UNIVERSITÉ LIBRE DE BRUXELLES
Ecole Polytechnique de Bruxelles
IRIDIA - Institut de Recherches Interdisciplinaires
et de Développements en Intelligence Artificielle

Engineering swarm systems: A design pattern for the best-of-n decision problem

Andreagiovanni REINA

Ph.D. Thesis

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

Co-Promoteur de Thèse:
Dr. **Vito TRIANNI**

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

Année académique 2015-2016

to dad and mum,
source of unconditional love

ACKNOWLEDGEMENTS

This thesis is the result of a long and hard path which I walked in the last five years. Here, I would like to acknowledge everyone who has supported me and that, in some way, has been important to reach the goal of this Ph.D.

First of all, I would like to thank Dr. Vito Trianni for guiding and supporting me in my doctoral studies. He became my co-supervisor in the midst of my studies and rescued me from a difficult period in which I was lost. I enjoyed the intense supervision relationship that we established based on very frequent interactions. Innumerate times, we discussed and reasoned about a problem together, we explored solutions together, we learned together and we overcame many difficulties together. Having him on my side for three intense years has been of utmost importance to the completion of this Ph.D. Thank you.

I would like to thank Prof. Marco Dorigo, who gave me the opportunity to conduct my studies at the prestigious IRIDIA Lab. Marco has been always available at any moment to answer any my question and to give his comments on any result. Though, I experienced a hard time when my questions or results were missing. I also thank Marco for granting me financial support through the European Research Council (ERC) Advanced Grant "E-SWARM: Engineering Swarm Intelligence Systems" (contract 246939). Thanks to Marco, I had the opportunity to practise a wide range of activities which helped me to improve my skills and accumulate experience of big importance for the survival in the academic world. I am sure I will benefit from these valuable lessons during my career. Additionally, his precision in organising and presenting ideas forced me to upgrade my exposition skills to meet his high standards.

For some very awkward and obscure reason, my supervisor and co-supervisors share the same birthday with my father and mother, respectively. Even if this information does not belong to the acknowledgements, I believe it has been worth to write a note about it. Future studies aim at investigating the cause of such a rare event.

Dr. Mauro Birattari supervised me in several teaching activities, ranging from the co-supervision of four M.Sc. theses to the direction of several M.Sc. projects. I owe to him much of my improvements in supervision skills. I hope that he enjoyed as much as I did the fruitful discussions we had. The part of joint work I enjoyed the most has been the one involving creativ-

ity, when, each year, we met to discuss and define together new theses and projects for students. Our interactions did not limit to work but we enjoyed tons of chats during our work breaks. Thank you, Mauro.

At IRIDIA, I had the luck to find a wonderful group of colleagues and friends. Being with *The iridians* has been very important for me and allowed me to *keep calm and carry on* even in the most critical periods. I imagine that bearing me all these years has not been easy for them. So, I did my best to return them all the positive energy that they shared with me. I avoid making the long list of names with the obvious result of disappointing somebody that I would forget in the rush of the submission. I am glad to have met each of you, iridians, and I am very happy to have found not only colleagues but many friends. Thank you.

Then, I would like to thank Prof. James Marshall, who welcomed me in his group at the University of Sheffield and allowed me to take the time necessary to complete the writing of this thesis. I look forward to making great research together!

In Brussels, I dedicated most of my free time to let discs fly. I would like to thank all the friends that shared with me the passion for Ultimate Frisbee and with which I shared many weekends. Disconnecting my mind from work has been vital. In particular, I thank the Mooncatchers family that welcomed me to be part of them and let me live years of great ultimate. Thank XLR8RS for the much fun we had together and for the one that I am sure we will keep having. And finally, thank Flying Rabbits, even if we met just a few months before my departure, knowing you has been great. Special thanks go to Vince which has been a great flatmate despite all the difficulties he is confronted with.

Then, I would like to thank my dearest, Leslie, for supporting and encouraging me during these years. She never gave up and has been there to listen to my complaints for failures, to motivate me and to celebrate my successes. Your love and affection have been very important to me and I feel very lucky to share my life with you.

Infine, ringrazio infinitamente la mia famiglia, specialmente mamma e papà i quali sono stati fonte di incondizionato amore. Anche se distante da casa, vi ho sempre sentito vicino.

ABSTRACT

The study of large-scale decentralised systems composed of numerous interacting agents that self-organise to perform a common task is receiving growing attention in several application domains. However, real world implementations are limited by a lack of well-established design methodologies that provide performance guarantees. Engineering such systems is a challenging task because of the difficulties to obtain the *micro-macro link*: a correspondence between the microscopic description of the individual agent behaviour and the macroscopic models that describe the system's dynamics at the global level. In this thesis, we propose an engineering methodology for designing decentralised systems, based on the concept of *design patterns*. A design pattern provides a general solution to a specific class of problems which are relevant in several application domains. The main component of the solution consists of a multi-level description of the collective process, from macro to micro models, accompanied by rules for converting the model parameters between description levels. In other words, the design pattern provides a formal description of the micro-macro link for a process that tackles a specific class of problems. Additionally, a design pattern provides a set of case studies to illustrate possible implementation alternatives both for simple or particularly challenging scenarios. We present a design pattern for the best-of-n, decentralised decision problem that is derived from a model of nest-site selection in honeybees. We present two case studies to showcase the design pattern usage in (i) a multiagent system interacting through a fully-connected network, and (ii) a swarm of particles moving on a bidimensional plane.

CONTENTS

ACRONYMS

ADR Advection-Diffusion-Reaction

CDCI Collective Decision Making through Cross Inhibition

CDF Cumulative Distribution Function

DDM Drift-Diffusion Model

EGG Embedded Graph Grammar

ODE Ordinary Differential Equation

PDE Partial Differential Equation

PFSM Probabilistic Finite State Machine

TVWS TV White Space

WSN Wireless Sensor Networks

Nothing in life is to be feared, it is only to be understood. Now is the time to understand more, so that we may fear less.

—Marie Skłodowska Curie

1

INTRODUCTION

The last decades witnessed a proliferation of interconnected smart devices that are able to operate in an autonomous way. This phenomenon will increase in the future leading to a multitude of large-scale decentralised systems (Helbing and Pournaras, 2015). These systems, which we refer to as *swarm systems* in analogy with insect swarms or other similar biological systems (see also Section 2.1), will be composed of numerous autonomous agents (e.g., smart devices, robots) interacting with each other. Understanding and controlling such swarm systems will be necessary in order to be able to exploit their potential. Several studies investigated methods to analyse, design and engineer swarm systems from different viewpoints (Liu et al., 2011; Helbing, 2014; Lee et al., 2014). However, general methodologies to design and engineer swarm systems are missing due to the difficulties to model (complex) systems composed of numerous interacting components and to predict their behaviour. The main cause of these difficulties is the gap between the behaviour at the system (swarm) level and the behaviour of each agent composing the swarm at the local, individual level. In several cases, even if the individual behaviour of each agent is known, the resulting swarm behaviour is hard to predict. The lack of tools to predict the behaviour and, thus, to guarantee the performance of the system is one of the main causes that prevents the diffusion of swarm systems in real world applications. In fact, nowadays, swarm systems applications are still bound within the research labs. We believe that an engineering methodology able to provide performance guarantees and to predict the swarm behaviour would foster the introduction of swarm systems in commercial applications and everyday life. Examples of large-scale distributed systems that would benefit from a more rigorous methodology are swarms of robots (Brambilla et al., 2013), cognitive radio networks (Akyildiz et al., 2011) and cyber-physical systems (Derler et al., 2012).

A general design methodology is difficult (and may be impossible) to achieve because of the difficulties in treating at the same time several complexity factors (e.g., heterogeneities in the interaction topology or in the

individual behaviour). Additionally, each application faces domain-specific challenges. For instance, in swarm robotics, any design methodology needs to deal with the inherent spatial factors (Hamann and Wörn, 2008b; Berman et al., 2011b; Sartoretti et al., 2014; Scheidler et al., 2015). Instead of a general design methodology, an alternative approach consists in providing general solutions to specific classes of problems. In this thesis, we propose the use of *design patterns* to organise and formalise such solutions. Each design pattern consists in a reusable package of methods and guidelines to design a swarm system allowing the prediction of the swarm behaviour and providing performance guarantees.

After introducing design patterns as a general methodology for the design of swarm systems, we present a design pattern for the cognitive ability of decentralised decision making in the best-of-n problem, that is the ability of a group of autonomous agents to select the best option among a set of n alternatives. We selected decentralised decision making in the best-of-n problem because it represents a crucial ability in several application domains (Halloy et al., 2007; Vigelius et al., 2014; Srivastava and Leonard, 2014; Valentini et al., 2016). The design pattern presented in this thesis is based on behavioural models of honeybee swarms selecting their nest site (Marshall et al., 2009; Seeley et al., 2012; Pais et al., 2013). Generally speaking, the ability of a biological system to select the most profitable option among the available alternatives is a fundamental adaptive response that can determine death or survival. This is why the decision-making strategy observed in honeybee swarms presents a near-optimal speed-accuracy tradeoff: given a desired accuracy level, a decision is taken in the shortest time possible (Marshall et al., 2009; Seeley et al., 2012). Additionally, inhibitory signals among bees allow the swarm to quickly break decision deadlocks and provide an adaptive mechanism to tune the decision dynamics as a function of the perceived quality of the available options (Seeley et al., 2012; Pais et al., 2013). These properties are desirable in many practical decision-making scenarios and justify the selection of the honeybee model as the reference.

To achieve a complete formalisation of the design pattern, we extend the model for binary decisions presented in (Seeley et al., 2012) to a more general case that deals with decisions among n options. We report a description of the best-of-n decision process at various levels: (i) a macroscopic description through an infinite-size, deterministic, time/state-space continuous model, (ii) a macroscopic description through a finite-size, stochastic, time continuous, state-space discrete model, and (iii) a microscopic, agent-based, stochastic, time/state-space discrete description. We also present rules to convert the model parameters from one description level to another. Domain-specific challenges (such as spatial factors and heterogeneity between agents composing the swarm) are tackled through a set of guidelines. Additionally, as a necessary intermediate step to the formalisation of the

design pattern, we implement a preliminary case study using a bottom-up design approach. This case study consists in the implementation of a swarm robotics system (in physics-based simulation) with its performance quantitatively predicted by the macroscopic model. Finally, we present two case studies (implemented with a top-down approach) that showcase the design pattern usage following various types of implementation strategies, selecting various system parameterisations and in a particularly challenging spatial scenario.

1.1 ORIGINAL CONTRIBUTIONS

In this section, we present the original contributions of this thesis and the corresponding scientific publications.

I. We perform an accurate analysis of the state of the art of design methods for swarm systems aimed at providing performance guarantees (see Section 2.2).

II. We propose design patterns as an engineering methodology for swarm systems. A design pattern provides a well-formalised design solution for swarm systems with performance guarantees. In other words, the design pattern provides instructions to design a swarm system so that a quantitative match between the dynamics of the macroscopic models and of the implemented swarm system is attained (see Chapter 3).

III. We present the necessary steps to formalise a design pattern for swarm systems (see Section 3.3).

IV. We attain a precise quantitative match between the macroscopic model dynamics and the dynamics of a swarm robotics system performing a shortest path discovery/selection task, implemented in physics-based simulation (see Chapter 4).

V. We present a design pattern for decentralised decision making in the best-of-n problem (see Chapter 5).

VI. We extend the binary decision-making model presented by Seeley et al. (2012) to a more general case with n options (see Section 5.5.1).

VII. We introduce a master equation to study how the system dynamics change as a function of the swarm size. We approximate the master equation solution through numerical simulations using the Gillespie algorithm (Gillespie, 1976) (see Section 5.5.2).

