mathbf{x}_i and orientation θ_i of each particle i are governed by the following set of overdamped² dynamical equations

$$\dot{\mathbf{x}}_i = U \mathbf{o}_i + K_1 [(\mathbf{F}_i + \xi_r \mathbf{n}_r) \cdot \mathbf{o}_i] \mathbf{o}_i, \quad (7.1)$$

$$\dot{\theta}_i = K_2 [(\mathbf{F}_i + \xi_r \mathbf{n}_r) \cdot \mathbf{o}_i^\perp] + \xi_\theta n_\theta. \quad (7.2)$$

Here, U is the forward biasing speed that induces self-propulsion (by injecting energy at the individual particle level) and parameters K_1 and K_2 are inverse translational and rotational damping coefficients, respectively. Unit vector \mathbf{o}_i points parallel to the orientation of particle i and unit vector \mathbf{o}_i^\perp points perpendicular to it. The total force over particle i is given by the sum of linear spring-like forces³

$$\mathbf{F}_i = \sum_{j \in \mathcal{N}_i} -\frac{\epsilon}{\sigma_{ij}} (|\mathbf{x}_i - \mathbf{x}_j| - \sigma_{ij}), \quad (7.3)$$

²We consider a dynamical equation as overdamped when its damping coefficient are very high, so that it can effectively be considered as a kinematic equation.

³We use \mathbf{F} instead of \mathbf{p} as in the previous chapters since here we are using linear, spring-like forces rather than the non-linear ones obtained from the Lennard-Jones potential.

with equilibrium distances σ_{ij} and stiffness ϵ/σ_{ij} . We chose to define the stiffness of each spring as inversely proportional to its natural length σ_{ij} in order to mimic the elastic response of equivalent physical springs of different lengths. Each set \mathcal{N}_i contains all particles that interact with particle i . All \mathcal{N}_i sets are chosen at $t = 0$ from the local neighbourhood of particle i and remain constant throughout the integration. As in MDMC, for zero noise ($\xi_r = \xi_\theta = 0$), each particle simply turns at a rate proportional to the projection of elastic forces perpendicular to its orientation and moves forward or backwards driven by the projection of these forces parallel to the direction of its orientation and by the self-propulsion term U . We introduce *sensing noise* (errors in the measured forces) by adding $\xi_r \mathbf{n}_r$ to \mathbf{F}_i , where ξ_r is the noise strength coefficient and \mathbf{n}_r is a randomly oriented unit vector. We introduce *actuation noise* (fluctuations in the particle dynamics) by adding $\xi_\theta n_\theta$ to the orientation of each particle, where ξ_θ is the noise strength coefficient and n_θ is a normally distributed random variable. We also tested cases where ξ_θ followed a uniform distribution in the $[-\pi, \pi]$ interval, finding equivalent results (data not shown).

The AES model presents many differences with the Vicsek model. First, in the Vicsek case particles exchange information on their relative orientation, while in the AES model they only know the relative positions of their neighbours. Second, the Vicsek model displays shuffling of particles that allow them to interact with different neighbours over time. In fact, long range order cannot be achieved by the Vicsek dynamics if particles always interact with the same neighbours (Czirók et al., 1997; Toner and Tu, 1995). This result is established for equilibrium systems by the Mermin-Wagner theorem (Mermin and Wagner, 1966) and has been shown to extend to this nonequilibrium case: even relatively small systems do not display flocking if interactions are only local and neighbours do not change over time (Aldana and Huepe, 2003). In contrast, we show below that the AES model achieves flocking despite having particles that interact with a fixed set of neighbours. Finally, while both models describe overdamped systems, in the AES model the angular equation of motion (7.2) gradually changes the orientation instead of instantaneously switching it to the next desired orientation (Vicsek et al., 1995; Grégoire and Chaté, 2004; Grégoire et al., 2003). We show below that this, apparently small, difference is essential for achieving flocking in the AES model.

7.3 Numerical dynamics and stationary solutions

In this Section we present numerical simulations that characterize the typical dynamics and stationary solutions of the AES model.

We integrated Eqs. (7.1) and (7.2) numerically using a standard Euler method, which

yields the expressions

$$\mathbf{x}_i^{t+1} = \mathbf{x}_i^t + \left\{ U \mathbf{o}_i + K_1 \left[\left(\mathbf{F}_i^t + \frac{\xi_r}{\sqrt{\Delta t}} \mathbf{n}_r \right) \cdot \mathbf{o}_i \right] \mathbf{o}_i \right\} \Delta t, \quad (7.4)$$

$$\theta_i^{t+1} = \theta_i^t + \left\{ K_2 \left[\left(\mathbf{F}_i^t + \frac{\xi_r}{\sqrt{\Delta t}} \mathbf{n}_r \right) \cdot \mathbf{o}_i^\perp \right] + \frac{\xi_\theta}{\sqrt{\Delta t}} n_\theta \right\} \Delta t, \quad (7.5)$$

where Δt is the numerical time-step. Note that ξ_r and ξ_θ are divided here by $\sqrt{\Delta t}$ in order to properly take account of the accumulation of noise over time (Kloeden and Platen, 2011). The degree of alignment in the system is monitored by computing the usual polarization order parameter, defined as⁴.

$$\psi = \frac{1}{N} \left\| \sum_{i=1}^N \mathbf{o}_i \right\|. \quad (7.6)$$

If all particles are perfectly aligned, we have $\psi = 1$; if they are instead randomly oriented, we have $\psi = 0$.

All simulations are carried out using: $K_1 = 0.01$, $K_2 = 0.12$, $U = 0.002$, and $\Delta t = 0.1$. We tested other parameters without finding any significant qualitative difference in the resulting dynamics.

Figure 7.1 presents three different simulation runs of the AES model. Row A displays the dynamics of $N = 91$ particles forming an hexagonal active crystal. This configuration is the equivalent to the one robots achieve using the proximal control method described in Section 3.3. At time $t = 0$ (panel A1), particles are placed with random orientation and separated by $d_A = 0.65$ on a perfect hexagonal lattice. Nearest neighbours are connected by springs of natural length $\sigma = d_A$ and spring constant $\epsilon/\sigma = 5/0.65$. We include here only sensing noise by setting $\xi_\theta = 0$ and $\xi_r = 0.5\sqrt{0.1} \approx 0.158$. Results remain qualitatively unchanged for other types of noise. As time advances, individual self-propulsion deforms the lattice, producing elastic forces that in turn affect the particle dynamics. Growing regions of coherent motion develop, eventually deforming the whole structure (A2 and A3) until the group starts translating or rotating collectively. The displayed case converges to a rotating state (A4) with an axis of rotation that does not coincide with its barycenter. The group therefore translates while rotating. As it will better explained in the following, the rotating state has higher elastic energy compared to the translating one, since the inner and outer shells cannot move at the same U and must be sped up or slowed down by elastic forces. Consequently, the rotating case shown on panel A4 is metastable and less frequently observed. If we integrate long enough and with high enough noise, it will eventually relax to the lower-energy, translating state. We chose to display this case on the figure, however, to illustrate its rich rotating dynamics, which cannot be attained by the Vicsek model.

Row B presents the dynamics of an elastic rod comprised of $N = 118$ particles arranged

⁴This is equivalent to the order metric defined in Section 4.2 and used in Chapter 4 and Chapter 6

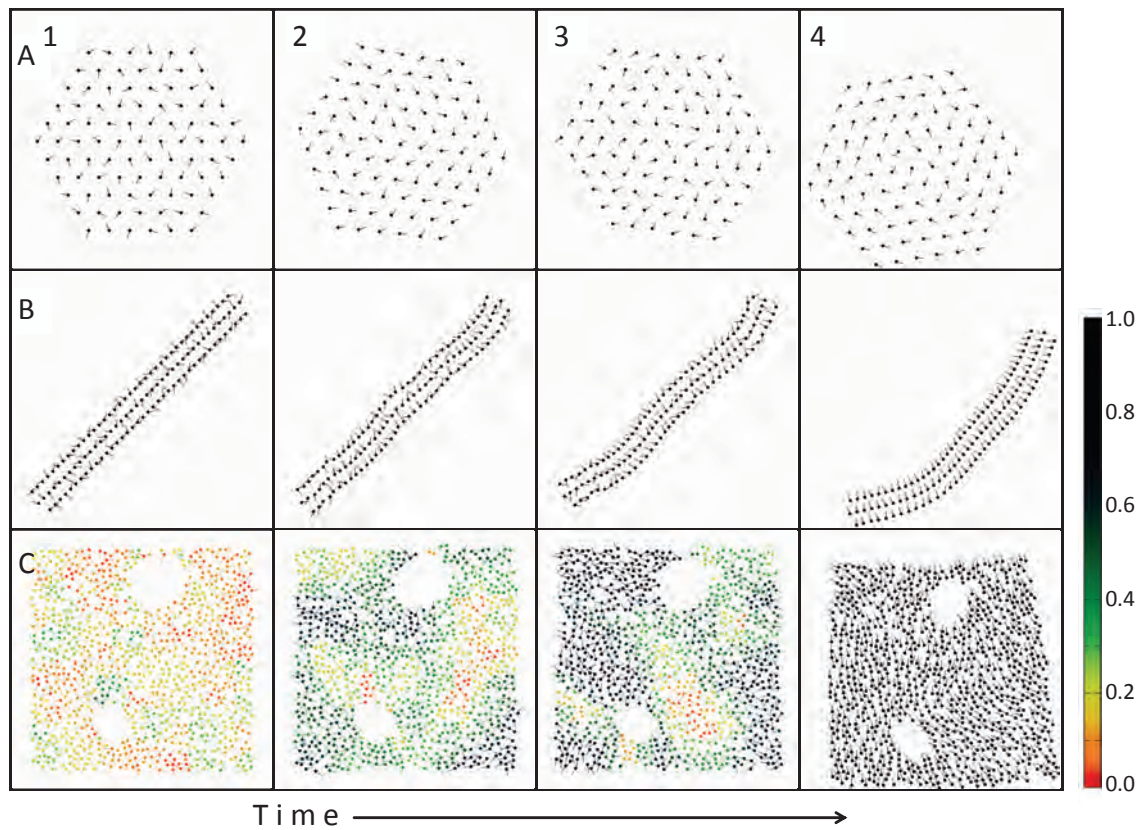


Figure 7.1: Snapshots of three different active elastic sheet simulations obtained by integrating Eqs. (7.4) and (7.5). A: Hexagonal active crystal at $t = 0$ (A1), 240 (A2), 400 (A3), and 1700 (A4). B: Rod-like active crystal at $t = 0$ (B1), 240 (B2), 400 (B3), and 1700 (B4). C: Arbitrarily shaped active solid (here a square with two holes) at $t = 0$ (C1), 240 (C2), 400 (C3), and 1300 (C4). In row C, each particle is coloured according to the degree of local alignment of its local neighbourhood, with $\psi^{\text{loc}} = 1$ indicating full alignment and $\psi^{\text{loc}} = 0$, no alignment.

into three rows. At $t = 0$ (B1), randomly oriented particles are positioned with distances to nearest neighbours $d_B = 0.32$ (within each row) and $d_B^* = 0.58$ (between rows). All particles separated by a distance $d < 1$ are then linked by springs of natural length d and spring constant $\epsilon/\sigma = 1/d$. Noise is the same as in row A. Here again, after the initial transient larger and larger regions of coherent deformation emerge (B2 and B3) until flocking is reached and the rod starts moving (B4). By computing the spectral decomposition of this elastic structure (more details in Section 7.4.1), we find that the first elastic mode, that is, the direction of maximum relaxation of the elastic lattice, is parallel to the rod's axis. This first mode has here the largest final deformation (Row B of the figure, last two panels), thus favouring a collective direction that is perpendicular to the rod's axis. This observation opens the interesting possibility of controlling the direction of self-organized flocking by changing the shape of the particle formation. This could be achieved in swarm robotics by using attraction-repulsion rules that are able to control the shape of robot formation (Barnes et al., 2009).