VIII. We showcase the design pattern implementation through two case studies. The first case study allows the investigation of a wide set of implementation choices, parameterisations and option qualities. The second case study showcases the design in a more challenging scenario that involves agents moving in a 2D environment (see Chapter 6).

IX. We provide a method to estimate macroscopic transition rates directly from experimental data through survival analysis (see Appendix A).

Part of the work and results presented in this thesis have been published in international peer-reviewed journals and conferences. Here, we present the list of these publications.

International journals

- **A. Reina**, G. Valentini, C. Fernández-Oto, M. Dorigo, and V. Trianni. A design pattern for decentralised decision making. *PLoS ONE*, 10 (10):e0140950, 2015.
- **A. Reina**, R. Miletitch, M. Dorigo, and V. Trianni. A quantitative micro-macro link for collective decisions: The shortest path discovery/selection example. *Swarm Intelligence*, 9 (23):75–102, 2015.

International conferences

- **A. Reina**, M. Dorigo, and V. Trianni. Towards a cognitive design pattern for collective decision-making. In M. Dorigo et al., editors, *Swarm Intelligence (ANTS 2014)*, volume 8667 of LNCS, pages 194–205. Springer, Berlin, Germany, 2014.
- **A. Reina**, M. Dorigo, and V. Trianni. Collective decision making in distributed systems inspired by honeybees behaviour. In *Proceedings of 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1421–1422. IFAAMAS, New York, 2014.

Additionally, part of the results of (Reina et al., 2015d) have been presented as an extended abstract at the *Biologically Distributed Algorithms Workshop (BDA 2015)* (Reina et al., 2015c).

1.2 OTHER CONTRIBUTIONS NOT RELATED TO THIS THESIS

In this section, we present the scientific contributions generated during the doctoral studies that are not related to the topic presented in this thesis.

I. DISTRIBUTED PATH PLANNING USING AN OVERHEAD CAMERA NETWORK. In this study, we have proposed a novel system for distributed path planning. The system is composed of an overhead camera network that computes in a completely decentralised way the precise sequence of roto-translations to guide a ground object of arbitrary shape from an initial location to a given destination through a large, cluttered, dynamic environment. In the investigated scenario, each camera has only a partial, local,

overhead view of the ground environment and can communicate only with the neighbouring cameras. Through local calculations and wireless message exchanges, the camera network cooperatively computes the full path. The task is made particularly challenging by alignment errors between the field of view of neighbouring cameras. We evaluated the performance of our system through an extensive set of simulation experiments and a more limited set of real-world robot experiments. The results showed that our distributed approach is more scalable, flexible and robust than a monolithic centralised computation. This work resulted in the following publication:

- **A. Reina**, L. M. Gambardella, M. Dorigo, and G. A. Di Caro. *zePPeLIN*: Distributed path planning using an overhead camera network. *International Journal of Advanced Robotic Systems*, 11(119), 2014.

II. TOOLS FOR IMPROVING SWARM ROBOTICS EXPERIMENTS. We implemented a set of tools to support a researcher while performing multi-robot experiments.

II.i **Augmented reality for robots.** In this work, we have proposed a novel virtual sensing technology in which robots are equipped with virtual sensors. Virtual sensors allow a robot to perceive a simulated environment around itself while it moves and operates in the real world. This technology is useful for (i) prototyping new sensors before producing them, and (ii) implementing and investigating robot experiments in environments that cannot be reproduced within a laboratory (e.g., because involving dangerous components such as fire or radiations). The architecture of the virtual sensing technology is based on an overhead multi-camera system that tracks the robots and communicates with a server running a physics-based robot simulator. The simulator sends the virtual sensor's information to each robot according to its sensing range. The usage of this tool has been showcased through an experiment involving 15 robots. This work resulted in the following publication:

- **A. Reina**, M. Salvaro, G. Francesca, L. Garattoni, C. Pinciroli, M. Dorigo, and M. Birattari. Augmented reality for robots: virtual sensing technology applied to a swarm of e-pucks. In *Proceedings of the IEEE 2015 NASA/ESA Conference of Adaptive Hardware and Systems (AHS)*, pages 1-6, Los Alamitos, CA, 2015. IEEE Computer Society Press. Paper ID sBp_3.

II.ii **IRIDIA's Arena Tracking System.** The main purpose of the tracking system we implemented at the IRIDIA Lab is to allow a researcher to record and control the state of the experiment throughout its complete execution. Other than experimental analysis, the tracking sys-

tem has also been extended to allow augmented reality for robots (see (Reina et al., 2015b)). The IRIDIA's Arena Tracking System details are presented in the following technical report:

- A. Stranieri, A.E. Turgut, M. Salvaro, G. Francesca, **A. Reina**, M. Dorigo, M. Birattari. IRIDIA's Arena Tracking System. *Technical Report TR/IRIDIA/2013-013*, IRIDIA, Université Libre de Bruxelles, Brussels, Belgium, revision r003 July 2014.

III. A SWARM ROBOTICS PERSPECTIVE TO THE DEPLOYMENT AND REDEPLOYMENT OF WIRELESS SENSOR NETWORKS. The deployment of wireless sensor networks (WSNs) is a significant problem that may determine the correct functioning of the whole sensing system. In recent WSN systems, nodes have a certain degree of autonomy and mobility which makes such systems very similar to swarm robotics systems in which relatively simple agents rely only on local sensing and communication to perform their task. In this work, we review the swarm robotics literature regarding coverage, exploration and navigation tasks to highlight the link between the WSN deployment problem and swarm robotics algorithms. We illustrate the challenges and opportunities offered by swarm robotics with respect to the deployment of mobile WSNs, and we identify relevant directions for a hybridization of WSN and swarm robotics research.

The work has been published as Chapter 7 of the book *Wireless Sensor and Robot Networks*:

- **A. Reina** and V. Trianni. Deployment and redeployment of wireless sensor networks: a swarm robotics perspective. In N. Mitton and D. Simplot-Ryl, editors, *Wireless Sensor and Robot Networks*, chapter 7, pages 143-162. World Scientific, Singapore, 2014.

IV. AUTOMODE: AUTOMATIC DESIGN OF CONTROL SOFTWARE FOR ROBOT SWARMS. In this work, the original idea has been contributed by colleagues of mine, while my contributions have been limited to: (i) give help in the definition of the experimental protocol, (ii) define one of the five experimental tasks, (iii) design and implement algorithms for two swarm robotics tasks, (iv) design and implement part of the infrastructure needed to perform the robot experiments, and (v) edit a limited part of the articles.

This work presents a novel tool to design and generate automatically the control software of robot swarms. This tool is named *AutoMoDe-Chocolate* and is the extension of a previous tool, *AutoMoDe-Vanilla*. The tool makes use of a set of user-defined code blocks which get combined via an optimisation algorithm to create control software for a robot swarm. The optimisation algorithm evaluates the control software

through a physics-based robot simulator and selects the control software that maximises a user-defined performance function. In this work, the generated control software is evaluated in real-robot experiments. The performance of the control software generated by AutoMoDe-Chocolate has been contrasted with the performance of robots that have been programmed by humans expert in swarm robotics. We show that, in this set of experiments, the automatic tool outperforms the human expert designers. In this work, I acted in the role of *human expert*. This work resulted in the following publications:

- G. Francesca, M. Brambilla, A. Brutschy, L. Garattoni, R. Miletitch, G. Podevijn, **A. Reina**, T. Soleymani, M. Salvaro, C. Pinciroli, F. Mascia, V. Trianni, and M. Birattari. AutoMoDe-Chocolate: automatic design of control software for robot swarms. *Swarm Intelligence*, 9 (2–3):125–152, 2015.
- G. Francesca, M. Brambilla, A. Brutschy, L. Garattoni, R. Miletitch, G. Podevijn, **A. Reina**, T. Soleymani, M. Salvaro, C. Pinciroli, V. Trianni, and M. Birattari. An experiment in automatic design of robot swarms: AutoMoDe-Vanilla, EvoStick, and human experts. In M. Dorigo et al., editors, *Swarm Intelligence (ANTS 2014)*, volume 8667 of LNCS, pages 25–37. Springer, Berlin, Germany, 2014.

v. CO-SUPERVISED MASTER OF SCIENCE THESES. During my doctoral studies, I co-supervised five Master of Science students during their final thesis work. All of them successfully defended their theses and completed their studies. In most cases (marked with •), I contributed to the conception of the original idea and I co-directed the research work. In the works marked with ◦, I contributed to the definition of the research problem and I assisted the student in the implementation of the experiments. Here, the list of students and thesis titles:

- ◦ Davide Brambilla. *Environment Classification: an Empirical Study of the Response of a Robot Swarm to Three Different Decision-Making Rules*. M.Sc. Thesis in Computer Science Engineering, Politecnico di Milano, Italy, 2015. Supervised together with Marco Dorigo, Gabriele Valentini and Anthony Antoun.
- Anthony Debruyn. *Human-Swarm Interaction: An Escorting Robot Swarm that Diverts a Human away from Dangers one cannot perceive*. M.Sc. Thesis in Computer Science Engineering, Université Libre de Bruxelles, Belgium, 2015. Supervised together with Mauro Birattari and Gaëtan Podevijn.
- Mattia Salvaro. *Virtual sensing technology applied to a swarm of autonomous robots*. M.Sc. Thesis in Computer Science Engineering, Uni-

versità di Bologna, Italy, 2015. Supervised together with Mauro Birattari and Gianpiero Francesca.

- Bernard Mayeur. *A tool for automatic robot placement in swarm robotics experiments*. M.Sc. Thesis in Information Technology, Université Libre de Bruxelles, Brussels, Belgium, 2014 (in French). Supervised together with Mauro Birattari and Gianpiero Francesca.

- Jacopo De Stefani. *Spatial Allocation in Swarm Robotics*. M.Sc. Thesis in Computer Science Engineering (Project TIME for double degree), Université Libre de Bruxelles, Brussels, Belgium, 2013. Supervised together with Mauro Birattari and Alessandro Stranieri.

1.3 THESIS LAYOUT

This thesis is organised in seven chapters and two appendices.

In Chapter 2, we situate our work in the literature and we present a thorough review of the investigated area. In Section 2.1, we define the swarm systems, their main characteristics and their significance. In Section 2.2, we illustrate which are the challenges and difficulties in designing swarm systems and we present the state of the art with respect to the design problem. In particular, we focus on design methods that, similarly to the methodology proposed in this thesis, provide guarantees on the system performance.

In Chapter 3, we introduce the idea of design patterns as a methodology for the design of swarm systems. In Section 3.1, we discuss previous works that have similarities to ours under the aspect of using a catalogue of reusable solutions. In Section 3.2, we present the components that constitute a design pattern, and, in Section 3.3, we report the road map for the formalisation of novel design patterns. Part of the work presented in this chapter has been published in (Reina et al., 2015d).

In Chapter 4, we present the intermediate steps we moved to reach the formalisation of our design pattern. In particular, in this chapter, we select the reference model of decentralised decision making from the literature (motivating this choice), and we implement a validation case study. The goal of the validation case study is to demonstrate the possibility of attaining a precise quantitative match between the dynamics of the selected macroscopic model and the dynamics of the implemented multiagent system. For doing so, in Section 4.1, we sketch the implementation guidelines (which later will be formalised in a design pattern). Then, in Section 4.2, starting from these guidelines, we implement in simulation an abstract particle system and a swarm robotics system. Finally, in Section 4.3, we show the attainment of a precise quantitative match between the macroscopic model dynamics and the dynamics of the two implemented systems. The work presented in this chapter has been published in (Reina et al., 2015a).