Row C displays $N = 891$ particles forming an arbitrary square-like structure with two holes. We refer to this case as an active solid due to the lack of regularity in particle positions. This case is only relevant to statistical physics as replicating it with robots would need them capable of identifying neighbors and to use different attraction-repulsion rules depending on the neighbor. Here, particles are initially distributed uniformly within the predetermined shape of the structure, with random positions and orientations. All particles separated by $d < 1$ are linked by springs of natural length d and spring constants $\epsilon/\sigma = 5/d$. Noise is set to zero, but we verified that the same qualitative dynamics is observed for small enough ξ_r and ξ_θ values. To highlight the ordered regions, we define the measure of local order

$$\psi_i^{\text{loc}} = \frac{1}{|\mathcal{N}_i|} \left\| \sum_{j \in \mathcal{N}_i} \mathbf{o}_j \right\|, \quad (7.7)$$

where $|\mathcal{N}_i|$ denotes the cardinality of set \mathcal{N}_i . If particle i and all the neighbours it interacts with are aligned, then $\psi_i^{\text{loc}} = 1$; if they are oriented in different directions, then $\psi_i^{\text{loc}} = 0$. We colour each particle in the elastic solid according to its ψ_i^{loc} value, following the scale displayed on the figure. At $t = 0$, particles are randomly oriented and ψ_i^{loc} values are typically small (C1). As time advances and the elastic sheet deforms, aligned regions of coherent motion (with ψ_i^{loc} close to 1) appear and grow (C2 and C3). Finally, the whole structure starts moving collectively when particles become sufficiently aligned (C3).

The AES model displays an order-disorder phase transition as a function of noise similar to that in the Vicsek model. Figure 7.2 examines this transition in the same hexagonal active crystal displayed on row A of Fig. 7.1 and in a larger ($N = 547$) hexagonal configuration with identical parameters. The leftmost column presents results as a function of sensing noise ξ_r and the central column, as a function of actuation noise ξ_θ . Top panels display the




Figure 7.2: Global order parameter ψ and Binder cumulant G as a function of positional sensory noise ξ_r (leftmost column) and angular actuation noise ξ_θ (central column) for the hexagonal active crystal with $N = 91$ particles displayed on Fig. 7.1 (panel A1) and for an equivalent but larger hexagonal system with $N = 547$ particles. The curves display the mean and local maxima of the distribution of values obtained in numerical simulations. The rightmost column shows the ψ distributions close to the critical noise. Both transitions are shown to be first order, displaying a discontinuity at the critical point and a bistable region.

mean and local maxima of the distribution of ψ values computed for the last 10^6 time-steps (after discarding the initial 10^6 steps to ensure convergence to a statistical steady state) in 30 (or 80, in the transition region) equivalent runs per noise value. Bottom panels display the corresponding values of the Binder cumulant (Grégoire and Chaté, 2004)

$$G = 1 - \langle \psi^4 \rangle / 3 \langle \psi^2 \rangle^2,$$

⁵ which is used to detect whether a distribution is unimodal or has more than one mode (in which case $G < 0$). As the level of either type of noise is increased, the system undergoes a discontinuous transition from an ordered state where particles self-organize, to a disordered state where they continue to point in random directions without achieving flocking. This is evidenced by the discontinuous drop of the order parameter at the critical point and by the Binder cumulant, which is known to become negative in the transition region for first order transitions with bistable solutions. In the sensing noise case, we find that G reaches negative

⁵ $\langle \psi^2 \rangle$ and $\langle \psi^4 \rangle$ denote the second and fourth moment of the distribution of ψ , respectively

values for $N = 547$, indicating that the transition is discontinuous for large enough systems. In the actuation noise case, $N = 547$ does not appear to be large enough to reach $G < 0$, but the dip at the transition region drops further and further below $G = 1/3$ (the expected value in the disordered phase) as the system size is increased, which strongly suggests a discontinuous transition (Binder, 1997; Grégoire and Chaté, 2004; Chaté et al., 2008). We confirm the presence of a bistable region (where ordered and disordered solutions coexist) for both cases by displaying on the rightmost column the distribution of ψ values in the transition region. It is apparent that these distributions become more bimodal as the system size is increased. Finally, we point out that we observed an equivalent discontinuous transition when using either Gaussian or uniformly-distributed noise distributions, for other spatial configurations (data not shown) as well as for much bigger system sizes, up to about 100000 particles. From the swarm robotics perspective, this suggests that the AES model, and consequently also MDMC, allows robots to achieve flocking in a common direction even for very large swarms, and we can use the above results to predict how robust sensors should be designed depending on the noise the system can sustain at a given scale.

7.4 Convergence dynamics

We focus in this section on the convergence dynamics of the AES model. First, we show that the mechanism that leads to flocking can be best understood by decomposing the energy of the system into its elastic modes. Then, we study how the convergence to the ordered state depends on the system size.

7.4.1 Energy cascading mechanism

We begin by computing the spectral decomposition of the energy into the elastic modes of the system. In order to do this, we first numerically determine the elasticity matrix K of the structure. We then find the eigenvalues and eigenvectors of the K matrix. As in standard elasticity, these eigenvectors, also called elastic modes, define an orthogonal base over which we can decompose the relative displacements (with respect to the equilibrium configuration) and the velocities. The projected displacements and velocities can then be used to do the spectral decomposition of the potential and of the kinetic energy, respectively.

We display in Fig. 7.3 the dynamics of the total kinetic and potential energy (panels A and B), and of the spectral decomposition of the potential energy (panel C) for the same system simulated in row A of Fig. 7.1, but with zero noise ($\xi_r = \xi_\theta = 0$) and for a run that converges to a translating solution. Note that the sum of potential plus kinetic energy is not conserved here due to the overdamped nature of the dynamics. For this system, we have 182 elastic modes, corresponding to 91 particles with two positional degrees of freedom per particle. The amplitude of these modes are displayed on Fig. 7.3-C as a function of time, numbered in




Figure 7.3: Kinetic energy (A), elastic energy (B), and spectral decomposition of the elastic energy (C) as a function of time for a zero noise simulation of the hexagonal active crystal displayed on Fig. 7.1 (panel A1) that converges to an aligned, translating state. Brighter points in C represent larger mode amplitudes. After a transient, both energies converge to their stationary values for collective translational motion, $E_K^* = 1.82 \times 10^{-4}$ and $E_V^* = 0$, respectively. All modes show decaying oscillations that dampen out faster for higher modes. After $t \approx 1500$, most elastic energy has become kinetic energy, where it either dissipates or flows to lower modes, eventually reaching the zero (translational) mode.

order of growing energy and smaller scales, without accounting for degeneracies. The initial condition is set with all particles randomly oriented and placed in an undeformed hexagonal lattice. At $t = 0$, the potential energy⁶ is therefore

$$E_V^* = \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N \frac{1}{2} \frac{\epsilon}{\sigma_{ij}} (|\mathbf{x}_i - \mathbf{x}_j| - \sigma_{ij})^2 = 0,$$

since $(|\mathbf{x}_i - \mathbf{x}_j| - \sigma_{ij}) = 0$ (particles at the equilibrium), while the kinetic energy is equal to

$$E_K^* = \frac{1}{2} NU^2 = 1.82 \times 10^{-4},$$

where we have set the particle mass to 1. At the beginning of the dynamics, kinetic energy drops and potential energy grows, due to elastic forces. Given the disorder in the system,

⁶Here, we are simply summing all potential energies between pairs of agent, each equal to $\frac{1}{2} \frac{\epsilon}{\sigma_{ij}} (|\mathbf{x}_i - \mathbf{x}_j| - \sigma_{ij})^2$. In doing so, we need to further divide by 2 in order to take into account the contribution of the pairs i, j and j, i only once.

this potential energy is initially broadly distributed throughout the different energy modes. As time advances, the system rearranges itself into states with decaying elastic energy and growing kinetic energy, until the former reaches values close to zero while the latter reaches again its E_K^* value. Figure 7.3-C allows us to visualize the mechanism that leads to self-organization. After the initial transient, all modes are excited and their amplitude oscillates while dampening, with higher modes decaying faster than lower ones. This results from a combination of a standard elasticity process and the coupling between elastic forces and orientation imposed by the model. Indeed, it is well-known that higher energy modes dampen at a faster rate in elastic systems, since they are more rigid and have intrinsic higher oscillation frequencies, which typically leads to faster dissipation. In this active system, however, each particle is continuously injecting energy at the individual level through its self-propulsion term, so motion cannot dampen out. Instead, elastic forces will steer particles away from higher modes more strongly than from lower modes. If while doing this particles do not re-excite higher elastic modes faster than these decay, the self-propulsion energy will be channeled to lower and lower modes until the first (rotational) mode or zero (translational) mode is reached and flocking is achieved. If instead particles feed too much self-propulsion energy to higher modes while turning, these modes will always remain excited and no flocking state will be reached. We conclude that not every active elastic system will achieve flocking. For example, we consider in Section 7.5 a constant-speed algorithm that does not display flocking despite having the same angular dynamics as the AES model. In the swarm robotics perspective, such an analysis can be used in order to determine the conditions under which a motion control method can produce self-organized flocking.

7.4.2 Dependence on system size

We study here how the convergence dynamics depends on the system size. Figure 7.4 displays the global order parameter ψ as a function of time for hexagonal active crystals with $N = 91, 547, 1027,$ and 5167 particles. Each simulation is started from a random initial condition, using the same parameters as in Fig. 7.3. Ten convergence curves are presented per system size. We observe that for $N = 91$ not all runs converge to the aligned ($\psi \approx 1$) state. Instead, four runs reach the metastable rotating ($\psi \approx 0$) state and remain trapped there until the end of our simulation time, which was set here at $t = 10^4$ (much longer than the displayed time frame). For the other sizes, however, no run ultimately converges to the rotating state. This is because the larger a rotating structure is, the faster outer particles must advance in order to maintain cohesion. For large enough systems, the drag introduced by these particles will be enough to destabilize the rotating solution.

Figure 7.4 also shows that convergence times have a large variability. Despite this, the figure readily provides a rough estimate of how these times scale with the system size. Indeed, the time frame displayed on each panel is proportional to $N^{1/2}$ (i.e., $\sim 100N^{1/2}$). Given

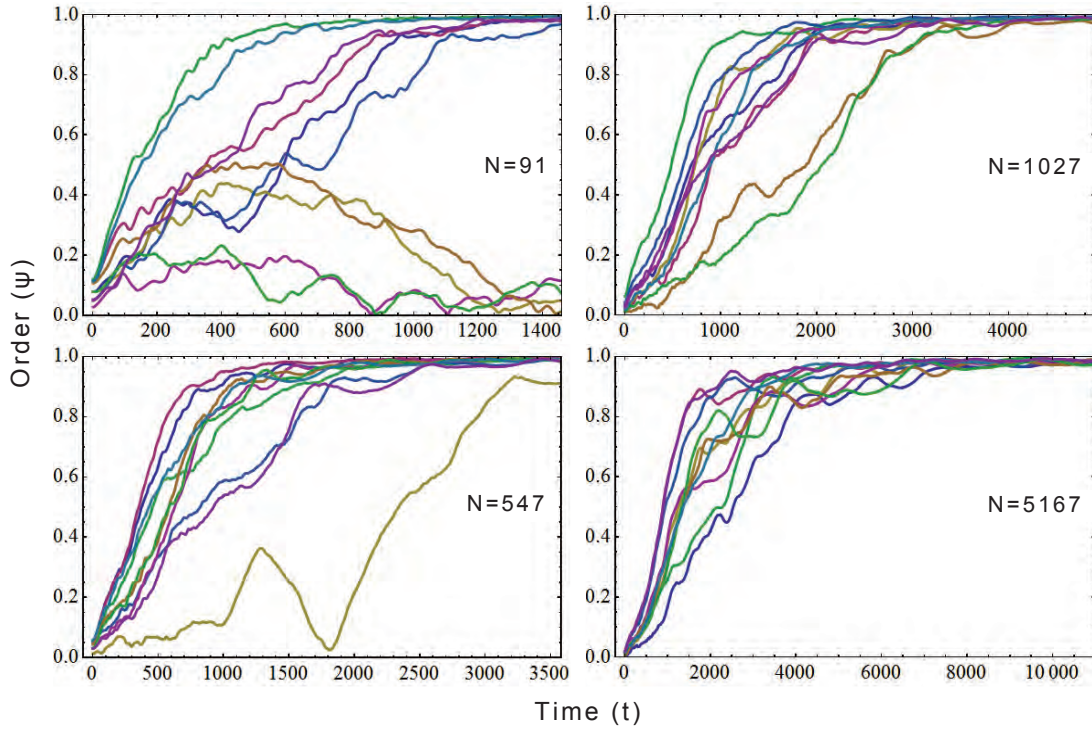


Figure 7.4: Convergence dynamics of the order parameter ψ for four hexagonal active crystals of different sizes. All parameters are the same as in Fig. 7.3. Ten runs with different, randomly oriented initial conditions are displayed for each size. When ψ approaches 1, the system is converging to an aligned translating state and when it approaches 0, to a rotating (metastable) state. All system sizes display a broad variability of convergence times, with larger systems typically taking longer to converge.

that even with this rescaled temporal axis curves seem closer to the ordinate axis for higher N , it is apparent that the typical convergence time grows here slower than $N^{1/2}$. The fact that the convergence dynamics is strongly non-monotonous suggests that its variability is a reflection of the complex dynamical landscape that the system navigates, where it can spend unpredictable amounts of time near local attractors and metastable states. Since the main metastable state is the rotating solution, a systematic scaling analysis of convergence times will require not only a much larger set of simulation runs, but also structures that suppress the rotating state, such as strongly elongated shapes. From the swarm robotics perspective, this suggests that controlling the shape of the swarm to suppress the rotating states can be a useful tool to keep convergence times bounded.