In Chapter 5 and 6, we present the complete design pattern for decentralised decision making in the best-of-n problem. A formal description of the design pattern is provided in Chapter 5, which discusses models and guidelines to support the implementation of best-of-n decision making. The chapter is composed of five sections, each presenting an attribute of the formalised design pattern. An important element of design patterns are implementation examples that showcase its usage. Chapter 6 is dedicated to describe two case studies that complement the formal description given in Chapter 5. The performance of the systems implemented in these case studies is evaluated through a set of metrics which are defined in Section 6.1. The first case study, presented in Section 6.2, shows decentralised decision making by static agents that interacts on a fully-connected network. Thanks to the simplicity of this case study, we are able to show results for several parameterisations, implementation strategies and option qualities. The second case study is presented in Section 6.3 and illustrates the design of a multiagent system in a more challenging scenario involving spatiality. The work presented in these chapters has been published in (Reina et al., 2015d).

Finally, in Chapter 7, we summarise the contributions of this thesis and we outline the future research work.

Appendix A introduces a methodology, based on survival analysis, to estimate from experimental data the transition rates of a macroscopic model. Appendix B shows the stability analysis and the mathematical details for selecting the parameterisation of the case study presented in Section 6.2.1.

Note that the contents of Chapters 4, 5, and 6 and Appendices A and B substantially match the text of the two papers (Reina et al., 2015a) and (Reina et al., 2015d).

Student: "Dr. Einstein, aren't these the same questions as last year's final exam?"
Dr. Einstein: "Yes; but this year the answers are different."

—Albert Einstein

2

THE DESIGN PROBLEM

This thesis describes an engineering methodology for the design of swarm systems. In this chapter, we situate our work in the literature. We first describe swarm systems, their characteristics and some of their possible applications in Section 2.1. Then, in Section 2.2, we motivate our work by illustrating the challenges related to the engineering and the design of swarm systems, and we present some of the main solutions in the literature. We organise these solutions in two opposite (although very coupled) approaches: the bottom-up approach in Subsection 2.2.1 and the top-down approach in Subsection 2.2.2.

2.1 SWARM SYSTEMS

Swarm systems are composed of a large number of autonomous agents that interact with their environment and with each other for achieving a common goal without relying on any central controller (Haykin, 2005; Hatziargyriou et al., 2007; Nakamoto, 2008; Dressler, 2008; Driesen and Katiraei, 2008; Rajkumar et al., 2010; Dorigo et al., 2014). Such systems are based on the principles of swarm intelligence (Bonabeau et al., 1999; Dorigo and Birattari, 2007) and their functioning is based on the underlying idea that a self-organisation process allows the swarm to display an intelligent collective behaviour. In contrast to classical monolithic systems composed of a single agent that must be endowed with all the capabilities necessary to conduct the designed task, swarm systems exploit self-organisation to display swarm-level capabilities beyond the ones of a single agent. Therefore, swarm systems can perform tasks that are more difficult than the tasks that a single agent could deal with. Based on this concept, swarm systems profit of a simpler individual agent design because the desired system capabilities are shifted from the agents to the swarm. In other words, the individual agent capabilities can be reduced, in exchange for a more complex collective behaviour that arises from self-organisation. This kind of approach results in reduced hardware costs as a consequence of fewer agent require-

ments, and could allow future swarm systems to operate in environments that are currently inaccessible due to either extremely reduced dimensions (e.g., nanoscale) or to particularly isolated environments (e.g., space or ocean floor). For instance, an envisioned application of swarm systems concerns the injection in human bodies of nanorobot swarms to perform healing operations (Hauert et al., 2013; Hauert and Bhatia, 2014). Here, the individual agent hardware is clearly constrained by its limited dimensions and has to rely on a minimalistic design (thus minimal capabilities). Such minimalistic agents may exploit swarm intelligence principles to achieve their goal. A second application example is underwater operations (Kalantar and Zimmer, 2006; Fiorelli et al., 2006; Leonard et al., 2007). In this case, the constraints on the agent capabilities are imposed by underwater communication limitations that do not allow a centralised human control, global communication or any global positioning system. Swarm systems are a potential approach to deal with such constraints.

Further advantageous characteristics that, in the literature, are normally ascribed to swarm systems are adaptivity, fault tolerance, parallelism and scalability (Brambilla et al., 2013; Hecker and Moses, 2015).

ADAPTIVITY. We refer to adaptivity as the capability of the swarm to adapt to changing environments and/or to unpredicted conditions. We expect this capability to derive from the self-organisation process which allows the swarm to independently organise itself in different ways according to the given environmental circumstances. Adaptivity can be useful in scenarios that cannot be known and accurately modelled in advance, such as disaster areas (Ghassemi et al., 2010; Athreya and Tague, 2012), space explorations (Truszkowski et al., 2006), or scenarios that may change rapidly over time, such as environments involving human crowds (which may be considered by the swarm as obstacles that continuously change shape and position) (Guzzi et al., 2013).

FAULT TOLERANCE. We refer to fault tolerance as the system property to keep functioning in case of failure of part of the agents. Given the decentralised nature of swarm systems, there is no central point of failure and none of the agents is indispensable to the functioning of the system. Agent failures result in a graceful deterioration of the system performance, yet without an overall task discontinuation. This property derives from the high redundancy of swarm systems which are composed of a large number of interchangeable agents (Christensen et al., 2009). Application areas that may benefit from fault tolerant systems are characterised by hazardous environments, such as war battlefields or disaster areas (Athreya and Tague, 2012; Dames et al., 2012), or by environments out of human reach where main-

tenance is impossible, such as space or underwater exploration missions (Truszkowski et al., 2006).

PARALLELISM. We refer to parallelism as the system property of processing multiple pieces of information and/or performing multiple operations at the same time. This property is attained exploiting the large number of agents (of which the swarm is composed) that can autonomously act in the environment. An agent may perform its tasks either individually or joining other agents to create work-teams. Parallelism may be particularly beneficial in applications that involve the repetition of the same operation several times. For instance, construction of buildings is an application that typically requires the sequential accumulation of building blocks (e.g., bricks) to create a solid structure; here, parallelism may have a crucial role on the system speed (Soleymani et al., 2015).

SCALABILITY. We refer to scalability as the system property to function with different swarm sizes (i.e., the number of agents in the swarm) without the need of re-programming or tuning the agent's behaviour. The local, decentralised, self-organising process allows the swarm to cope with varying swarm size and does not require the agents to have global knowledge of the state of the process or of the total number of agents. Scalability allows the employment of swarm systems in applications that require the increment (or decrement) of agents at runtime.

The above properties are desired in several domains. However, a real world application of swarm systems is hindered by the difficulty in modelling and designing such systems. In the literature, the modelling and the design of swarms have been the focus of several studies which we present in the next section.

2.2 ENGINEERING SWARM SYSTEMS

A recent survey (Brambilla et al., 2013) of swarm robotics states that the implementation of swarm systems is rarely assisted by engineering methods that guide the design process. More commonly, the design relies on the ingenuity of the human designer for finding a satisfactory solution through an unsupervised trial-and-error approach. This approach is clearly very unstructured and is solely based on the expertise and intuition of the designer.

An alternative and more structured approach, that has been the subject of several works in the last decade, consists in using automatic methods (Panait and Luke, 2005; Trianni, 2008; Francesca and Birattari, 2016). These methods, through metaheuristic optimisation, automatically generate and test a large number of possible individual agent's behaviours to determine the behaviour that optimises a task-specific target metric. This target metric

is defined by the designer and is tailored to the specific task. The performance of the generated solution is evaluated through a set of simulation experiments, and eventually tested on real systems. Even if these methods have empirically shown to be able to generate in an automatic way swarm systems for a variety of tasks (Francesca et al., 2014b), there are no guarantees on the generated system's performance.

Despite most swarm systems have been implemented through heuristic techniques (such as the trial-and-error process, or metaheuristic optimisation), we believe that engineering methodologies that provide performance guarantees are needed to enable the employment of swarm systems in real world applications. Our work follows this direction by proposing an engineering methodology that provides guarantees in statistical terms for the provided solution. Hereafter, we give an overview of the main works in swarm robotics that investigated the problem of ensuring a particular level of performance through engineering tools. We focus on swarm robotics systems because they exemplify the main challenges of swarm systems and a large number of works investigated their design. The design challenges discussed at the end of the section pertain also to other swarm systems, such as cyber-physical systems (Ilic et al., 2010; Lee et al., 2014; Lee, 2015), mobile wireless sensor networks (Leonard et al., 2007; Dressler, 2008), distributed smart grids (Rohden et al., 2012; Ma et al., 2013; Olivares et al., 2014) or cognitive radios (Haykin, 2005; Akyildiz et al., 2006; Trianni et al., 2016a).

2.2.1 *Micro to macro*

The most common approach to providing guarantees on the system performance is to create a model of the collective dynamics starting from the individual agent implementation, and then analyse such a model. As already mentioned in Section 2.1, the main challenge in modelling swarm systems is attaining an accurate and well-formalised relationship between the macroscopic system dynamics and the microscopic (individual-agent) behaviour. In the literature, this relationship is typically called the *micro-macro link* (Schillo et al., 2001; Hamann and Wörn, 2008b). Attaining the micro-macro link is a challenging task given the difficulties to develop and analyse nonlinear systems composed of numerous interacting parts. It may be even more difficult due to heterogeneities in the interaction topology, in the agent behaviour, and/or in the environmental features. In order to attain treatable models, several simplifications and abstractions are necessary. As a result, the macroscopic model usually describes the system dynamics only qualitatively, and extensive tuning may be required to obtain an appropriate parameterisation for the real system.

Several works successfully derived a micro-macro link by creating a macroscopic model from the microscopic implementation of the individual

agents (Parker and Zhang, 2009, 2010; Montes et al., 2010; Valentini et al., 2014; Wilson et al., 2014; Sartoretti et al., 2014; Scheidler et al., 2015). With this bottom-up approach, the resulting macroscopic model is tailored to the investigated problem and allows the analysis of only the system from which the model is derived. In the case of a change in the problem, the implementation is modified accordingly and thus the model must be consistently adjusted and the analysis repeated. Each iteration of this process may be very resource demanding due to the difficulty of attaining an accurate quantitative micro-macro link. For instance, the studies presented by Montes et al. (2010) and by Scheidler et al. (2015) present a differential latencies model to describe and analyse the dynamics of a binary decision making process of mobile agents. The investigated case study is very similar to the case study that we present in Chapter 4 and their solution is a valid alternative to solve that problem. However, their model is specific to differential latencies problems and cannot generalise to the other decision making case studies that we considered in this thesis (e.g., the case study II of Section 6.3).

In the following, we list notable modelling techniques that we think are worth discussing in more details because of the tractability of the produced models and of their widespread usage in a variety of swarm implementations. Note that in these works the macroscopic model is derived from the microscopic agent behaviour, therefore following a (tailored-to-solution) bottom-up approach.

Martinoli et al. (2004) and Lerman et al. (2005) present a methodology to mechanically develop microscopic and macroscopic models in good correspondence with each other—i.e., attaining the micro-macro link—for swarms of homogeneous robots in a homogeneous environment. The method consists in deriving from a probabilistic finite state machine (PFSM) a master equation and a rate equation (as a system of ODEs). The PFSM describes (or directly corresponds to) the individual robot control software. The master equation describes the temporal evolution of an individual robot's probability density. The rate equation describes the average proportion of agents in each state of the PFSM. This approximation assumes a well-mixed system (i.e., interactions between any couple of agents are equally probable) and an infinite size of the swarm (i.e., infinite number of robots, therefore, the model prediction are more accurate for large swarm size). Correll and Martinoli (2006) show that, through heuristic optimisation of the macroscopic model parameters, it is possible to increase the predictive accuracy of the model and to attain a quantitative match between dynamics at the microscopic and the macroscopic level. The authors of these works exploited their modelling methodology to analyse a variety of swarm robotics case studies among which collaborative stick pulling (Martinoli et al., 2004), collaborative object collection (Agassounon et al., 2004), task allocation (Ler-

man et al., 2006), area coverage (Winfield et al., 2008) and aggregation (Correll and Martinoli, 2011).