7.5 Linear stability analysis

One of the interesting aspects of our AES model is that we can use a continuous elastic sheet approximation to carry out analytical calculations. We follow this approach and perform a standard linear stability analysis (Fetter and Walecka, 2003) of the zero noise case to investigate which specific dynamical rules can sustain translating flocking solutions. We begin by writing the elastic forces $\mathbf{F} = (F_x, F_y)$ that result from small displacements $\mathbf{v} = (v_x, v_y)$ of points on the lattice with respect to their equilibrium position. These are given by the standard elasticity equations

$$F_x = (\lambda + 2S) \frac{\partial^2 v_x}{\partial x^2} + S \frac{\partial^2 v_x}{\partial y^2} + (\lambda + S) \frac{\partial^2 v_y}{\partial x \partial y}, \quad (7.8)$$

$$F_y = (\lambda + 2S) \frac{\partial^2 v_y}{\partial y^2} + S \frac{\partial^2 v_y}{\partial x^2} + (\lambda + S) \frac{\partial^2 v_x}{\partial x \partial y}, \quad (7.9)$$

where the elastic constants are the Lamé parameter λ and shear modulus S (Fetter and Walecka, 2003). By linearizing the equations of motion (7.1-7.2) around an equilibrium solution with undeformed lattice and all particles moving at speed U in the \mathbf{x} direction, we find the following expressions for v_x, v_y , and the perturbation field ζ of the orientation

$$\dot{v}_x = K_1 F_x, \quad \dot{v}_y = U \zeta, \quad \dot{\zeta} = K_2 F_y. \quad (7.10)$$

Replacing Eqs. (7.8) and (7.9) into (7.10) and casting the resulting expression in Fourier space with wavevector components k_x and k_y , we can express the perturbation dynamics in matrix form and compute its eigenvalues Λ to determine the linear stability of the system. We find that Λ satisfies the characteristic equation $\Lambda^3 + C_2 \Lambda^2 + C_1 \Lambda + C_0 = 0$, with

$$C_0 = K_1 K_2 S U (\lambda + 2S) [k_x^2 + k_y^2]^2 \quad (7.11)$$

$$C_1 = K_2 U [S k_x^2 + (\lambda + 2S) k_y^2] \quad (7.12)$$

$$C_2 = K_1 [(\lambda + 2S) k_x^2 + S k_y^2]. \quad (7.13)$$

Using Routh's stability criterion, here given by $C_1 C_2 > C_0$ (Ogata, 2001), we find that the system will be linearly stable if $K_1 K_2 U (\lambda + S)^2 k_x^2 k_y^2 > 0$, which is always verified. We conclude that the translating flocking solution is linearly stable for all parameter values.

Analytical calculations like those presented above allow us to determine, a priori, which elasticity-based equations of motion will be able to sustain flocking. We found that most variations of the AES model cannot support stable aligned solutions. Consider, for example, an algorithm where the orientation is determined by Eq. (7.2) but the forward speed is set constant to U ($K_1 = 0$). Given that the angular dynamics remains unchanged, one could naively think that this system will align like the AES model. We will now show, however,

that this is not the case. The characteristic equation for $K_1 = 0$ becomes

$$\Lambda^3 + K_2 U [S k_x^2 + (\lambda + 2S) k_y^2] \Lambda = 0, \quad (7.14)$$

which has solutions $\Lambda = 0$ and $\Lambda^2 = -UK_2 [Sk_x^2 + k_y^2(\lambda + 2S)]$. Since $\Lambda \in i\mathbb{R}$, linear perturbations will not dampen out, but produce instead permanent oscillations. We confirmed through numerical simulations that, even for zero noise and starting from an aligned initial condition, the group will lose order as particles rotate in place. In the context of robotics, this is exactly the same dynamics we observed in Chapter 6 using the MIMC motion control method, in which the forward speed u is constant. After testing several other models, we found only one other example that can sustain flocking: a variation of the model introduced in Szabó et al. (2006). However, this model is not applicable to non-holonomic robots, as it requires robots or particles to translate sideways. We conclude that, among the models and methods we tested, the AES model (and consequently the MDMC motion control method) is the only method capable of achieving flocking.

7.6 Exploring Active Elastic Sheet dynamics

We carry out in this section an initial exploration of three different dynamical setups of the AES model that allow us to better understand its typical behaviour and modelling possibilities. We first consider a rotating ring solution, then the propagation of perturbations on an aligned rod-like configuration, and finally a variation of the AES model where each particle has a different self-propulsion speed U . With the exception of the third one, the analysis in this section has been performed mainly to explore the dynamics of the AES model within the statistical physics context, and presents non-intuitive interpretations in the context of swarm robotics.

7.6.1 Dynamics of a ring-shaped configuration

Figure 7.5 displays the dynamics of $N = 100$ particles in a ring-shaped structure where nearest neighbours are connected by springs of natural length $\sigma = 0.65$ and stiffness $\epsilon/\sigma = 0.25/0.65$. Particles are initially placed on a circular configuration, separated by σ and pointing tangentially (clockwise). The system is then integrated forward in time with zero noise. Top panels display snapshots at $t = 5135$, $t = 7655$, $t = 10^5$, and $t = 4 \times 10^5$; the bottom panel shows the mean radius of the ring as a function of time. Initially, the ring expands, increasing elastic forces until a maximum radius is reached (panel 1) and the ring starts contracting. The contraction speed then increases until it reaches a maximum (panel 2) and particles start turning outwards until they move again tangentially, reaching a minimum radius. This breathing mode continues to oscillate with decaying amplitude until it fully dampens out.




Figure 7.5: Simulation of a ring-shaped active elastic system. Top panels display snapshots of the dynamics at $t = 5135$ (maximum radius, panel 1), 7655 (next maximal contraction speed, panel 2), 100000 (panel 3), and 400000 (panel 4). The arrow on panel 1 indicates the sense of rotation of the structure and the circle on all panels, the shape of the initial condition. At $t = 0$, particles are placed on this circle, pointing tangentially (clockwise). The bottom panel presents the mean radius of the structure (with respect to its barycenter) as a function of time. After the initial oscillations of the breathing mode dissipate, higher elastic modes develop and deform the circular structure.

The circular configuration then loses stability, exciting higher modes that deform it (panel 3). This state survives with different levels of deformation until the end of our integration time (panel 4).

7.6.2 Propagation of perturbations

Figure 7.6 shows the propagation dynamics of a local perturbation of the orientation on an aligned flocking state. We set up a three-row rod-like structure similar to that on Fig. 7.1-B, with the same parameters but in a longer configuration with $N = 499$ particles. All particles are initially placed aligned and pointing on the same direction as the rod axis, which we define as $\theta = 0$. At $t = 0$, we perturb the orientation of the frontmost particle, rotating it by $\pi/18$ (small perturbation) or $\pi/2$ (large perturbation). The system is then integrated forward in time with zero noise. We plot the orientation of all particles on the central row of the rod-like structure at four different moments in time, indexed in order of their position from back to front. Both small and large perturbations display here a wake of persistent

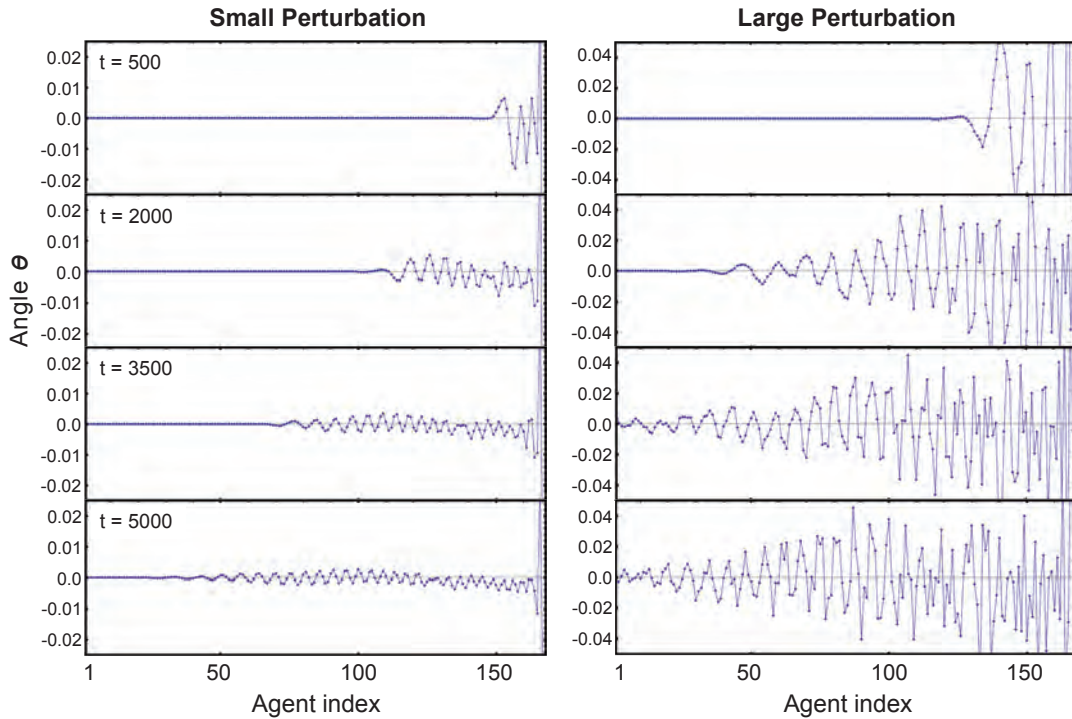


Figure 7.6: Propagation of localized angular perturbations of the aligned state. Particles are initially placed on a three-row rod-like structure similar to that on Fig. 7.1-B but longer (containing $N = 499$ particles), with all particles oriented in the same direction ($\theta = 0$) and aligned to the rods main axis. Simulations are carried out with the same parameters used in Fig. 7.1-B, but for zero noise. At $t = 0$, the frontmost particle is rotated by $\pi/18$ (left column) or $\pi/2$ (right column). The plots display the orientation of all 167 particles on the central row, numbered from back to front. The large perturbation propagates faster and produces a wake of longer wavelength than the small one.

angular oscillations behind them. However, note that other preliminary simulations that we have carried out on less elongated structures show this wake rapidly decaying after the passage of the initial perturbation. For the long rod-like case presented here, we observe that small angular perturbations propagate faster and leave a wake of shorter wavelength and smaller amplitude than large ones. These results illustrate the rich dynamics exhibited by the propagation of perturbations in the AES model. Their study will require further systematic analyses that are left for future work.

7.6.3 Heterogeneous self-propulsion speeds

We now consider a variation of the AES model where all particles have different preferred speeds. Figure 7.7 displays the dynamics of an active elastic hexagon with the same parameters as in Fig. 7.1-A, but where instead of fixing all self-propulsion speeds to $U = 0.02$ we

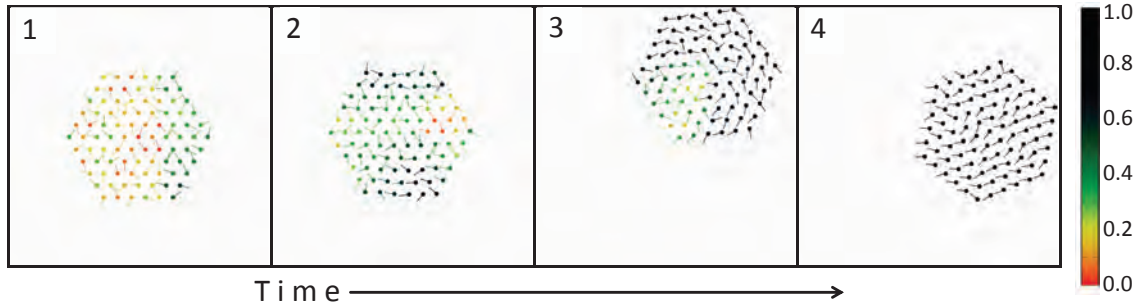


Figure 7.7: Snapshots of a hexagonal active elastic sheet simulation with heterogeneous preferred speeds at $t = 0$ (1), 7300 (2), 12000 (3), and 22000 (4). Each particle's self-propulsion speed U is chosen at random between 0 and 0.004. After a very long integration time, the system reaches a quasi-ordered state of flocking where the orientation of faster particles oscillates broadly so they can remain cohesive with their slower neighbours. Particles are colored based on their degree of local alignment ψ^{loc} , as on Fig. 7.1-C.

select a different U for each particle, at random, from the interval $0 \leq U \leq 0.04$. Particles are coloured based on their local alignment, as in Fig. 7.1-C. Interestingly, even for this highly heterogeneous system particles manage to self-organize and achieve flocking, albeit after a very long relaxation time. As time advances, growing regions of coherent motion emerge (panel 2), but a localized part of the hexagonal structure remains persistently disordered (the lower-left quadrant of the hexagon on panel 3). This is the area that displays the highest differences of preferred speeds, which makes flocking harder to reach. Eventually, the system finds a way to fully self-organize and the groups starts moving collectively. In the case shown, it advances in a curved trajectory due to the random accumulation of faster particles at one side of the structure. A salient feature of the final quasi-ordered state is that it can never reach the stationary solution where all particles are fully aligned. This is because the orientation of faster particles must oscillate strongly so they can remain cohesive with their slower neighbours. The ability displayed here by the AES model to self-organize even in highly heterogeneous situations opens the possibility of constructing a rich variety of active solids that produce diverse collective dynamics by assembling groups of particles with different individual characteristics, or the possibility to control a swarm of heterogeneous robots whose hardware design or control rules allow to exhibit the desired dynamics.

7.7 Summary and discussion

We have identified in this chapter the elasticity-based mechanism at the basis of MDMC, the motion control method we propose for flocking with implicit information transfer, and we introduced the AES model to illustrate it. Up to now, the only existing theoretical framework in statistical physics that explains how systems achieve flocking was either based on

the Vicsek model (Vicsek et al., 1995), which uses explicit information transfer, or on the active hydrodynamic theory introduced by Toner and Tu (1995). Our work develops a very different theoretical framework, providing a simple alternative mechanism for flocking that is based on implicit information transfer and elasticity instead of explicit information transfer or momentum transfer and requires no exchange of orientation information. We found only one other system that can display flocking under similar conditions: the model introduced in (Szabó et al., 2006) to study the collective migration of tissue cells. As we explained in Section 7.5, though, this model is not applicable to non-holonomic robots as it requires robots sliding sideways.

The relevance of aligning interactions for achieving flocking has been a long-standing issue in the statistical physics field. Only a few studies have considered systems where particles align and move collectively without exchanging orientation information, and their underlying mechanism remained unexplained. In fact, seminal work by Grégoire et al. (2003); Grégoire and Chaté (2004) suggested that flocking is not possible in minimal models without aligning interactions. Our studies show that their model does not converge to flocking when aligning interactions are turned off because of two reasons. First, it considers particles with constant speed. As shown above, this is enough to prevent the elasticity-based mechanism from achieving flocking in the AES case, and this is the reason why motion control methods such as MIMC do not achieve flocking in robotics (see Chapter 6). Second, their particles switch to the next desired orientation in one time-step, instead of integrating Eq. (7.2), an assumption that obviously is not valid in robotics. This instantaneous relaxation can stop energy from smoothly flowing to lower energy modes and from developing growing regions of coherent motion.

To the best of our knowledge, this work is the first to combine the theories of elasticity and flocking. An appealing aspect of the AES approach is that it is well-suited for analytical studies. In addition to the stability calculations described in this chapter, we can envision stochastic differential equation analyses determining the critical noise and Kolmogorov-style energy cascading arguments describing the self-organization dynamics. These could in turn help develop our fundamental understanding of flocking and of the more general class of nonequilibrium, self-organizing systems where energy is injected at the smallest scales, which are highly relevant for the design of self-organizing artificial systems such as swarm of robots as well as for the behavioral study of biological systems.

Chapter 8

Conclusions and future work

In this chapter we summarize the main contributions of this dissertation and we present future direction concerning the study of information transfer in large swarms of autonomous robots.

8.1 Contributions

In this dissertation, we performed small but significant steps in the context of flocking, which studies the cohesive and coordinated motion of large swarms of robots. The coordinated motion of robots can lead to several advantages compared to having robots moving individually, such as: improved exploration and data acquisition, promoted by the capability of a swarm to cover larger areas and to share information acquired by individual robots; energy-efficient navigation, promoted by the realization of specific aerial or underwater robot formations (Newton, 2010; Hoare et al., 2000); and fault-tolerant exploration of dangerous locations, promoted by redundancy of hardware and behaviors.

We focused on information transfer, a key concept for the understanding of collective motion (Sumpter et al., 2008) and for its design. We showed that, through the design and the analysis of information transfer, it is possible to obtain collective behaviors (in our case flocking and collective transport) with the desired macroscopic characteristics. We developed design methods to achieve both explicit and implicit information transfer within a swarm of robots. In explicit information transfer, robots share directional information using a communication device; in implicit information transfer, we only require robots capable of sensing each others' relative position and bearing. We thoroughly studied the proposed methods in simulation by comparing them with state-of-the-art methods, and we performed empirical validation on real robots. Additionally, we also analyzed implicit information transfer, whose underlying mechanism is the most complex of the two, using tools of statistical physics.

In the next section, we highlight the other directions that can be taken following the

footsteps of the work carried out in this dissertation.

8.2 Future work

We believe that the study of information transfer is important for better understanding and designing flocking as well as other swarm robotics collective behaviors. We now highlight which future work can be done under this perspective.

Information-theoretic study Information-theoretic metrics have been proposed for measuring information transfer in complex networks ([Lizier et al., 2008](#)) and more recently in swarms ([Wang et al., 2012](#)). These metrics measure the mutual information between the state (in the case of flocking, the orientation) of one agent at a given time and the state of another agent at a time in the future, and they incorporate a sense of directionality to represent the transfer from a source to a target. Using these metrics, it is possible to assess which agents transfer more information, and at which spatial (local, global, scale-free, etc . . .) and time scales the process takes place. One possible investigation involves using information transfer metrics to assess at which spatial and time scales the process takes place for each of the different mechanisms we used in this dissertation and in each environmental conditions. For instance, we hypothesize that information transfer can be maximized and linked to a global (collective) scale when the swarm is close to a phase transition between disordered and ordered motion. [Sumpter et al. \(2008\)](#) has shown this is the case for the collective motion of living organisms, and corresponding to these situations the swarm exhibit its maximal potential in terms of flexibility and robustness.

Automatic design of flocking mechanisms In this dissertation, we used hand-coded design methods to achieve flocking with a swarm of mobile robots. Future research could aim at using automatic design method, such as genetic programming / grammatical evolution ([Ferrante et al., 2013a](#)) or standard evolutionary robotics ([Nolfi and Floreano, 2004](#)), to investigate whether different control components and information transfer mechanisms would lead to optimal swarm-level flocking, and to compare the evolved solutions with the ones considered in this dissertation. It would be interesting also to design an evolutionary robotics experiment in order to answer to the following, more biological question. As reported in Section 2.1.1, biological systems that organize into flocking (such as birds) are composed by non-related individuals that maximize their individual fitness rather than the collective fitness. Thus, it would be interesting to investigate what are the factors that promote the evolution of flocking behaviors given the hypothesis that group fitness (such as the order considered in this dissertation) cannot be maximized, but only relying on the maximization of individual fitness. This involves designing experiments in which robots have to perform some individual tasks,

and investigate for which tasks/conditions collective motion can evolve. One example of this fitness function can be the one that rewards robots that can “catch” other robots, and punishes robots that are “caught” by other robots. This setting would validate the biological system consisting of locusts (Bazazi et al., 2008; Romanczuk et al., 2009), that are shown to evolve into collective motion due to their cannibalistic instinct and escape-pursuit types of interactions.

Automatic design of information transfer Another interesting direction could be to use artificial design method to find individual-level strategies that maximize information transfer. Indeed, Sumpter et al. (2008) and other recent studies (Cavagna et al., 2010) have shown that information transfer can lead to flexibility in moving animal groups, such as to the capability of quickly reacting to local perturbations (for instance a predator) at the level of the swarm. They showed that those systems have this property when information transfer is maximized. Cavagna et al. (2010) also showed that these system have powerlaw-type correlations among the states (i.e., the orientations) of the different individuals in the swarm. From the engineering perspective, it can be very useful to be able to replicate artificially the same levels of flexibility found in natural swarm. In an initial exploration, this could be achieved for instance by using the correlation functions used by Cavagna et al. (2010), or the information-theoretic metrics introduced by Lizier et al. (2008), as functions to be maximized by an artificial design method.

Modeling A further direction of study is modeling the mechanisms used in this dissertation using known theoretical frameworks, such as stochastic differential equations, in order to complement the analysis presented in Chapter 7. As an example, it could be possible to use Langevin and Fokker-Planck equations (see Section 2.2.3) or, as suggested in Chapter 7, Kolmogorov-style modeling. Using the former, it would be possible to model both the mesoscopic (between micro and macro) and the macroscopic process. These models could give us useful predictions in terms of whether, for instance, a new implicit information transfer would produce flocking and under which conditions, or could also predict the convergence times of the different methods. In general, modeling would be useful to further develop our understanding of flocking as a non-equilibrium, self-organizing system, to predict its dynamics under different conditions and to improve our capability to design successful methods.

Information transfer in other collective behaviors We believe that, beyond flocking, information transfer is a central process that can lead to other self-organized collective behaviors with the desired macroscopic properties. To test this hypothesis, information transfer experiment can be performed, for instance, in synchronization and/or in aggregation. In synchronization, robots need to agree on a common phase (as done in

Christensen et al. (2009)) or frequency. In the latter case, we can study a setting in which few informed robots could establish a communication channel by choosing a transmission frequency and imposing it to the rest of the swarm. In aggregation (Garnier et al., 2005), robots need to agree on a common gathering area in the environment. Typically, these experiments consider either a fully informed swarm (Garnier et al., 2005), in which all robots detect the presence of a shelter in the environment, or fully uninformed swarms (Soysal et al., 2007), in which no robot can detect features of the environment but where they have to aggregate in a random location. A possible experiment here would involve designing a swarm in which few individuals can detect the presence of a shelter, as in Garnier et al. (2005), while the majority of the swarm is uninformed and can aggregate only based on social clues, as in Soysal et al. (2007). In this case, a possible research question is how can we design an implicit information transfer strategy in aggregation in order to make the swarm meet at a given goal location perceived only by few informed robots.

Appendix A

The robot and the simulation platform

In this appendix, we describe the tools utilized to carry out the experiments described in Chapter 4, Chapter 5 and Chapter 6.

A.1 The foot-bot and the implementation

Here, we first introduce the hardware used to carry out the experiments. We then explain how the sensors are used to obtain the sensing capabilities needed by the different control components.

A.1.1 The foot-bot

The foot-bot robot is a modular robot composed of many sensors and actuators, collectively referred to as modules. Each module is controlled by a dedicated dsPIC micro-controller. The foot-bot is equipped with a main processor board, a Freescale i.MX31 ARM 11 low-energy 533 MHz processor running linux. The main board features 128 MB of DDR RAM and 64 MB of flash. The dsPICs on the modules communicate with the central processor asynchronously using a common bus and the ASEBA software platform (S. Magnenat, 2011).