Hamann and Wörn (2008b) obtain the micro-macro link by modelling robot swarm dynamics through statistical physics methods, in particular through Fokker-Planck equations. This modelling approach allows an explicit representation of space, and thus, takes into account heterogeneous spatial distribution of agents. Their method results in a good qualitative correspondence between the predictions of the macroscopic dynamics and multiagent simulations. However, this method presents a few issues: (i) achieving quantitative correctness may require a considerably big amount of computational resources, (ii) Fokker-Plank modelling requires a significant effort for the designer, and (iii) only basic communication can be modelled. Despite these points, Fokker-Plank modelling has been successfully employed to support the design of various case studies: aggregation (Schmickl et al., 2009; Hamann et al., 2008; Hamann and Wörn, 2008a), collective perception (Hamann and Wörn, 2008b), phototaxis (Hamann and Wörn, 2008b), complex environment exploration (Prorok et al., 2011) and foraging (Hamann and Wörn, 2007).

In (Hamann et al., 2014), the authors devise a mathematical representation of a decentralised decision-making system in the form of a drift-diffusion model (DDM). Through recursive equations, they find in an automatic way the drift coefficient. This method as been validated in modelling a single case study of the binary decision of locusts motion. At the current stage, this method allows the attainment of only a qualitative micro-macro link and works with the assumption of a memoryless (i.e., Markovian) and well-mixed system. However, it has the potential of being extended to become a more powerful modelling tool.

Vigelius et al. (2014) propose a framework to derive from a multiagent implementation a macroscopic description of the system. In particular, the system is described through a master equation which allows the analysis of the finite-size effects (i.e., how the systems behaves when the swarm is composed of a specific number of agents). In the paper, the framework is validated through a decentralised decision-making case study and the microscopic implementation includes both simulation and robot experiments. The attainment of a good agreement between the two levels (micro-macro) is achieved only for memoryless and well-mixed systems. Since spatiality hinders the satisfaction of these two properties, the micro-macro link is not accurate in some of the scenarios that involve spatiality. Similar modelling techniques have been employed in other works investigating decentralised decision making (Massink et al., 2013; Valentini et al., 2014, 2016; Scheidler et al., 2015). Also in these works, the decision process has been modelled both through an ODE system using a mean-field approximation, and a master equation to investigate the finite-size effects. The results have

been validated through robot experiments (Scheidler et al., 2015; Valentini et al., 2016). Similarly, other tools of statistical physics have been employed to analyse the coordinated motion of swarms of simple, reactive agents (Ferrante et al., 2013). In this work, the authors employ non-linear elasticity theory to investigate the phase transitions of their system.

Khaluf and Dorigo (2016) show how some swarm systems can be modelled as a linear birth-death process. This modelling technique allows the prediction of the swarm performance in terms of the average activity time (that is, how long the swarm will be operative before stopping to function) and the swarm energy consumption.

A recent work proposes a novel model checking procedure based on the mean-field approximation to facilitate the modelling of swarm behaviours by increasing the level of abstraction (Latella et al., 2015). This procedure assumes infinite swarm size (thus it only approximates results of finite size swarms) and requires a well-mixed system.

2.2.2 *Macro to micro*

An alternative approach consists in deriving a microscopic implementation from a model that describes the target macroscopic dynamics. The main advantage of this top-down approach is that the engineer can, at first, design the desired swarm system behaviour at a macroscopic level, and then, (in an automated way) convert the swarm behaviour to individual agent behavioural rules. The methodology must provide the engineer with tools that guarantee the attainment of a quantitative accurate micro-macro link, thus a predictable global behaviour.

Unfortunately, a general top-down methodology to automatically provide the micro-macro link for any given macroscopic description of the desired collective behaviour is currently unavailable. However, some works in the literature have proposed various solutions that follow a top-down design paradigm. Hereafter, we organise these works in two categories with respect to the approach used to provide performance guarantees: population dynamics models and control theory.

POPULATION DYNAMICS. Through this approach, the macroscopic model describes the evolution over time of the proportions of swarm subgroups (subpopulations) with respect to the total swarm size (population). Agents are grouped into subpopulations according to either similar internal state or proximal location. Usually, the model is a set of rate equations described as an ODE system or chemical reactions, and the subpopulation variations are determined by transition rates. The micro-macro link is established by mapping these rates into probabilities to be used in the control software of the individual agents.

In (Berman et al., 2009), the authors formalise a design method to allocate mobile robots to tasks using rate equations. The rate equations are in the form of an ODE system that describes how the number of agents allocated to each task varies; while the transition rates describe how frequently robots change their task. Previous to this work, this method had been successfully employed to model a set of foraging scenarios (Berman et al., 2007b,a; Halász et al., 2007; Hsieh et al., 2008). Afterwards, the design method has been used to implement robot swarms control software for two simulated task allocation scenarios (Berman et al., 2011b; Dantu et al., 2012). However, there are some criticisms of this methodology. A comparative experimental analysis of this methodology evidenced that the underlying assumptions —such as infinite number of robots, perfect localisation, global communication— may be unrealistic for most swarm systems (Mermoud et al., 2014). The analysis shows that the performance of a robot system implemented through this methodology does not match the performance predicted by the macroscopic model, at least in the investigated experiments.

The methodology described above is extended in (Berman et al., 2011a) to deal with spatially inhomogeneous swarms, that is, considering the spatial distribution of agents throughout the space. Here, the swarm is modelled as a system of advection-diffusion-reaction (ADR) partial differential equations (PDEs), which is numerically solved through the technique of smoothed particle hydrodynamics.

Brambilla et al. (2015) propose an high-level method for top-down design of robot swarms based on model checking. The method consists in organising the design in four phases: requirements, model, simulation and implementation on robots. In the first phase, the designer formally specifies the system requirements; in the second phase, (s)he creates a (population dynamics) model of the robot swarm's dynamics; in the third phase, (s)he implements physics-based computer simulations of the systems; and, finally, the fourth phase concerns the implementation on a swarm of robots. Additionally, similarly to our work, the authors report two case studies to illustrate the usage of the proposed method. The advantage of this design method is its generality because, in principle, it can be employed to design any swarm robotics system. Shortcomings of this method are that it gives guidelines at a very-high level, it provides limited support to guide the system implementation and it requires a big creative effort from the designer to follow the four design phases. Instead, the methodology that we propose in this thesis is less general, in the sense that each design pattern gives a solution to a specific class of problems. However, our methodology provides a precise set of rules to derive, in an automated way, the individual agent behavioural rules from the macroscopic model. Finally, we believe that the method proposed by Brambilla et al. (2015) might be a good method to as-

sist the formalisation of a novel design pattern (see the second step defined in Section 3.3).

CONTROL THEORY. A number of works make use of a control-theoretic approach to define control laws for each agent and derive formal proofs of the global system dynamics. Through such proofs, the engineer may attain formal guarantees on the swarm behaviour (typically regarding the system stability). This approach is grounded in solid mathematical basis; however, its usage in the design of swarm systems presents big challenges. The control-theory approach is normally based on closed-loop systems, in which the system behaviour is modified through time as a function of the current measured state. In swarm systems, a measure of the current state is not easy to attain. Most of the early works assume each agent of the swarm capable of attaining at each timestep a measure of the global system state and of adapting its behaviour accordingly (Desai et al., 2001; Leonard and Fiorelli, 2001; Ögren et al., 2001, 2004; Bachmayer and Leonard, 2002; Gazi and Passino, 2002, 2003, 2004; Lawton et al., 2003; Cortes et al., 2004; Gazi, 2005; Sepulchre et al., 2007; Leonard et al., 2007; Michael and Kumar, 2009; Bhattacharya et al., 2014). A review of most of these methods and their limitations is presented in (Gazi and Fidan, 2007). Additionally, all these works focus on behaviours aimed at achieving some form of spatial organisation of the agents (e.g., pattern formation, coverage), where the measured state is the relative position of each agent. The main limitation of these works comes from the assumption of an unbounded sensing range. That is, every agent is capable of locating any other agent at any distance. This global information assumption may be unrealistic in several applications with real robots (or mobile sensors), and may present scalability issues when the number of robots is very large.

Accounting for only local sensing and communication while ensuring guarantees has been possible only by reducing the attainable guarantees (e.g., Ji and Egerstedt, 2007; Kennedy et al., 2015) or for simpler tasks (e.g., McNew et al., 2007; Sepulchre et al., 2008).

In (Ji and Egerstedt, 2007), the agents can only perceive neighbours within a local range, (while they still have access to a common axis of reference for global orientation). Under such assumption, the authors propose a solution to the formation control problem that provides guarantees only for the connectedness of the neighbour network.

Kennedy et al. (2015) propose a distributed control algorithm for the collective transport of a target object by a swarm of robots. The robots can only locate and communicate with neighbours within a local range, however, each robot knows its absolute orientation and its position with respect to the target object, and can estimate the net wrench being exerted on the object. The authors provide the formal proof of the asymptotic convergence

to a force distribution with zero interaction forces, i.e., without exerting opposite (and conflicting) forces on the object.

McNew et al. (2007) propose an algorithm, that is based only on the relative positioning of neighbour agents (i.e., taking in consideration a limited sensing and communication range), to organise the agents in a triangular lattice. The algorithm is specified in Embedded Graph Grammars (EGGs) which allow having guarantees of convergence to the target spatial configuration. This deployment and coordination algorithm has been validated through the implementation on a multi-robot system (Smith et al., 2009). The attainment of guarantee proofs is very difficult with algorithms based on EGGs. For this reason, this approach has been used only for relatively simple problems.

In (Sepulchre et al., 2008), the authors propose the use of Lyapunov functions as individual agent control laws to stabilise steering particles moving on a plane. This work extends a previous approach by considering local communication among agents. As the authors say, extending the approach to more complicated models may be intractable: *"An open question of particular interest is the robustness of the proposed approach to more complicated models for individuals [...]"* (Sepulchre et al., 2008).

Srivastava and Leonard (2013, 2015) propose a very interesting study of decentralised decision making in which they formally prove the expected decision time and its variance as a function of the individual agent parameters. Each agent behaviour is modelled as a drift-diffusion model (DDM), the system dynamics as a coupled drift-diffusion equation, and the interaction network is defined through a Laplacian matrix. Here, the micro-macro link is attained without the need of a mean-field approximation. This design method provides a solid basis for the engineering of swarm systems; however, the current work still presents some limitations. In particular, the study is limited to a binary decision problem, the interaction network is static and the DDMs are linear. Notwithstanding, this work paves the way for a promising direction of research for the design of swarm systems.

2.2.3 *Discussion of engineering methods for swarm systems*

In Sections 2.2.1 and 2.2.2, we have seen a series of different endeavours to obtain an accurate quantitative micro-macro link to provide performance guarantees for swarm systems. These works proposed different methods, explored the use of various techniques and tackled problems of various difficulty and complexity. Despite the numerous efforts, a general methodology is missing. We can conclude that the attainment of such a micro-macro link is a very challenging task. The difficulties are due to the fact that: (i) macroscopic modelling is limited by the available theory and methodologies to deal with the non-linearity and stochasticity of the system; and (ii) the de-

sign of the individual behaviour should account for the decentralised nature of the system based on limited agent communication and sensing capabilities, and on the asynchronous, parallel processing of partial information.

While a general method to derive individual agent rules from any macroscopic model is currently unworkable, general solutions to specific classes of problems can be provided exploiting the concept of design patterns. Design patterns provide formal guidelines to deal with recurring problems in a specific field. For the swarm systems field, design patterns prescribe the individual-level microscopic behaviour required to obtain desired system-level macroscopic properties. In the next chapter, we present a design methodology based on design patterns.

The important thing in science is not so much to obtain new facts as to discover new ways of thinking about them.