Figure A.1 shows a picture of the foot-bot where only the sensors and the actuators used in this dissertation are marked. The foot-bot is 29 cm tall and has a radius of 8.5 cm. Its weight is 1.8 Kg. It uses a lithium polymer battery with very long duration and that can be hot-swapped during an experiment thanks to the presence of a super capacitor. The complete list of robot's sensor and actuators is the following:

- Two differential drive *treels*, a combination of tracks and wheels, are used for locomotion in normal and rough terrains.
- A turret actuator allows a plastic ring with 12 RGB LEDs and a gripper to rotate almost 360 degrees around the Z axis of the robot. The ring is used both to emit light with




Figure A.1: The foot-bot robot with its used sensors and actuators. The distance scanner, the gripper and the turret are only used in collective transport (see Chapter 5) whereas the other three sensors and actuators are used also in flocking (see Chapter 4 and Chapter 6).

different colors and as a docking mechanism for another foot-bot. In fact, the gripper is especially designed to fit this specific plastic ring and to hold into it while open.

- 24 infrared sensors, that are evenly distributed around the foot-bot's body, have two functionalities. First, as a proximity sensor, they can detect obstacles in close range (5 cm). Second, as a light sensor, they can measure the intensity and the direction of the ambient light, even when placed far away from the robot. In our experiments, we use these sensors only as a light sensor.
- 4+8 additional infrared sensors, also referred to as ground sensors, are located underneath the robot, between the two tracks and on the outside part, respectively. They are used to detect the color of the ground in the gray scale.
- Two Pixelplus 2.0 MegaPixels CMOS cameras provide basic visual information. The first camera is located in the above part of the robot, in the center. It point upwards towards a mirror located at the top of a glass tube. This provides the foot-bot with omni-directional vision capabilities. The second camera is also located on the top but in a non-central position, and can be installed to either look upwards towards the ceiling or forward along the X - Y plane.

- A rotating scanner, composed of two short and two long distance scanners, is used to measure distances. It can detect distance to obstacles very precisely, at the expenses of a poor 360 degrees resolution due to the limited rotation speed.
- A range and bearing (RAB) sensing and communication board is the device used for explicit information transfer in this dissertation. It is composed of 20 infrared transmitters, of 12 receivers, and of a radio communication device. The infrared transmitters and receivers have two functionalities. First, they are used to detect the range and the bearing of neighboring foot-bots. Second, they are used to manifest the presence of the focal robot to neighbors, in order to emulate local communication. In fact, the radio mechanism resorts to broadcast communication, which in principle violates swarm robotics principles. However, local communication is achieved by combining radio with infra-red sensors: a radio message is accepted by the focal robot's neighbors only if an infrared beam is also received by the neighbor during the same time slot, as this indicates that the two robots are next to each other. The communication device has a limited bandwidth, but it is relatively noise free.
- A three-axes accelerometer and a three-axes gyroscope installed on the left tree.

A.1.2 Flocking with foot-bots

We explain here how the different algorithms described in Chapter 3 have been implemented on both simulated and real foot-bots.

Proximal control This component is realized using the RAB sensing and communication board to measure the relative range and bearing of the focal robot's neighbors.

Alignment control This component requires the focal robot to sense its own orientation. Since the foot-bot is not equipped with a compass sensor, the same functionality is realized as follows. A light source, placed in a fixed position in the environment, is used to define a reference frame common to all robots. The orientation of a robot is defined as relative to this common light source. Thus, the focal robot measures its own orientation with respect to this reference frame by using the on-board light sensor. It then communicates to its neighbors either its own orientation or a different angle by using the communication unit present in the RAB, depending on the communication strategy that is used (see Section 3.4).

Goal direction In the experiment with informed robots, the goal direction is externally provided and always expressed with respect to the fixed light source.

Note that, in some experiments described in Chapter 6, the light sensor has not been used at all as some robots (non-informed robots that do not use alignment control) do not need this type of information.

A.1.3 Collective transport with foot-bots

In collective transport, the irregularly shaped object that foot-bots must transport is the hand-bot robot (Bonani et al., 2009), which remains passive during the entire duration of the experiment.

We use more sensors and actuators compared to flocking. More precisely, we used: i) the light sensor to measure the focal robot's orientation as in flocking; ii) the distance scanner to obtain the distance and the angle from the focal robot to objects (i.e., obstacles) (Magnenat et al., 2010); iii) the RAB system for explicit information transfer; iv) the gripper to physically connect to the hand-bot to be transported; v) the turret actuator that, when set to active mode, can actively rotate the gripper and that, when set to passive mode, can passively rotate in accordance with the speed of the wheels and the dynamics of the compound; vi) the wheels actuator to independently control the speed of the left and right wheels.

To implement motion control, we considered the robot attached to the left of the hand-bot as the left wheel of the compound system and the robot attached to the right as the right wheel. This assumes that the two robots have always the direction of the wheels axis parallel to each other, and this is achieved by setting the turret to active mode, which blocks the rotational degree of freedom of the gripper around its axis. Hence, the robot attached to the left of the compound sets both wheels speed to N_L (left wheel speed computed by motion control), whereas the robot to the right sets both wheels speed to N_R (right wheel speed computed by motion control). The robot at the center, instead, sets its left and right wheels consistently with the values computed by motion control (N_L and N_R , respectively). Thus, we set the turret of the central robot to passive mode, as to allow its gripper to freely rotate in order to follow the dynamics of the compound and the ones imposed by the wheels.

A.2 The ARGoS simulator

The experiments described in Chapter 4, Chapter 5 and Chapter 6 have been carried out using a simulation tool called ARGoS, which stands for *Autonomous Robots Go Swarming* (Pinciroli et al., 2012). Argos is an open source simulator¹ especially developed in the context of the Swarmanoid project.

The overall architecture of ARGoS is shown in Figure A.2. ARGoS has been developed with two key concepts in mind: *flexibility* and *efficiency*. Flexibility refers to the possibility to tune the experiments according to the needs. Efficiency refers to the possibility to perform experiments with large numbers of robots by keeping satisfactory run-time performance. Flexibility has been achieved through a modular design of the simulator: all components, such as robots, sensors, actuators, physic engines, visualization engines, etc ... are plugins that can be freely selected. It is also relatively easy to design new plugins in order to achieve

¹ARGoS simulator, <http://iridia.ulb.ac.be/argos>, February 2013.

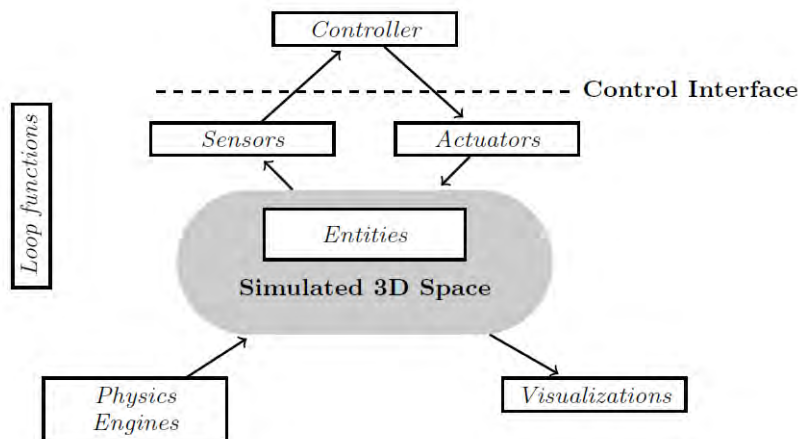


Figure A.2: Architecture of the ARGOS simulator. The image has been extracted from the work of Pincirolì et al. (2012), with permission of the author.

the desired trade-off between realism of simulation and efficiency. In addition, efficiency in ARGOS has been pursued through parallelization, that is, multiple sensors or multiple physics engine can run in parallel on multiple cores. Parallelization of the physics engine has been achieved via partitioning of the simulated space into non-overlapping sub-spaces and assignment of each sub-space to a separate physics engine.

ARGOS has multiple physics engines, spanning from 2D to 3D, from kinematic-based to dynamic-based. In the experiments of this dissertation, a two-dimensional dynamics engine called Chipmunk² has been used.

In ARGOS, it is also possible to seamlessly port controllers from simulation to real robots. This is achieved via a control interface (Figure A.2), which is an abstraction layer between the controller and the sensors/actuators. Cross-compilation is then used to compile the controller on the different platforms: simulation (PC) and real robot. As a matter of facts, all the experiments in this dissertation have been performed with controllers that were identical in simulation and on the real robots.

²Chipmunk physics, <http://code.google.com/p/chipmunk-physics/>, February 2013.

Bibliography

- Aldana, M. and Huepe, C. (2003). Phase transitions in self-driven many-particle systems and related non-equilibrium models: A network approach. *Journal of Statistical Physics*, 112:135–153. [2.4.2](#), [7.2](#)
- Amé, J., Halloy, J., Rivault, C., Detrain, C., and Deneubourg, J. L. (2006). Collegial decision making based on social amplification leads to optimal group formation. *Proceedings of the National Academy of Sciences*, 103(15):5835–5840. [2.2.1](#), [2.2.2](#)
- Ampatzis, C., Tuci, E., Trianni, V., and Dorigo, M. (2008). Evolution of signaling in a multi-robot system: categorization and communication. *Adaptive Behavior*, 16(1):5–26. [2.2.2](#)
- Anderson, C. and Ratnieks, F. (2000). Task partitioning in insect societies: novel situations. *Insectes Sociaux*, 47:198–199. [2.1.3](#)
- Bachrach, J., Beal, J., and McLurkin, J. (2010). Composable continuous-space programs for robotic swarms. *Neural Computation and Applications*, 19(6):825–847. [2.2.2](#)
- Baldassarre, G., Nolfi, S., and Parisi, D. (2003). Evolving mobile robots able to display collective behaviors. *Artificial Life*, 9(3):255–267. [2.2.2](#), [2.4.4](#)
- Baldassarre, G., Parisi, D., and Nolfi, S. (2006). Distributed coordination of simulated robots based on self-organization. *Artificial Life*, 12(3):289–743. [5.1](#), [5.5](#)
- Ballerini, M., Cabibbo, N., Candelier, R., Cavagna, A., Cisbani, E., Giardina, I., Lecomte, V., Parisi, A. O. G., Procaccini, A., Viale, M., and Zdravkovic, V. (2008). Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study. *Proceedings of the National Academy of Sciences, USA*, 105(4):1232–1237. [2.4.1](#), [2.4.1](#)
- Barnes, L., Fields, M., and Valavanis, K. (2009). Swarm Formation Control Utilizing Elliptical Surfaces and Limiting Functions. *IEEE Transactions on Systems, Man, and Cybernetics - Part B*, 39(6):1434–1445. [2.4.4](#), [7.3](#)

- Bazazi, S., Buhl, J., Hale, J. J., Anstey, M. L., Sword, G. A., Simpson, S. J., and Couzin, I. D. (2008). Collective motion and cannibalism in locust migratory bands. *Current Biology*, 18(10):735 – 739. [2.1.1](#), [8.2](#)
- Beckers, R., Deneubourg, J., Goss, S., and Pasteels, J. (1990). Collective decision making through food recruitment. *Insectes Sociaux*, 37:258–267. [2.1.1](#)
- Beer, R. D. and Gallagher, J. C. (1992). Evolving dynamic neural networks for adaptive behavior. *Adaptive Behavior*, 1(1):91–122. [2.2.2](#)
- Berman, S., Halász, Á. M., Hsieh, M. A., and Kumar, V. (2009). Optimized stochastic policies for task allocation in swarms of robots. *IEEE Transactions on Robotics*, 25(4):927–937. [2.2.3](#)
- Bert Hölldobler, E. and Wilson, E. (1990). *The Ants*. Harvard University Press, Cambridge, MA. [2.1.3](#), [2.3.1](#)
- Binder, K. (1997). Applications of monte carlo methods to statistical physics. *Reports on Progress in Physics*, 60(5):487. [7.3](#)
- Binney, J. J., Dowrick, N. J., Fisher, A. J., and Newman, M. E. J. (1992). *The Theory of Critical Phenomena: An Introduction to Renormalization Group*. Oxford University Press, New York, NY. [7.1](#)
- Bodi, M., Thenius, R., Szopek, M., Schmickl, T., and Crailsheim, K. (2012). Interaction of robot swarms using the honeybee-inspired control algorithm beeclust. *Mathematical and Computer Modelling of Dynamical Systems*, 18(1):87–100. [2.2.1](#)
- Bonabeau, E., Sobkowski, A., Theraulaz, G., and Deneubourg, J.-L. (1997). Adaptive task allocation inspired by a model of division of labor in social insects. In Narayana, A., editor, *Biocomputing and emergent computation: Proceedings of BCEC97*, pages 36–45, London, UK. World Scientific Press. [2.2.1](#), [2.2.2](#)
- Bonani, M., Magnenat, S., Rétornaz, P., and Mondada, F. (2009). The hand-bot, a robot design for simultaneous climbing and manipulation. In M. Xie et al., editor, *Proceedings of the Second International Conference on Intelligent Robotics and Applications*, number 5928 in LNAI, pages 11–22, Berlin, Germany. Springer. [A.1.3](#)
- Brambilla, M., Ferrante, E., Birattari, M., and Dorigo, M. (2013). Swarm robotics: a review from the swarm engineering perspective. *Swarm Intelligence*, 7(1):1–41. [1](#), [2](#), [2.1.3](#), [2.2](#), [2.2.1](#), [1](#), [2.2.2](#)
- Brambilla, M., Pinciroli, C., Birattari, M., and Dorigo, M. (2012). Property-driven design for swarm robotics. In Conitzer, W. and van der Hoek, P., editors, *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, pages

- 139–146. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC. [2.2.2](#), [2.2.3](#)
- Brutschy, A., Pini, G., Pinciroli, C., Birattari, M., and Dorigo, M. (2012). Self-organized task allocation to sequentially interdependent tasks in swarm robotics. *Autonomous Agents and Multi-Agent Systems*, 2012:1–25. [2.2.1](#)
- Buhl, J., Sumpter, D. J. T., Couzin, I., Hale, J., Despland, E., Miller, E., and Simpson, S. J. (2006). From disorder to order in marching locusts. *Science*, (312):1402–1406. [2.4.1](#)
- Camazine, S., Franks, N. R., Sneyd, J., Bonabeau, E., Deneubourg, J.-L., and Theraulaz, G. (2001). *Self-Organization in Biological Systems*. Princeton University Press, Princeton, NJ. [2.1.1](#), [2.2.1](#)
- Campo, A. (2011). *On the Design of Self-Organized Decision Making in Robot Swarms*. PhD thesis, IRIDIA, Université Libre de Bruxelles, Belgium. [2.2.1](#)
- Campo, A., Garnier, S., Dédriché, O., Zekkri, M., and Dorigo, M. (2011). Self-organized discrimination of resources. *PLoS ONE*, 6(5). [2.2.1](#)
- Campo, A., Nouyan, S., Birattari, M., Groß, R., and Dorigo, M. (2006). Negotiation of goal direction for cooperative transport. In Dorigo, M., Gambardella, L. M., Birattari, M., Martinoli, A., Poli, R., and Stützle, T., editors, *Ant Colony Optimization and Swarm Intelligence (ANTS) 2006*, volume 4150 of *Lecture Notes in Computer Science*, pages 191–202, Berlin, Germany. Springer. [2.4.4](#), [4.8](#), [5.1](#)
- Cao, Y. and Ren, W. (2012). Distributed coordinated tracking with reduced interaction via a variable structure approach. *IEEE Transactions on Automatic Control*, 57(1):33–48. [2.4.3](#)
- Castellano, C., Fortunato, S., and Loreto, V. (2009). Statistical physics of social dynamics. *Reviews of Modern Physics*, 81(2):591–646. [1.2](#)
- Cavagna, A., Cimarelli, A., Giardina, I., Parisi, G., Santagati, R., Stefanini, F., and Viale, M. (2010). Scale-free correlations in starling flocks. *Proceedings of the National Academy of Sciences*, 107(26):11865–11870. [2.4.1](#), [8.2](#)
- Çelikkanat, H. and Şahin, E. (2010). Steering self-organized robot flocks through externally guided individuals. *Neural Computing and Applications*, 19(6):849–865. [2.4.4](#), [3.5.2](#), [4.2](#), [4.6](#), [6.1](#), [6.3.1](#)
- Chaté, H., Ginelli, F., Grégoire, G., and Raynaud, F. (2008). Collective motion of self-propelled particles interacting without cohesion. *Physical Review E*, 77:046113. [7.3](#)
- Christensen, A. L., Grady, R. O., and Dorigo, M. (2009). From Fireflies to Fault-Tolerant Swarms of Robots. *IEEE Transactions on Evolutionary Computation*, 13(4):754–766. [8.2](#)

- Clark, C. W. and Mangel, M. (1986). The evolutionary advantages of group foraging. *Theoretical Population Biology*, 30(1):45 – 75. [2.1.3](#)
- Couzin, I. D. and Franks, N. (2002). Self-organized lane formation and optimized traffic flow in army ants. *Proceedings-Royal Society of London B*, 270(1511):139–146. [2.4.1](#)
- Couzin, I. D., Krause, J., Franks, N. R., and Levin, S. A. (2005). Effective leadership and decision-making in animal groups on the move. *Nature*, 433:513–516. [2.1.1](#), [2.4.1](#), [2.4.4](#), [3.1](#), [4.2](#), [4.8](#), [6.2](#), [6.2](#)
- Couzin, I. D., Krause, J., James, R., Ruxton, G. D., and Franks, N. R. (2002). Collective memory and spatial sorting in animal groups. *Journal of Theoretical Biology*, 218(1):1–11. [2.4](#), [2.4.1](#), [7.1](#)
- Şahin, E. (2005). Swarm robotics: From sources of inspiration to domains of application. In Şahin, E. and Spears, W., editors, *Swarm Robotics Workshop: State-of-the-art Survey*, number 3342 in Lecture Notes in Computer Science, pages 10–20, Berlin, Germany. Springer. [1](#), [2.2](#)
- Cucker, F. and Huepe, C. (2008). Flocking with informed agents. *MathematicS in Action*, 1:1–25. [2.4.2](#)
- Czirok, A., Barabasi, A.-L., and Vicsek, T. (1999). Collective motion of self-propelled particles: Kinetic phase transition in one dimension. *Physical Review Letters*, 82(1):209–212. [2.4.1](#)
- Czirók, A., Stanley, H. E., and Vicsek, T. (1997). Spontaneously ordered motion of self-propelled particles. *Journal of Physics A: Mathematical and General*, 30(5):1375. [7.1](#), [7.2](#)
- Czirok, A. and Vicsek, T. (2000). Collective behavior of interacting self-propelled particles. *Physica A*, 281:17–29. [2.4.1](#)
- Deutsch, A., Theraulaz, G., and Vicsek, T. (2012). Collective motion in biological systems. *Interface Focus*, 2(6):689–692. [7.1](#)
- Di Caro, G. and Dorigo, M. (1998). AntNet: Distributed stigmergetic control for communications networks. *Journal of Artificial Intelligence Research (JAIR)*, 9:317–365. [2.1.2](#)
- Diwold, K., Schaerf, T., Myerscough, M., Middendorf, M., and Beekman, M. (2011). Deciding on the wing: in-flight decision making and search space sampling in the red dwarf honeybee *Apis florea*. *Swarm Intelligence*, 5(2):121–141. [3.4.3](#)
- Ducatelle, F., Di Caro, G. A., Pinciroli, C., Mondada, F., and Gambardella, L. M. (2011). Communication assisted navigation in robotic swarms: self-organization and cooperation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2011)*, pages 4981–4988. IEEE Computer Society Press, Los Alamitos, CA. [2.2.1](#)

- Ferrante, E., Brambilla, M., Birattari, M., and Dorigo, M. (2010a). "Look-out!": Socially mediated obstacle avoidance in collective transport: Complete data. Supplementary information page at <http://iridia.ulb.ac.be/supp/IridiaSupp2010-005/>. 5.4
- Ferrante, E., E. Duéñez Guzmán, Turgut, A. E., and Wenseleers, T. (2013a). Geswarm: Grammatical evolution for the automatic synthesis of collective behaviors in swarm robotics. In *Proceedings of the fifteenth international conference on Genetic and evolutionary computation conference companion*, pages 17–24. ACM, New York, NY, USA. 8.2
- Ferrante, E., Turgut, A. E., Dorigo, M., and Huepe, C. (2013b). Collective motion dynamics of active solids and active crystals. Technical Report TR/IRIDIA/2013-007, IRIDIA, ULB. Submitted to *New Journal of Physics*. 7
- Ferrante, E., Turgut, A. E., Dorigo, M., and Huepe, C. (2013c). Elasticity-driven collective motion in active solids and active crystals. Technical Report <http://arxiv.org/abs/1301.2620>, IRIDIA, ULB. Submitted to *Physical Review Letters*. 7
- Ferrante, E., Turgut, A. E., Huepe, C., Stranieri, A., Pinciroli, C., and Dorigo, M. (2012a). Self-organized flocking with a mobile robot swarm: a novel motion control method. *Adaptive Behavior*, 20(6):460–477. 3.5.1
- Ferrante, E., Turgut, A. E., Huepe, C., Stranieri, A., Pinciroli, C., and Dorigo, M. (2012b). Self-organized flocking with a mobile robot swarm: a novel motion control method. Complete data. Supplementary information page at <http://iridia.ulb.ac.be/supp/IridiaSupp2012-007/>. 6.5.1, 6.5.1, 6.6.1, 6.6.1, 6.6.1, 6.6.1
- Ferrante, E., Turgut, A. E., Mathews, N., Birattari, M., and Dorigo, M. (2010b). Flocking in stationary and non-stationary environments: A novel communication strategy for heading alignment. In Schaefer, R., Cotta, C., Kołodziej, J., and Rudolph, G., editors, *Parallel Problem Solving from Nature – PPSN XI*, volume 6239 of *Lecture Notes in Computer Science*, pages 331–340, Berlin, Germany. Springer. 4.4.1
- Ferrante, E., Turgut, A. E., Stranieri, A., Pinciroli, C., Birattari, M., and Dorigo, M. (2011). A self-adaptive communication strategy for flocking in stationary and non-stationary environments: Complete data. Supplementary information page at <http://iridia.ulb.ac.be/supp/IridiaSupp2011-025/>. 4.4.1, 4.5.1, 4.7.1, 4.7.2
- Fetter, A. and Walecka, J. (2003). *Theoretical Mechanics of Particles and Continua*. Dover Books on Physics Series. Dover, Mineola, NY. 7, 7.1, 7.5, 7.5
- François, G., Sophie, B., Nicolas, B., and Gutierrez, M. (2006). Waves of agitation inside anchovy schools observed with multibeam sensor: a way to transmit information in response to predation. *ICES Journal of Marine Science*, 63:1405–1417. 4.1

- Franks, N. R., Hooper, J. W., Gumn, M., Bridger, T. H., Marshall, J. A. R., Gross, R., and Dornhaus, A. (2007). Moving targets: collective decisions and flexible choices in house-hunting ants. *Swarm Intelligence*, 1(2):81–94. [4.1](#)
- Garnier, S., Gautrais, J., and Theraulaz, G. (2007). The biological principles of swarm intelligence. *Swarm Intelligence*, 1(1):3–31. [2.1.3](#)
- Garnier, S., Jost, C., Jeanson, R., Gautrais, J., Asadpour, M., Caprari, G., and Theraulaz, G. (2005). Aggregation behaviour as a source of collective decision in a group of cockroach-like robots. In Capcarrère, M. S., Freitas, A., Bentley, P. J., Johnson, C. G., and Timmis, J., editors, *Advances in Artificial Life*, volume 3630 of *Lecture Notes in Artificial Intelligence*, pages 169–178. Springer, Berlin, Germany. [2.2.1](#), [8.2](#)
- Gazi, V. and Passino, K. M. (2004). A class of attractions/repulsion functions for stable swarm aggregations. *International Journal of Control*, 77(18):1567–1579. [2.2.3](#)
- Gjondrekaj, E., Loreti, M., Pugliese, R., Tiezzi, F., Pinciroli, C., Brambilla, M., Birattari, M., and Dorigo, M. (2012). Towards a Formal Verification Methodology for Collective Robotic Systems. In Aoki, T. and Taguchi, K., editors, *Formal Methods and Software Engineering*, volume 7635 of *Lecture Notes in Computer Science*, pages 54–70. Springer, Berlin, Germany. [5.5](#)
- Gökçe, F. and Şahin, E. (2010). The pros and cons of flocking in the long-range “migration” of mobile robot swarms. *Theoretical Computer Science*, 411(21):2140–2154. [2.4.4](#)
- Goldberg, D. E. (1989). *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley, Reading, MA. [2.2.2](#)
- Grégoire, G. and Chaté, H. (2004). Onset of collective and cohesive motion. *Physical Review Letters*, 92:025702. [7](#), [7.1](#), [7.2](#), [7.3](#), [7.7](#)
- Grégoire, G., Chaté, H., and Tu, Y. (2003). Moving and staying together without a leader. *Physica D*, 181(34):157 – 170. [2.4.2](#), [3.3](#), [7.1](#), [7.2](#), [7.7](#)
- Groß R. and M.Dorigo (2009). Towards group transport by swarms of robots. *International Journal of Bio-Inspired Computation*, 1(1-2):1–13. [5.1](#)
- Grossman, D., Aranson, I. S., and Jacob, E. B. (2008). Emergence of agent swarm migration and vortex formation through inelastic collisions. *New Journal of Physics*, 10:023036. [2.4.2](#), [7.1](#)
- Guilford, T. and Dawkins, M. S. (1991). Receiver psychology and the evolution of animal signals. *Animal Behaviour*, 42(1):1 – 14. [2.3.1](#)