—William Lawrence Bragg

3

DESIGN PATTERNS AS AN ENGINEERING METHODOLOGY FOR DECENTRALISED SYSTEMS

In Section 2.2, we have seen that, for most swarm systems, obtaining a quantitative micro-macro link to provide guarantees on the system performance is very complicated and may require much effort. While a universal design methodology is currently out of reach, a viable approach consists in having reusable methods to top-down design swarm systems that tackle specific classes of problems. To be reusable, a method should aim to (i) be general, (ii) tackle a recurrent (relevant) problem, and (iii) follow a reusable formalisation. Formalising and organising the methods on a common underlying structure allows an easier spread and adoption of them. Finally, a collection of such methods would represent a methodology which would guide engineers in the swarm system design. We propose to formalise each reusable method as a *design pattern*. The underlying idea is to have a consistent way of representing packages of knowledge that are mainly composed of a problem, a solution and the micro-macro link.

The idea of having a catalogue of reusable solutions in the form of design patterns has been originally proposed in the domain of architecture (Alexander et al., 1977) and it has been then exported in various domains ranging from pedagogy (Jones et al., 1999) to software engineering (Gamma et al., 1995). We believe that the concept of design patterns as an engineering methodology should be exported to the swarm system field too. This idea has been acknowledged in a number of works in the literature, which we cover in Section 3.1. Differently from previous work, our methodology pivots on the formalisation of the micro-macro link. We present our interpretation of design patterns for swarm systems in Section 3.2. Finally, in Section 3.3, we present the procedure to formalise a novel design pattern.

3.1 A CATALOGUE OF REUSABLE SOLUTIONS

A first attempt to organise design solutions for swarm systems in a catalogue has been made by Nagpal (2004). In this work, she lists a series of solutions to attain various collective behaviours. However, we can consider this as a preliminary work in this direction as it does not include any formal problem/solution description nor tackles the problem of guaranteeing the attainment of desired performance levels.

Babaoğlu et al. (2006) propose a first design pattern catalogue for decentralised systems. The work consists in the definition of five design patterns for decentralised systems and a set of case studies to illustrate their usage in various scenarios. Similar works have extended the catalogue or redefined some of the existing patterns (Gardelli et al., 2007; De Wolf and Holvoet, 2007). The main limitations of these works are the lack of a precise formalisation of problems and solutions, and of the micro-macro links that connect the macroscopic behaviours to the individual agent rules. The absence of these components does not allow having any guarantee on the system performance. With our formalisation of the design pattern, we aim to overcome these limitations and we propose methods to attain an accurate quantitative micro-macro link.

The work of Fernandez-Marquez et al. (2013) extends the design pattern approach by introducing a more accurate formalisation of the solutions and providing a description of the effects of the individual parameters on the global outcome. Despite these improvements, this work does not provide tools to guarantee performance levels, yet. The work gives guidelines to the design and the implementation of swarm systems without providing formal rules to attain a micro-macro link, which we believe is a necessary requirement for having performance guarantees. An interesting aspect of this work is the study of interactions—dependencies and interference—among design patterns. We believe that interactions between patterns are a very important aspect that needs to be deeply investigated. In our work, we do not tackle this aspect since we present only a single design pattern. However, we plan to perform studies in this direction once a set of patterns will be developed.

Field calculus (Beal and Viroli, 2015) is a framework that allows programming swarm systems through an amorphous medium abstraction. Through this abstraction, the system (that is composed of a finite number of agents) is modelled as a space continuous medium (Beal, 2005). The field calculus employs Proto (Beal and Bachrach, 2006) as spatial language, and provides a set of *building blocks* as basic macroscopic functions that define simple behaviours at the swarm level. The main advantage of this programming paradigm is that the designer deals only with building blocks at the macroscopic level, while the implementation of the microscopic behaviour is hidden and is managed automatically and implicitly by the aggregate program-

ming framework. These building blocks can be used as language operators to compose algorithms for collective behaviours. Similarly to design patterns, these algorithms for collective behaviours, called *libraries*, represent reusable solutions to recurrent collective problems. Differently from our work, field calculus focuses primarily on the aspects of the implementation to provide an aggregate programming framework and allows a mathematical validation of only a few general properties, such as spreading process stabilisation over a network (see (Viroli and Damiani, 2014) for further details). Instead, in our work, the design pattern is meant to provide mathematical tools to guarantee certain levels of system performance. At the moment, this aspect is missing in field calculus. Furthermore, field calculus is still at a preliminary stage, no libraries have been proposed yet and the current number of available building blocks is quite limited; for instance, there is no support (yet) for functionalities to control the agent movement (that is of utmost importance in domains such as swarm robotics).

3.2 THE DESIGN PATTERN AND ITS ATTRIBUTES

While the concept of design patterns in swarm systems has been already employed, previous work has devoted limited attention to guaranteeing the attainment of desired performance levels. Instead, we believe that providing performance guarantees, at least in statistical terms, is a key requirement for the engineering of swarm systems. Therefore, in our definition of design patterns, the key component is the micro-macro link formalisation. A quantitative micro-macro link allows the analysis of the macroscopic model to design and predict the swarm behaviour.

For consistency with the literature, we organise the design pattern in the same six attributes as Babaoğlu et al. (2006): name, problem, context, design rationale, solution, and case-studies. However, the definition of these attributes (especially the solution) differs from previous work. The *name* identifies the design pattern, and should be informative enough to summarise its objective. The *problem* formally describes the purpose of the design pattern, and possibly gives constraints on the system that limit the domain of the possible solutions (e.g., constraints on agent's memory or communication capabilities). The *context* determines the domain of applicability of the design pattern, and presents a set of preconditions that must be fulfilled for its usage. The *design rationale* explains the what and the how of the design pattern, that is, its origin and working principles. Typically, the design rationale contains a description of the basic principles underlying the proposed solution, together with some insight about why it efficiently functions. The *solution* provides tools and guidelines for the implementation of the decentralised system. The core of the solution is represented by the formal description of the decentralised process at different abstraction

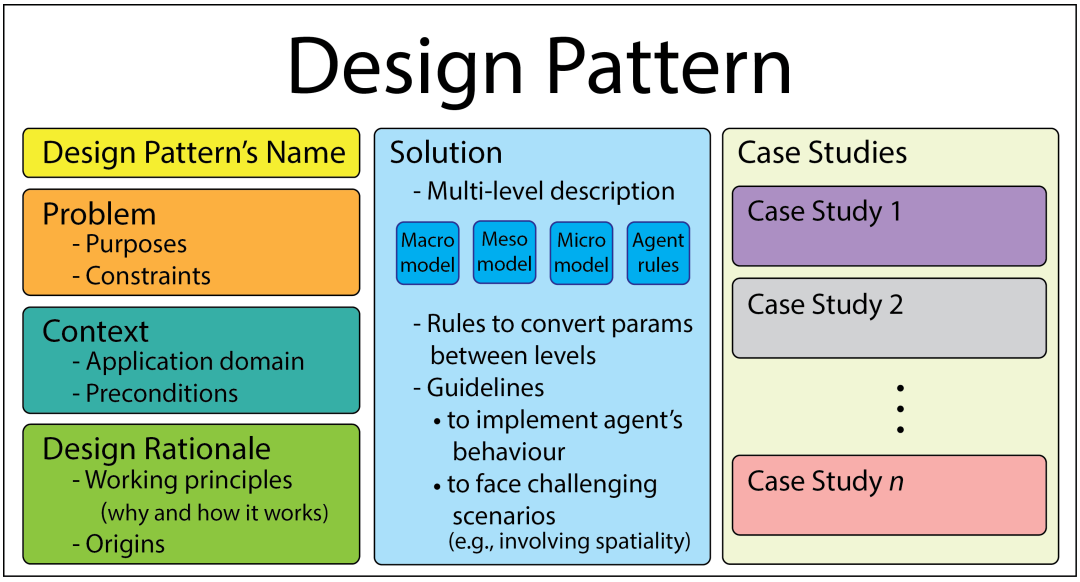

Figure 1: Schematic representation of a design pattern for swarm systems composed of six attributes: name, problem, context, design rationale, solution, and case-studies.

levels—from macroscopic to microscopic—together with the relationship between them. Additionally, a set of recommended implementation strategies describe how to deal with problems recurring in practical application scenarios, which may have a bearing on the relationship between microscopic and macroscopic description levels (i.e., the micro-macro link). Finally, a design pattern includes a set of *case studies* along with a thorough evaluation to showcase the functioning of the design pattern. Figure 1 shows a graphical representation of the components of a design pattern.

3.3 HOW TO BUILD A DESIGN PATTERN

In Chapter 2, we contrasted analysis with design methods. In general, analytical models are tailored on a specific system implementation and developed once the system is implemented. Instead, design methods typically rely on models that are independent of a specific implementation or scenario. However, in several cases, analysis and design share the same modelling tools (e.g., master equations, Markov chains), and the distinction between analysis and design methods in some cases may be very subtle.

These similarities are evident in the process of formalisation of a design pattern. As defined above, a design pattern is a design method that provides the engineer with tools to attain an accurate micro-macro link for the implemented solution. The micro-macro link is typically the outcome of an analysis process. Therefore, defining a design pattern consists in formalising the analysis of a general system as a reusable design methodology.

We organise the process of formalisation of a novel design pattern for swarm systems in three steps. First, the identification of well conceived macroscopic models is performed. These models should have been deeply analysed, have an engineering relevance (i.e., potentially useful for something) and generally well understood. In fact, one of the main underlying ideas of design patterns for swarm systems is to leverage the principled understanding of theoretical models to engineer artificial systems. Second, at least one validation case study should be implemented to demonstrate the possibility of attaining a micro-macro link based on the target model. Here, the micro-macro link is attained through analysis tools following a bottom-up approach. During the third, final step, the analysis must be formalised and extended to obtain rules to convert parameters between models at different abstraction levels (macro-micro). In addition, the knowledge must be organised in the six attributes (see Figure 1), which include a set of case-studies. In this thesis, we present the identification of the macroscopic model and the validation case study in Chapter 4, a complete design pattern for the best-of-n, decentralised decision making in Chapter 5 and two further case studies in Chapter 6.

Well-conceived case studies are a necessary milestone for the development of a design pattern, and they are as important as the theoretical characterisation of the micro-macro link that serves as guidance for the implementation. Indeed, case studies demonstrate the viability of the methodology by showing practical situations in which the micro-macro link is obtained quantitatively, and also offer a representative case that can be used as a guideline for future implementations. Preferably, they tackle a particularly challenging scenario, for instance involving spatiality which may hinder a match between the two abstraction levels.

The beginning is the most important part of the work.

—Plato

4

TOWARDS THE FORMALISATION OF A DESIGN PATTERN FOR DECENTRALISED DECISION MAKING

Among the classes of cognitive processes that can be addressed through well-conceived design patterns, the best-of-n, decentralised decision making represents a fundamental ability in several contexts and application domains (Halloy et al., 2007; Vigelius et al., 2014; Srivastava and Leonard, 2014; Valentini et al., 2016). In this thesis, we propose a design pattern for this type of decision making based on the nest-site selection behaviour of honeybee swarms (Marshall et al., 2009; Seeley et al., 2012; Pais et al., 2013). Previous experimental and theoretical studies have demonstrated near-optimal speed-accuracy tradeoffs in the selection of the most profitable option among a set of alternative nesting sites by honeybees (Marshall et al., 2009; Seeley et al., 2012). Most importantly, inhibitory signals among bees provide an adaptive mechanism to quickly break deadlocks and tune the decision dynamics according to the perceived quality of the discovered options (Seeley et al., 2012; Pais et al., 2013). The above properties of the nest-site selection process are relevant for many practical decision-making scenarios in decentralised systems, and justify its choice in this study.

In this chapter, we study how to obtain a quantitative correspondence between the dynamics of the microscopic implementation of a robot swarm and the dynamics of the macroscopic model of nest-site selection in honeybees. We do so by considering a decentralised decision-making case study: the shortest path discovery/selection problem (Gutiérrez et al., 2010; Montes et al., 2010; Scheidler et al., 2015). This problem requires the identification and collective selection of the shortest path between target areas, to be performed by a (possibly large) swarm of autonomous robots. In this case study, without loss of generality, we limit the study to a binary decision problem, i.e., the swarm has to select between two alternative paths. This choice allows us to investigate several system parameterisations while keeping the analysis and result interpretation simple. Path discovery/selection is particularly demanding from a decision-making standpoint because spa-

tial factors strongly influence the collective dynamics, as they determine the interaction patterns among the robots. As a consequence, finding a micro-macro link between a non-spatial macroscopic model of decision making and a strongly-spatial decision problem is not trivial. A preliminary study and partial results have been presented in Reina et al. (2014a,b) for a very abstract scenario involving agents moving in a one-dimensional space. Here, we present the complete study of a more realistic scenario for agents/robots moving on a flat surface. The attainment of the micro-macro link presented in this chapter is a prerequisite and a necessary step towards the formal characterisation of the design pattern for the general case of decentralised decisions in the best-of-n problem, which is presented in its complete form in Chapter 5.

This chapter is organised as follows. In Section 4.1, we introduce a set of high-level implementation guidelines inspired by studies of honeybee nest-site selection, which describe the target macroscopic dynamics (Seeley et al., 2012; Pais et al., 2013). In Section 4.2, we report a solution to the shortest path discovery/selection problem both in an idealised multiagent simulation that retains the relevant spatial factors but neglects physical interactions (see Section 4.2.2), and in a physics-based simulation of a swarm of e-pucks (Mondada et al., 2009), therefore accounting for limited swarm size and physical interferences (see Section 4.2.3). In Section 4.3, we analyse the implemented behaviour under a variety of different parameterisations, and show that in every case there is a precise correspondence between the microscopic and the macroscopic descriptions. We extend the analysis towards varying group size, and also show the correspondence with Monte Carlo simulations of a macroscopic finite-size model. As discussed in Section 4.4, the case study presented in this chapter contributes to the formalisation of a design pattern for decentralised decision making by extending the high-level guidelines presented in Section 4.1 with specific recommendations for dealing with issues deriving from spatiality.

4.1 FROM MACROSCOPIC DESCRIPTIONS TO IMPLEMENTATION GUIDELINES

As mentioned above, how collectives can achieve consensus is widely studied in many different contexts, and several models have been proposed in the literature (e.g., Castellano et al., 2009; Vicsek and Zafeiris, 2012; Kao et al., 2014). In this work, we selected a macroscopic model inspired by the nest-site selection behaviour observed in honeybees (Seeley et al., 2012; Pais et al., 2013). We selected this model because it possesses properties that are desirable in artificial decentralised systems: (i) it attains near-optimal speed-accuracy tradeoffs in the selection of the best option (Marshall et al., 2009), and (ii) it exploits adaptive mechanisms to tune decision speed and

to break symmetry deadlocks (e.g., caused by same-quality options). In this section, we first discuss the macroscopic model; then, we present a high-level description of the implementation path prescribed by a design pattern to obtain a quantitative micro-macro link.

4.1.1 *The macroscopic model*

An analytical model of the nest-site selection process in honeybee colonies has been developed and confronted with empirical results, confirming the existence of both positive and negative feedback loops that determine the collective decision (Seeley et al., 2012). The model describes a decision-making process in which only two options are available, referred to as A and B. Each option i is characterised by an objective quality v_i. The collective decision problem consists in identifying and selecting the best option, or any of the equal-best options. The model treats a population of agents that can be either uncommitted (sub-population U with fraction Ψ_U of the total population) or committed to one of the two options (sub-populations A and B, respectively with fraction Ψ_A and Ψ_B). Populations dynamics can be easily described by a system of two coupled ordinary differential equations, plus a mass conservation term:

$$\begin{cases} \dot{\Psi}_A = \gamma_A \Psi_U - \alpha_A \Psi_A + \rho_A \Psi_A \Psi_U - \sigma_B \Psi_A \Psi_B \\ \dot{\Psi}_B = \gamma_B \Psi_U - \alpha_B \Psi_B + \rho_B \Psi_B \Psi_U - \sigma_A \Psi_A \Psi_B \\ \Psi_U + \Psi_A + \Psi_B = 1 \end{cases}, \tag{1}$$

The variation of the population fraction $\Psi_i, i \in \{A, B\}$ results from four concurrent processes, which correspond to the four terms of each differential equation in (1):

(i) Ψ_i increases as uncommitted individuals spontaneously *discover* and become committed to the option i at the rate γ_i;

(ii) Ψ_i decreases as individuals committed to option i spontaneously *abandon* it and get uncommitted at the rate α_i;

(iii) Ψ_i increases as individuals from population i actively *recruit* uncommitted ones at the rate ρ_i;

(iv) Ψ_i decreases as individuals from population i are *inhibited* by individuals of population $j \neq i$ at the rate σ_j.

All transition rates—$\gamma_i, \alpha_i, \rho_i, \sigma_i$—are greater than zero. It is worth noting that this model does not require any explicit comparison of the option qualities. The quality value v_i of the two options is instead encoded in the transition rates (i.e., implementing a value-sensitive decision making, see Pais et al., 2013): different-quality options correspond to biased transition rates,

while same-quality options correspond to unbiased transition rates. Overall, the collective decision is based purely on the system dynamics resulting from individual-to-individual interactions.

4.1.2 *Implementation guidelines*

WORKING REGIME The choice of the parameters of the macroscopic model determines the working regime for the decision-making process. Understanding the macroscopic dynamics leads to a principled choice of the desired parameterisation, which ultimately translates in prescriptions for the implementation.

When the two options have different quality (e.g., $v_A > v_B$), a biased population distribution is obtained thanks to similarly biased commitment rates. Everything else being equal, a population distribution biased for the better option can be obtained thanks to a higher discovery rate (e.g., $\gamma_A > \gamma_B$, see Figure 2(A)) or similarly through recruitment (e.g. $\rho_A > \rho_B$, see Figure 2(B)). Abandonment and cross-inhibition instead reduce the size of a population committed to a given alternative. Abandonment should be small enough to avoid that a large fraction of the population remains uncommitted, possibly biased toward the lower quality (i.e., $\alpha_A < \alpha_B$). Cross-inhibition instead, being proportional to the size of the inhibiting population, contributes to the creation of an unbalanced distribution of individuals between committed populations even for unbiased inhibition rates (i.e., $\sigma_A = \sigma_B$, see Figure 2(C)). This is true also for same-quality alternatives (i.e., $v_A = v_B$). In this symmetric case, discovery, abandonment and recruitment are equal and are therefore not sufficient to break the symmetry. However, a sufficient level of cross-inhibition makes the equilibrium point unstable, therefore leading to a symmetry breaking, as shown in Figure 2(D). Through linear stability analysis, it is possible to identify the cross-inhibition level for which the system breaks the deadlock and converges to the choice of one option (see Seeley et al., 2012). The working region is $\{\rho > \alpha, \sigma > \sigma^*\}$, with critical value:

$$\sigma^* = \frac{4\alpha\gamma\rho}{(\rho - \alpha)^2}. \tag{2}$$

As a general guideline for the choice of the transition rates determining the working regime of the macroscopic model, it is advisable to have a parameterisation linked to the option quality, so that commitment is proportional to quality, abandonment inversely proportional, while cross-inhibition can be proportional to the option quality or independent of it, but in any case it must be sufficiently large to ensure convergence (for a detailed analysis, see Pais et al., 2013):

$$\gamma_i, \rho_i \propto v_i, \qquad \alpha_i \propto 1/v_i, \qquad \sigma_i > \sigma^*, \qquad i \in \{A, B\} \tag{3}$$

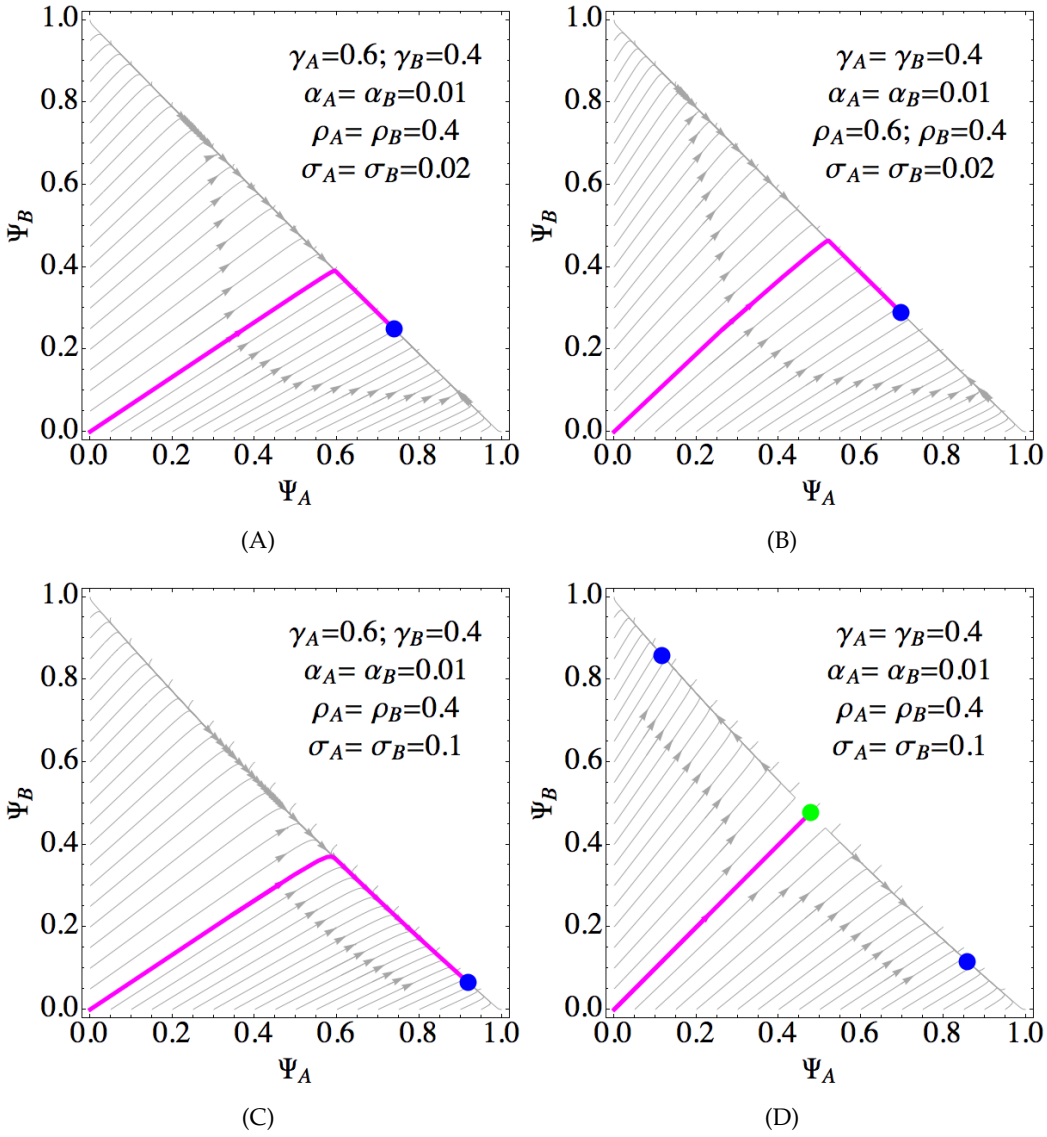

Figure 2: Effects of different parameterisations on the phase portrait of the macroscopic system of Equation (1). System trajectories in the phase plane $0 \leq \Psi_A + \Psi_B \leq 1$ are displayed as grey arrows, stable equilibria are displayed as dark blue dots and unstable saddle equilibria as light green dots. The bold magenta curve, in overlay, is the trajectory starting from the fully uncommitted state, $\Psi_A = \Psi_B = 0$. See text for a detailed discussion on the effects of the different parameterisations on the macroscopic dynamics.