- Hamann, H. and Wörn, H. (2008). A framework of space-time continuous models for algorithm design in swarm robotics. *Swarm Intelligence*, 2(2–4):209–239. [2.2.3](#)
- Hamilton, W. D. (1964). The genetical evolution of social behaviour. I. *Journal of Theoretical Biology*, 7(1):1–16. [2.1.1](#)
- Hamilton, W. D. (1971). Geometry for the selfish herd. *Journal of Theoretical Biology*, 31(2):295–311. [2.1.1](#)
- Hanada, Y., Geunho, L., and Chong, N. (2007). Adaptive flocking of a swarm of robots based on local interactions. In *Proceedings of the IEEE Swarm Intelligence Symposium*, pages 340–347, Piscataway, NJ. IEEE Press. [2.4.3](#)
- Handl, J. and Meyer, B. (2007). Ant-based and swarm-based clustering. *Swarm Intelligence*, 1(2):95–113. [2.1.2](#)
- Hayes, A. and Dormiani-Tabatabaei, P. (2002). Self-organized flocking with agent failure: Off-line optimization and demonstration with real robots. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 3900–3905, Piscataway, NJ. IEEE Press. [2.4.4](#), [6.1](#)
- Henkes, S., Fily, Y., and Marchetti, M. C. (2011). Active jamming: Self-propelled soft particles at high density. *Physical Review E*, 84(4):040301(R). [2.4.2](#), [7.1](#)
- Hoare, D. J., Krause, J., Peuhkuri, N., and Godin, J.-G. J. (2000). Body size and shoaling in fish. *Journal of Fish Biology*, 57(6):1351–1366. [1](#), [8.1](#)
- Holland, O., Woods, J., Nardi, R., and Clark, A. (2005). Beyond swarm intelligence: the ultraswarm. In *Proceedings of the IEEE Swarm Intelligence Symposium*, pages 217–224, Piscataway, NJ. IEEE Press. [2.4.4](#), [6.1](#)
- Howard, A., Matarić, M. J., and Sukhatme, G. S. (2002). Mobile sensor network deployment using potential fields: a distributed, scalable solution to the area coverage problem. In *Proceedings of the 2002 International Symposium on Distributed Autonomous Robotic Systems (DARS 2002)*, pages 299–308, Piscataway, NJ. IEEE Press. [2.2.1](#)
- Hsieh, M. A., Halász, Á., Berman, S., and Kumar, V. (2008). Biologically inspired redistribution of a swarm of robots among multiple sites. *Swarm Intelligence*, 2(2–4):121–141. [2.2.3](#)
- Jadbabaie, A., Lin, J., and Morse, A. S. (2003). Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Transactions on Automatic Control*, 48(6):988–1001. [2.4.3](#)

- Jones, J. E. (1924). On the Determination of Molecular Fields. I. From the Variation of the Viscosity of a Gas with Temperature. *Proceedings of the Royal Society of London. Series A*, 106(738):441–462. [7](#)
- Judd, T. M. and Sherman, P. W. (1996). Naked mole-rats recruit colony mates to food sources. *Animal Behaviour*, 52(5):957–969. [2.3.1](#)
- Kaelbling, L. P., Littman, M. L., and Moore, A. W. (1996). Reinforcement learning: a survey. *Journal of Artificial Intelligence Research*, 4:237–285. [1.2](#), [2.2.2](#)
- Kelly, I. and Keating, D. (1996). Flocking by the fusion of sonar and active infrared sensors on physical autonomous robots. In *Proceedings of the Third International Conference on Mechatronics and Machine Vision in Practice*, pages 14–17, Piscataway, NJ. IEEE Press. [2.4.4](#)
- Kennedy, J. and Eberhart, R. (1995). Particle swarm optimization. In *Proceedings of the IEEE International Conference on Neural Networks*, volume 4, pages 1942–1948, Piscataway, NJ. IEEE Press. [1.2](#), [2.1.2](#)
- Khatib, O. (1986). Real-time obstacle avoidance for manipulators and mobile robots. *The International Journal of Robotics Research*, 5(1):90–98. [2.2.2](#)
- Kloeden, P. E. and Platen, E. (2011). *Numerical Solution of Stochastic Differential Equations*. Springer, Berlin, Germany. [7.3](#)
- Krause, J. and Ruxton, G. (2002). *Living in groups*. Oxford University Press, New York, NY. [2.1.1](#)
- Langer, J. S. (1980). Instabilities and pattern formation in crystal growth. *Reviews of Modern Physics*, 52(1):1–28. [2.2.1](#)
- Lerman, K. and Galstyan, A. (2002). Mathematical model of foraging in a group of robots: effect of interference. *Autonomous Robots*, 13(2):127–141. [2.2.3](#)
- Lindhe, M., Ogren, P., and Johansson, K. (2005). Flocking with obstacle avoidance: A new distributed coordination algorithm based on Voronoi partitions. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1785–1790, Piscataway, NJ. IEEE Press. [2.4.3](#)
- Liu, W., Winfield, A. F. T., Sa, J., Chen, J., and Dou, L. (2007). Towards energy optimization: emergent task allocation in a swarm of foraging robots. *Adaptive Behavior*, 15(3):289–305. [2.2.1](#)
- Liu, Y. and Passino, K. M. (2004). Stable social foraging swarms in a noisy environment. *IEEE Transactions on Automatic Control*, 49(1):30–44. [2.2.3](#)

- Lizier, J. T., Prokopenko, M., and Zomaya, A. Y. (2008). Local information transfer as a spatiotemporal filter for complex systems. *Physical Review E*, 77(2):026110. [8.2](#)
- Magenat, S., Longchamp, V., Bonani, M., Rétornaz, P., Germano, P., Bleuler, H., and Mondada, F. (2010). Affordable slam through the co-design of hardware and methodology. In Snyder, W. and Kumar, V., editors, *IEEE International Conference on Robotics and Automation*, Piscataway, NJ. IEEE Press. [ii](#)
- Makinson, J. C., Oldroyd, B. P., Schaerf, T. M., Wattanachaiyingcharoen, W., and Beekman, M. (2011). Moving home: nest-site selection in the red dwarf honeybee (*apis florea*). *Behavioral Ecology and Sociobiology*, 65(5):945–958. [3.4.3](#)
- Martinoli, A., Easton, K., and Agassounon, W. (2004). Modeling swarm robotic systems: a case study in collaborative distributed manipulation. *The International Journal of Robotics Research*, 23(4–5):415–436. [2.2.3](#)
- Martinoli, A., Ijspeert, A. J., and Mondada, F. (1999). Understanding collective aggregation mechanisms: from probabilistic modelling to experiments with real robots. *Robotics and Autonomous Systems*, 29(1):51–63. [2.2.3](#)
- Massink, M., Brambilla, M., Latella, D., Dorigo, M., and Birattari, M. (2012). Analysing robot swarm decision-making with bio-pepa. In *Swarm Intelligence*, volume 7461 of *Lecture Notes in Computer Science*, pages 25–36. Springer, Berlin, Germany. [2.2.3](#)
- Matarić, M. J. (1994). *Interaction and Intelligent Behavior*. PhD thesis, MIT, MA. [2.4.4](#), [6.1](#)
- Matarić, M. J. (1998). Using communication to reduce locality in distributed multi-agent learning. *Journal of Experimental and Theoretical Artificial Intelligence*, 10(3):357–369. [2.2.2](#)
- Meinhardt, H. (1982). *Models of biological pattern formation*, volume 6. Academic Press, London, UK. [2.2.1](#)
- Menzel, A. M. and Ohta, T. (2012). Soft deformable self-propelled particles. *Europhysics Letters*, 99:154101. [2.4.2](#), [7.1](#)
- Mermin, N. D. and Wagner, H. (1966). Absence of ferromagnetism or antiferromagnetism in one- or two-dimensional isotropic heisenberg models. *Physical Review Letters*, 17:1133–1136. [7.2](#)
- Minsky, M. (1967). *Computation: Finite and Infinite Machines*. Prentice-Hall, Upper Saddle River, NJ. [2.2.1](#), [2.2.2](#)
- Monteiro, S. and Bicho, E. (2010). Attractor Dynamics Approach to Formation Control: Theory and Application. *Autonomous Robots*, 29(3-4):331–355. [2.4.4](#)

- Montes de Oca, M. A., Ferrante, E., Scheidler, A., Pinciroli, C., Birattari, M., and Dorigo, M. (2011). Majority-rule opinion dynamics with differential latency: a mechanism for self-organized collective decision-making. *Swarm Intelligence*, 5(3–4):305–327. [2.2.1](#), [2.2.3](#)
- Moshtagh, N., Jadbabaie, A., and Daniilidis, K. (2006). Vision-based control laws for distributed flocking of nonholonomic agents. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 2769–2774, Piscataway, NJ. IEEE Press. [2.4.3](#)
- Moslinger, C., Schmickl, T., and Crailsheim, K. (2009). A minimalistic flocking algorithm for swarm robots. In Kampis, G., Karsai, I., and Szathmàry, E., editors, *European Conference of Artificial Life (ECAL)*, Berlin, Germany. Springer. [2.4.4](#)
- Moussaïd, M., Helbing, D., and Theraulaz, G. (2011). How simple rules determine pedestrian behavior and crowd disasters. *Proceedings of the National Academy of Sciences*, 108(17):6884–6888. [2.1.1](#)
- Moyle, P. and Cech, J. (2004). *Fishes: an introduction to ichthyology*. Prentice-Hall, Upper Saddle River, NJ. [2.4.1](#)
- Nembrini, J., Winfield, A. F. T., and Melhuish, C. (2002). Minimalist Coherent Swarming of Wireless Networked Autonomous Mobile Robots. In *From Animals to Animats*, volume 7, pages 273–382, Berlin, Germany. Springer. [2.4.4](#)
- Newton, I. (2010). *The Migration Ecology of Birds*. Elsevier, Amsterdam, Netherlands. [1](#), [8.1](#)
- Nolfi, S. and Floreano, D. (2004). *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. Intelligent Robots and Autonomous Agents. MIT Press, Cambridge, MA. [2.2.1](#), [2.2.2](#), [8.2](#)
- Ogata, K. (2001). *Modern Control Engineering (4th Edition)*. Prentice Hall, Upper Saddle River, NJ. [7](#), [7.5](#)
- O’Grady, R., Christensen, A., and Dorigo, M. (2009a). SWARMORPH: multi-robot morphogenesis using directional self-assembly. *IEEE Transactions on Robotics*, 25(3):738–743. [2.2.2](#)
- O’Grady, R., Pinciroli, C., Christensen, A. L., and Dorigo, M. (2009b). Supervised group size regulation in a heterogeneous robotic swarm. In *9th Conference on Autonomous Robot Systems and Competitions, Robótica 2009*, pages 113–119. IPCB-Instituto Politécnico de Castelo Branco, Castelo Branco, Portugal. [2.2.3](#)
- Olfati-Saber, R. (2006). Flocking for multi-agent dynamic systems: Algorithms and theory. *Transactions on Automatic Control*, 51(3):401–420. [2.4.3](#)
- Panait, L. and Luke, S. (2005). Cooperative multi-agent learning: the state of the art. *Autonomous Agents and Multi-Agent Systems*, 11(3):387–434. [2.2.2](#)