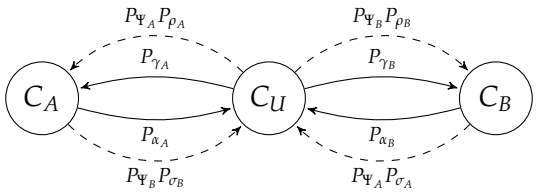

Figure 3: A compact representation as a probabilistic finite state machine of the individual agent behaviour. Solid arrows are spontaneous individual transitions, while dashed arrows represent interactive transitions that are triggered when an agent encounters an agent of another population (which happens with probability P_{Ψ_i} for population i).

INDIVIDUAL BEHAVIOUR Once a suitable working regime is identified, the four concurrent processes resulting in the macroscopic dynamics of (1) need to be implemented as a multiagent system. As a general guideline, the agent behaviour should be implemented as the probabilistic finite state machine shown in Figure 3. Here, the agent can be in three different *commitment* states that indicate whether the agent is uncommitted (C_U) or committed to either option A or B (C_A or C_B). Two types of transition need to be implemented: spontaneous transitions and interactive transitions. Spontaneous transitions correspond to discovery and abandonment, and pertain to the individual agent behaviour. Interactive transitions depend instead on the result of the interaction among agents, and are regulated by the probability of encountering agents of population i, which we refer to as P_{Ψ_i}: the larger the proportion of agents committed to population i, the larger the probability of encountering one of its members. The transition probabilities between different states determine the outcome of the decision-making process, and should be influenced by the option quality v_i to restrict the system dynamics within the working regime discussed above. Therefore, P_{γ_i} and P_{ρ_i}— respectively the discovery and recruitment probability for option i— should be biased toward the option of higher quality; the abandonment probability P_{α_i} should be small and possibly inversely proportional to the option quality v_i; finally, cross-inhibition should be governed by a high-enough probability P_{σ_i}. Prior to formalisation of the link between microscopic transition probabilities and macroscopic transition rates, the above high-level guidelines need to be verified for their sufficiency in providing a quantitative micro-macro link. In Section 4.2, we show how these guidelines translate to an actual implementation. Additionally, we discuss the inclusion of spatial factors influencing the interactions between agents, and we indicate possible extensions of the guidelines to deal with dynamic interaction topologies.

REQUIRED PROPERTIES To obtain a quantitative micro-macro link, two properties are fundamental: state transitions must be memoryless and the system must be well-mixed.

Memorylessness is required by the macroscopic model, which can be derived as a mean-field approximation of a population-level continuous-time Markov process (see Seeley et al., 2012, for details). Memorylessness implies that, at any time, the probability that any agent undergoes a state transitions depends only on the current system state, and is independent of the previous state of the system (thus the macroscopic process fulfils the Markov property). This can be achieved by ensuring a memoryless behaviour at the microscopic level, that is, ensuring that state transitions are governed by fixed probabilities per time unit, which results in an exponential distribution of the time intervals between entrance and exit from a state. Note that spatiality may interfere with memorylessness, and particular care must be given to the microscopic implementation, as exemplified in the case study presented in this chapter.

The well-mixed property ensures that the probability of interaction between any two agents is constant. This can be achieved by a fully connected interaction topology, or by a uniformly random interaction topology. If different populations were segregated, the well-mixed property would not hold anymore, therefore resulting in altered system dynamics (e.g., creation of islands of agents with different opinions). To overcome this, interaction among agents should take place only when an unbiased sample of all populations is available.

4.2 CASE STUDY: SHORTEST PATH DISCOVERY/SELECTION

The high-level prescriptions discussed above need to be reified through experimentation in challenging case studies. Among the several factors that may hinder a macroscopic micro-macro link in a swarm robotics context, spatiality and physical interferences are probably the most important ones. Spatial features constrain the ability of interaction among different populations and may easily lead to departures from the ideal well-mixed condition, while the physical embodiment of specific robotic platforms constrains both motion and robot-robot interactions. In some cases, embodiment may also impose a limit on the maximum number of robots that can operate in a given scenario, as beyond this limit physical interferences might impede the robots correct functioning.

A well-conceived case study should exemplify the challenges introduced by spatiality and embodiment and propose strategies to address them. We have chosen a decision-making scenario that strongly depends on spatial factors and in which physical interferences may have strong effects since the robots share the same space: the discovery and selection of the shortest path between two areas in a foraging context (see Section 4.2.1). We first implement an abstract multiagent simulation (presented in Section 4.2.2) to isolate the challenges introduced by the spatial distribution of target areas

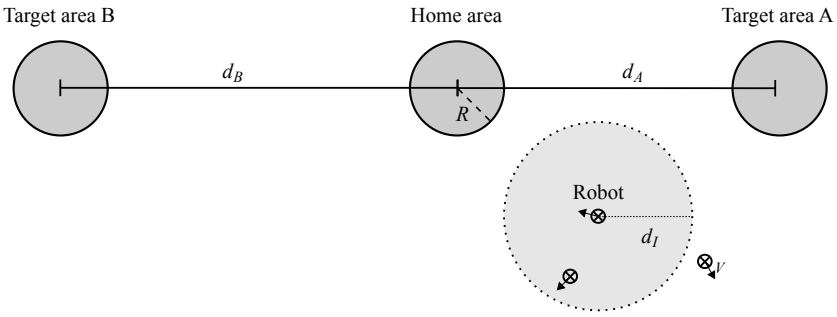

Figure 4: Graphical representation of the environment in the the shortest path discovery/selection case study. Target areas A and B are located at distance respectively d_A and d_B from the home area (in this figure, $d_A = 2\,\text{m}$ and $d_B = 2.5\,\text{m}$). All the three areas have radius $R = 0.3\,\text{m}$. Robots move at constant speed $v = 0.1\,\text{m/s}$ and can communicate with neighbours within a range $d_I = 0.6\,\text{m}$.

while ignoring the target robot embodiment and the resulting physical interactions. Here, agents are modelled as dimensionless particles that do not interfere with each other, so that we can analyse the performance of large swarms, study how the system dynamics vary as a function of the group size, and propose specific solutions to tackle spatiality effects. The robotic simulation instead addresses both spatiality and embodiment, and requires specific strategies to limit the physical interferences which are discussed in Section 4.2.3. Here, we have chosen a robotic platform—the e-puck robot (Mondada et al., 2009)—that stresses the challenges given by physical interferences due to their relatively small perceptual and interaction range.

Overall, the goal of this study consists in obtaining a good quantitative match between the non-spatial macroscopic model of Section 4.1 and the multiagent and robotic implementation, as discussed in Section 4.3.

4.2.1 *Problem definition*

Foraging is a classic problem in swarm robotics (see Brambilla et al., 2013), and often solutions take inspiration from the food-gathering behaviour observed in social insects. Broadly speaking, agents involved in a foraging task are required to carry out search and retrieval activities. They explore the environment to locate *target areas* that contain the objects to be retrieved (e.g., a food patch). Then, they exploit the chosen path to retrieve objects to a given *home area* (e.g., the nest). Foraging encompasses a wide range of activities, such as exploration, navigation between target areas, retrieval and clustering of objects, and collective decision making. All these activities are key components in several real world applications envisioned as potential swarm robotics scenarios — e.g., search and rescue in disaster zones, collective construction in hazardous environments, nuclear disaster cleaning, etc.

Here, we focus on decision making and we neglect other aspects (e.g., object retrieval) that are out of the scope of the present work.

In this study, the goal of the swarm is to identify and exploit the shortest path between a home area and any of two target areas, referred to as target area A and target area B (see Figure 4). The environment is a 2D infinite plane without obstacles, and all areas have circular shape with radius $R = 0.3\,\text{m}$. The targets are located at a distance from home that varies in the range $d_i \in [1.5, 3.5]\,\text{m}$, with $i \in \{A, B\}$. The swarm is composed of agents/robots with local sensing and communication. An agent perceives home and targets only when within the corresponding areas, moves at constant speed of $v = 0.1\,\text{m/s}$ and communicates only with neighbours at a maximum distance $d_I = 0.6\,\text{m}$ (indicated by the dotted circle around the robot in Figure 4). Agents have perfect knowledge of the home location, while target areas need to be located through exploration. The collective decision is taken when the number of agents that have chosen a single target area reaches a quorum Q.

4.2.2 *Implementation of abstract multiagent simulations*

The swarm is composed of dimensionless agents, so that no collisions or physical interferences are possible. Error-free odometry is exploited to track the position of known areas. Agents can be either committed to a specific target, or uncommitted. The committed agents—state C_A and C_B in Figure 3—keep moving back and forth between home and target thanks to odometry. The uncommitted agents—state C_U—explore the environment to discover potential target areas. Agents can communicate with neighbours and share information about their commitment state and the location of discovered target areas.

INTERACTIVE AND LATENT AGENTS Following the guidelines presented in Section 4.1, we have implemented a microscopic behaviour paying particular attention to the effects of spatiality. Indeed, given the locality of communication and the distance between target areas, agents committed to different targets and uncommitted agents cannot always interact with each other. In particular, different populations of agents are spatially segregated, and come into contact only when in the home area. To ensure the well-mixed property, we have therefore decided to let the agents interact with each others only when they are within the home area. When agents leave home, interactions are disallowed. More formally, agents can be in one of two *activity* states: an interactive state I (i.e., when inside the home area) and a latent state L (i.e., when outside the home area). Switching between activity states follows the dynamics prescribed by the implemented behaviour, and can be modelled by the simple PFSM shown in Figure 5 left. We introduce a

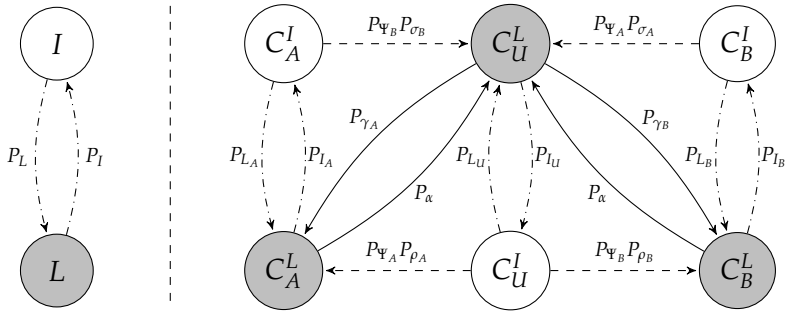

Figure 5: Microscopic description for the shortest path selection/discovery problem. Left: PFSM describing the switch between interactive and latent states. Right: PFSM describing the complete agent behaviour. White circles represent interactive states, while grey circles represent latent states in which no interaction is possible. Solid arrows correspond to spontaneous transitions, while dashed arrows represent interactive transitions, and they both correspond to transitions described in the PFSM of Figure 3. Dash-dotted arrows instead identify the transition between activity states, i.e., the interactive and latent states. See text for further details.

constant probability P_L for an agent to become latent. Conversely, agents become interactive in different ways depending on their commitment state, as will be detailed below. We refer to P_I as the average probability to become interactive. We can describe the change in the activity state through the PFSM model in the left part of Figure 5, which predicts that a fraction $P_I/(P_I + P_L)$ of agents can be found on average in the interactive state. We exploit this prediction to tune the value of P_L: to ensure that on average 10% of the agents are interactive within the home area, we set $P_L = 9P_I$.

Given the three possible commitment states prescribed by the design pattern (Figure 3) and the two activity states, the multiagent implementation can be described by a PFSM with six states as shown in the right part of Figure 5. Therefore, the agent committed to target A is interactive when in state C_A^I and latent when in state C_A^L (respectively, C_B^I and C_B^L for the agent committed to target B). Otherwise, the uncommitted agent is interactive in state C_U^I and latent in state C_U^L. The prescribed PFSM of Figure 3 can be easily obtained aggregating states by commitment.