- Partridge, B. (1982). The structure and function of fish schools. *Scientific American*, 246(6):114–123. [2.4.1](#)
- Pillot, M.-H., Gautrais, J., Arrufat, P., Couzin, I. D., Bon, R., and Deneubourg, J.-L. (2011). Scalable rules for coherent group motion in a gregarious vertebrate. *PLoS ONE*, 6(1):e14487. [2.1.1](#)
- Pinciroli, C., Trianni, V., O’Grady, R., Pini, G., Brutschy, A., Brambilla, M., Mathews, N., Ferrante, E., Di Caro, G., Ducatelle, F., Birattari, M., Gambardella, L. M., and Dorigo, M. (2012). ARGoS: a modular, parallel, multi-engine simulator for multi-robot systems. *Swarm Intelligence*, 6(4):271–295. [A.2](#)
- Pitcher, T. J., Magurran, A. E., and Winfield, I. J. (1982). Fish in larger shoals find food faster. *Behavioral Ecology and Sociobiology*, 10:149–151. [2.4.1](#)
- Press, W., Teukolsky, S., Vetterling, W., and Flannery, B. (1992). *Numerical Recipes in C*. Cambridge University Press, Cambridge, UK. [4.2](#)
- Prorok, A., Correll, N., and Martinoli, A. (2011). Multi-level spatial modeling for stochastic distributed robotic systems. *The International Journal of Robotics Research*, 30(5):574–589. [2.2.3](#)
- Ratcliffe, J. M. and ter Hofstede, H. M. (2005). Roosts as information centres: social learning of food preferences in bats. *Biology Letters*, 1(1):72–74. [2.3.2](#)
- Regmi, A., Sandoval, R., Byrne, R., Tanner, H., and Abdallah, C. (2005). Experimental implementation of flocking algorithms in wheeled mobile robots. In *Proceedings of the American Control Conference*, volume 7, pages 4917–4922, Piscataway, NJ. IEEE Press. [2.4.3](#)
- Reif, J. H. and Wang, J. (1999). Social potential fields: a distributed behavioral control for autonomous robots. *Robotics and Autonomous Systems*, 27(3):171–194. [2.2.2](#)
- Reynolds, C. (1987). Flocks, herds and schools: A distributed behavioral model. In Stone, M. C., editor, *SIGGRAPH ’87: Proceedings of the 14th annual conference on computer graphics and interactive techniques*, pages 25–34, New York, NY. ACM Press. [1](#), [2.4](#), [2.4.3](#), [3.1](#)
- Reynolds, C. (2008). Boids (flocks, herds and schools: a distributed behavioral model). <http://www.red3d.com/cwr/boids/>. [2.2](#)
- Riley, J. R., Greggers, U., Smith, A. D., Reynolds, D. R., and Menzel, R. (2005). The flight paths of honeybees recruited by the waggle dance. *Nature*, 435:205–207. [2.3.1](#)
- Romanczuk, P., Couzin, I. D., and Schimansky-Geier, L. (2009). Collective motion due to individual escape and pursuit response. *New Journal of Physics*, 102(1):010602. [2.4.1](#), [2.4.2](#), [6.1](#), [7.1](#), [8.2](#)

- S. Magnenat, P. Rtornaz, M. B. V. L. F. M. (2011). Aseba: A modular architecture for event-based control of complex robots. *IEEE/ASME Transactions on Mechatronics*, 16(2):321–329. [A.1.1](#)
- Schmickl, T., Hamann, H., Wörn, H., and Crailsheim, K. (2009). Two different approaches to a macroscopic model of a bio-inspired robotic swarm. *Robotics and Autonomous Systems*, 57(9):913–921. [2.2.3](#)
- Seeley, T. (2010). *Honeybee Democracy*. Princeton University Press. Princeton University Press, Princeton, NJ. [2.3.1](#)
- Simpson, S. J., Sword, G. A., Lorch, P. D., and Couzin, I. D. (2006). Cannibal crickets on a forced march for protein and salt. *Proceedings of the National Academy of Sciences*, 103(11):4152–4156. [2.4.1](#), [6.1](#)
- Smith, T. and Harper, D. (2003). *Animal Signals*. Oxford Series in Ecology and Evolution. Oxford University Press, New York, NY. [2.3.1](#)
- Sonerud, G. A., Smedshaug, C. A., and Brthen, . (2001). Ignorant hooded crows follow knowledgeable roost-mates to food: support for the information centre hypothesis. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 268(1469):827–831. [2.3.2](#)
- Soysal, O., Bahçeci, E., and Şahin, E. (2007). Aggregation in swarm robotic systems: evolution and probabilistic control. *Turkish Journal of Electrical Engineering and Computer Sciences*, 15(2):199–225. [8.2](#)
- Spears, W. M., Spears, D. F., Hamann, J. C., and Heil, R. (2004). Distributed, physics-based control of swarms of vehicles. *Autonomous Robots*, 17:137–162. [2.2.2](#), [2.4.4](#), [3](#), [3.3](#), [6.1](#)
- Sperati, V., Trianni, V., and Nolfi, S. (2011). Self-organised path formation in a swarm of robots. *Swarm Intelligence*, 5(2):97–119. [2.2.2](#)
- Stabentheiner, A., Kovac, H., and Brodschneider, R. (2010). Honeybee colony thermoregulation regulatory mechanisms and contribution of individuals in dependence on age, location and thermal stress. *PLoS ONE*, 5(1):e8967. [2.1.1](#)
- Stewart, K. J. and Harcourt, A. H. (1994). Gorillas' vocalizations during rest periods: signals of impending departure? *Behaviour*, 130(1):29–40. [2.3.1](#)
- Stranieri, A., Ferrante, E., Turgut, A., Trianni, V., Pinciroli, C., Birattari, M., and M., M. D. (2011a). Self-Organized flocking with a heterogeneous mobile robot swarm. In Lenaerts, T., Giacobini, M., Bersini, H., Borgine, P., Dorigo, M., and Doursat, R., editors, *Proceedings of ECAL 2001*, pages 789–796, Cambridge, MA. MIT Press. [3.5.1](#)

- Stranieri, A., Ferrante, E., Turgut, A. E., Trianni, V., Pinciroli, C., Birattari, M., and Dorigo, M. (2011b). Self-organized flocking with a heterogeneous mobile robot swarm. Supplementary information page at <http://iridia.ulb.ac.be/supp/IridiaSupp2011-011/>. 6.3.1, 6.4.1, 6.4.2
- Sumpter, D. J. T. (2010). *Collective Animal Behavior*. Princeton University Press, Princeton, NJ. 2.3
- Sumpter, D. J. T., Buhl, J., Biro, D., and Couzin, I. D. (2008). Information transfer in moving animal groups. *Theory in Biosciences*, 127(2):177–186. 1, 2.3, 8.1, 8.2
- Szabó, B., Szöllösi, G. J., Gönci, B., Selmeczi, Z. J. D., and Vicsek, T. (2006). Phase transition in the collective migration of tissue cells: Experiment and model. *Physical Review E*, 74(6):061908. 2.4.2, 6.1, 7.1, 7.5, 7.7
- Tanner, H. G., Jadbabaie, A., and Pappas, G. J. (2003a). Stable flocking of mobile agents part i: Fixed topology. In *Proceedings of the 42nd IEEE Conference on Decision and Control*, volume 2, pages 2010–2015, Piscataway, NJ. IEEE Press. 2.4.3
- Tanner, H. G., Jadbabaie, A., and Pappas, G. J. (2003b). Stable flocking of mobile agents part II: Dynamic topology. In *Proceedings of the 42nd IEEE Conference on Decision and Control*, volume 2, pages 2016–2021, Piscataway, NJ. IEEE Press. 2.4.3
- Tarcai, N., Virgh, C., bel, D., Nagy, M., Vrkonyi, P. L., Vsrhelyi, G., and Vicsek, T. (2011). Patterns, transitions and the role of leaders in the collective dynamics of a simple robotic flock. *Journal of Statistical Mechanics: Theory and Experiment*, 2011(04):P04010. 2.4.2, 2.4.4
- Templeton, J. J. and Giraldeau, L.-A. (1996). Vicarious sampling: the use of personal and public information by starlings foraging in a simple patchy environment. *Behavioral Ecology and Sociobiology*, 38:105–114. 2.3.2
- Toner, J. and Tu, Y. (1995). Long-range order in a two-dimensional dynamical XY model: How birds fly together. *Physical Review Letters*, 75:4326–4329. 7.1, 7.2, 7.7
- Toner, J. and Tu, Y. (1998). Flocks, herds, and schools: A quantitative theory of flocking. *Physical Review E*, 58:4828–4858. 7.1
- Trianni, V. and Dorigo, M. (2006). Self-organisation and communication in groups of simulated and physical robots. *Biological Cybernetics*, 95:213–231. 5.1
- Trianni, V. and Nolfi, S. (2011). Engineering the evolution of self-organizing behaviors in swarm robotics: A case study. *Artificial Life*, 17(3):183–202. 2.1.3, 2.2.2

- Turgut, A. E., Çelikkanat, H., Gökçe, F., and Şahin, E. (2008a). Self-organized flocking in mobile robot swarms. *Swarm Intelligence*, 2(2):97–120. [2.4](#), [2.4.2](#), [2.4.4](#), [3](#), [3.2](#), [3.5.2](#), [3.5.3](#), [4.2](#), [4.3.1](#), [6.1](#), [6.3.1](#), [6.3.1](#), [6.5.1](#)
- Turgut, A. E., Huepe, C., Çelikkanat, H., Gökçe, F., and Şahin, E. (2008b). *Modeling Phase Transition in Self-organized Mobile Robot Flocks*, volume 5217 of *Lecture Notes in Computer Science*, pages 108–119. Springer, Berlin, Germany. [2.4.2](#)
- Vicsek, T., Czirok, A., Ben-Jacob, E., Cohen, I., and Shochet, O. (1995). Novel type of phase transition in a system of self-driven particles. *Physical Review Letters*, 75(6):1226–1229. [2.4](#), [2.4.1](#), [2.4.2](#), [2.4.3](#), [4.2](#), [4.3.1](#), [6.1](#), [6.2](#), [6.3.1](#), [7](#), [7.1](#), [7.2](#), [7.7](#)
- Vicsek, T. and Zafeiris, A. (2012). Collective motion. *Physics Reports*, 517(3–4):71–140. [2.4.2](#), [7.1](#)
- Voronoi, G. (1908). Nouvelles applications des paramètres continus à la théorie des formes quadratiques. *Journal für die Reine und Angewandte Mathematik*, 133(97-178). [3.3](#)
- Waibel, M., Keller, L., and Floreano, D. (2009). Genetic team composition and level of selection in the evolution of cooperation. *IEEE Transactions on Evolutionary Computation*, 13(3):648–660. [2.2.2](#)
- Wang, X., Miller, J., Lizier, J., Prokopenko, M., and Rossi, L. (2012). Measuring information storage and transfer in swarms. In Christoph, A., David M., B., Charles, O., and Robert T., P., editors, *Artificial Life 13*, volume 13, pages 838–845. MIT Press, Cambridge, MA. [8.2](#)
- Ward, P. and Zahavi, A. (1973). Importance of certain assemblages of birds as information-centers for food-finding. *Ibis*, 115(4):517–534. [2.1.1](#)
- Winfield, A. F. T., Sa, J., Fernandez-Gago, M. C., Dixon, C., and Fisher, M. (2005). On formal specification of emergent behaviours in swarm robotic systems. *International Journal of Advanced Robotic Systems*, 2(4):363–370. [2.2.3](#)
- Yu, C., Werfel, J., and Nagpal, R. (2010). Collective decision-making in multi-agent systems by implicit leadership. In van der Hoek, Kaminka, L. L. and Sen, editors, *Proceedings of 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS2010)*, Richland, SC. International Foundation for Autonomous Agents and Multiagent Systems. [2.4.1](#)