INTERACTION PATTERNS When agents are in the interactive state, they can exchange short communication messages with neighbours, and on the basis of this communication they can change their commitment state. To obtain the well-mixed property, the distribution of interactive agents in the different commitment states must provide an unbiased representation of the entire population, including latent agents. More precisely, given the fraction Ψ^I (Ψ^L) of agents in the interactive state I (latent state L), we require that:

$$\frac{\Psi_i^I}{\Psi^I} \approx \frac{\Psi_i^L}{\Psi^L} \approx \Psi_i, \qquad i \in \{A, B, U\}, \tag{4}$$

where Ψ_i^I and Ψ_i^L represent the fractions of agents that are found in state C_i^I and C_i^L within the entire population. In fact, if changes in the commitment state within the interactive sub-population (fraction Ψ^I) are much faster than changes in the activity state (i.e., agents switching between states I and L), the distribution of commitment states among interactive agents would misrepresent the global population distribution, and therefore the microscopic and macroscopic dynamics would diverge. We have therefore decided to bind any change in the commitment state resulting from agent-agent interactions to the dynamics of activity change, which ensures the requirements given in (4). In other words, interactions are allowed only upon transitions from the interactive to the latent state: whenever an agent decides to become latent (following the constant probability per time unit P_L), it engages in an interaction with another agent in state I and updates its commitment state accordingly. In this way, changes in commitment state happen at the same rate as changes in the activity state.

Upon interaction, the probability of selecting a partner belonging to each population is proportional to the corresponding fraction of interactive agents. We refer to these probabilities as P_{Ψ_A}, P_{Ψ_B} and P_{Ψ_U}. Thanks to the attentive design of the interactive/latent dynamics discussed above, such probabilities closely represent the global fractions Ψ_A, Ψ_B and Ψ_U. Table 1 indicates all possible transitions from interactive to latent states and also links the transition probabilities to the control parameters P_ρ and P_σ that are introduced in the following. Here, it is worth noting that the sum of all outgoing transitions from any interactive state equals to P_L, which implements the link between changes in activity and commitment state discussed above. The actual change in the commitment state depends on the randomly selected partner, as detailed in the following.

MOTION PATTERNS Agents always move at constant speed $v = 0.1\,\mathrm{m/s}$, and their motion direction is determined by a motion vector $\mathbf{m}$ whose values depends on the agent state. Uncommitted agents explore the environment in search of target areas through a correlated random walk (Bartumeus et al., 2005; Codling et al., 2008). When in the latent state C_U^L, uncommitted agents compute their motion vector $\mathbf{m}$ as the sum of an inertia and a random vector as follows:

$$\mathbf{m} = r_i \angle \theta_i + 1 \angle \theta_r \tag{5}$$

where the notation $r_i \angle \theta_i$ indicates a vector in polar coordinates, with length r_i and angle θ_i. The angle θ_i is the agent's heading direction, θ_r is an angle uniformly drawn in the range $[-\pi, \pi]$, and $r_i = 2$ is the relative strength of the inertia vector determining the random walk. Uncommitted agents switch between latent and interactive states with fixed probabilities per time unit. While searching for target areas (state C_U^L), uncommitted agents stop exploration to return home with constant probability P_{I_U}. The actual transition

Table 1: Correspondence between the transition probabilities of the PFSM of Figure 5 and the control parameters of the implemented behaviour, for each interactive or latent state. This correspondence is shown in the central column of the table: the arrows indicate how each probability of the PFSM (left of the arrow) is implemented in the agent behaviour (right of the arrow). Considering that $P_{\Psi_A} + P_{\Psi_B} + P_{\Psi_U} = 1$, the sum of all outgoing transitions from any interactive state equals to P_L.

from	**transition probability**	**to**
C^I_U	$P_{\Psi_A} P_{\rho_A} \rightarrow P_L P_{\Psi_A} P_\rho$	C^L_A
	$P_{L_U} \rightarrow P_L(P_{\Psi_U} + (1 - P_\rho)(P_{\Psi_A} + P_{\Psi_B}))$	C^L_U
	$P_{\Psi_B} P_{\rho_B} \rightarrow P_L P_{\Psi_B} P_\rho$	C^L_B
C^I_A	$P_{\Psi_B} P_{\sigma_B} \rightarrow P_L P_{\Psi_B} P_\sigma$	C^L_U
	$P_{L_A} \rightarrow P_L(P_{\Psi_A} + P_{\Psi_U} + P_{\Psi_B}(1 - P_\sigma))$	C^L_A
C^I_B	$P_{\Psi_A} P_{\sigma_A} \rightarrow P_L P_{\Psi_A} P_\sigma$	C^L_U
	$P_{L_B} \rightarrow P_L(P_{\Psi_B} + P_{\Psi_U} + P_{\Psi_A}(1 - P_\sigma))$	C^L_B
C^L_U	$P_{\gamma_A} \rightarrow P_{\gamma_A}$	C^L_A
	$P_{I_U} \rightarrow P_I$	C^I_U
	$P_{\gamma_B} \rightarrow P_{\gamma_B}$	C^L_B
C^L_A	$P_\alpha \rightarrow P_\alpha$	C^L_U
	$P_{I_A} \rightarrow P_I$	C^I_A
C^L_B	$P_\alpha \rightarrow P_\alpha$	C^L_U
	$P_{I_B} \rightarrow P_I$	C^I_B

to the interactive state C_U^I takes place as soon as the agent enters the home area. When in the interactive state C_U^I, agents remain within the home area and move by correlated random walk (5). An interactive uncommitted agent becomes latent and resumes exploration with probability P_{L_U} (see Table 1).

Committed agents move back and forth between home and target exploiting odometry, which is used to update the motion vector $\mathbf{m}$ toward the stored area locations. Similarly to uncommitted agents, committed interactive agents (state C_A^I or C_B^I) remain within the home area and move with a correlated random walk (5). They become latent with probability P_{L_A} (P_{L_B}, see also Table 1). When latent (state C_A^L or C_B^L), they travel toward the target area and return home after a full round trip. This is modelled in the PFSM of Figure 5 by the transition probability P_{I_A} (P_{I_B}).

As already mentioned, to provide the well-mixed property, the probability to switch between active and interactive states must be comparable across the agents' commitments states: in this way, the agents in state I are an unbiased sample of the populations A, B and U. To this end, we tune the probability P_{I_U} to obtain a match between the average time of return to the home area of uncommitted agents with the average round-trip time of the committed agents. Given the agent speed v and the distance range of target areas, the average round-trip time is $\tau_m = 2\bar{d}/v = 50\,\mathrm{s}$, where $\bar{d} = 2.5\,\mathrm{m}$ is the average target area distance. Therefore, we fix the probability to become interactive per time unit to $P_{I_U} = 1/\tau_m = 0.02\,\mathrm{s}^{-1}$.

Note that in the present case study different distances of the target correspond to different round-trip times: the closer the target area, the shorter the time needed to return home, the higher the frequency with which agents become interactive. This bias toward closer target areas can be exploited for decision purposes, as the rate of becoming interactive positively bias the system dynamics towards the selection of the shortest path. Starting from these considerations, the implementation of the four concurrent processes prescribed by the design pattern—discovery, abandonment, recruitment and cross-inhibition—can be easily obtained.

DISCOVERY Uncommitted agents discover target areas through random exploration while in the latent state C_U^L. Whenever an agent stumbles upon a target area, it stores the area location and becomes committed to it. In the PFSM model of Figure 5, this event corresponds to the transition between C_U^L and C_A^L (C_B^L), which happens with a probability P_{γ_A} (P_{γ_B}). Discovery events result from correlated random walks that start from the home area: the closer the target area, the higher the discovery probability. Therefore, at a macroscopic level, the discovery rate is biased toward shorter paths as prescribed by the design pattern.

ABANDONMENT Committed agents may spontaneously abandon their commitment and revert to an uncommitted state with constant probability P_α. Only agents in the latent state C_A^L (C_B^L) are allowed to abandon commitment, and become uncommitted in order to resume exploration in state C_U^L, as shown in Figure 5. Given that longer paths imply longer travel times, the macroscopic abandonment rate is larger for longer paths, which is in agreement with the prescriptions of the design pattern. After abandonment, agents return home and from there retrieve exploration. In this way, the dynamics of abandonment do not interfere with discovery, ensuring that every discovery event results from a random exploration that originates from the home location. This allows to preserve the memoryless property of the agent behaviour as required by the design pattern.

RECRUITMENT An uncommitted agent (state C_U^I), upon interaction with a committed agent (state C_A^I or C_B^I), gets recruited and thus committed to the other agent's target area with constant probability P_ρ. Given the interaction pattern discussed above, uncommitted agents get recruited only when becoming latent, therefore with a constant probability P_L. Given that the interactions with committed agents are bound to the probability P_{Ψ_A} (P_{Ψ_B}), the overall recruitment probability is given by $P_L P_{\Psi_A} P_\rho$ ($P_L P_{\Psi_B} P_\rho$), as shown in Table 1[1]. Given that the interactive populations are slightly biased by shorter paths, at the macroscopic level recruitment for closer targets is slightly higher, as prescribed by the design pattern. The commitment message sent by a committed agent contains the information relative to the corresponding target area location, which can be used by the recruited agent for navigation. The transferred information is the angle and distance of the target location relative to the recruiter. The receiver combines this information with the recruiter's relative location and orientation, and through triangulation computes the target location.

CROSS-INHIBITION An interactive agent committed to target A (B) in state C_A^I (C_B^I), upon interaction with an agent in state C_B^I (C_A^I), gets cross-inhibited and reverts to an uncommitted state with constant probability P_σ. Recall that interactions take place with probability P_L, and that the probability of interacting with an agent from population A (B) is P_{Ψ_A} (P_{Ψ_B}). It follows that the overall cross-inhibition probability is $P_L P_{\Psi_B} P_\sigma$ ($P_L P_{\Psi_A} P_\sigma$, see also Table 1). Upon interaction, an agent recognises that the partner is committed to a different target area by measuring the distance between the target area location internally stored and the area location communicated by the partner. If the distance is greater than the target area radius R, cross-inhibition takes place.

1 Table 1 also shows that with probability P_{L_U} uncommitted agents do not get recruited and become uncommitted latent agents (C_U^L). In this way, the sum of all outgoing transitions from the interactive uncommitted state C_U^I is equal to P_L.

Table 2: Control parameters and chosen value or value range. In this case study, all parameters are fixed but the recruitment and cross-inhibition probabilities.

Parameter	Description	Value
v	agent speed	$0.1\,\mathrm{m/s}$
d_I	interaction radius	$0.6\,\mathrm{m}$
r_i	relative strength of the inertia vector	2
P_{I_U}	probability of abandoning exploration	$v/2\bar{d} = 0.02\,\mathrm{s}^{-1}$
P_L	probability of becoming latent	$3P_{I_U}$
P_α	probability of abandoning commitment	$0.005\,\mathrm{s}^{-1}$
P_ρ	probability of recruitment	$[0, 1]$
P_σ	probability of cross-inhibition	$[0, 1]$

CONTROL PARAMETERS The implemented behaviour has several control parameters that can be varied to obtain slightly different macroscopic dynamics. Some parameters have been arbitrarily chosen, while others are tuned to preserve the properties prescribed by the design pattern. In this study, we fix all parameters but P_ρ and P_σ, which determine the interaction pattern between agents. Table 2 summarises the parameters we introduced and the values that have been chosen.

4.2.3 *Implementation of physics-based swarm robotics simulations*

To study how the agent embodiment influences the system dynamics, we implemented the decision process on a swarm of simulated robots. The robotic platform that is simulated is the e-puck robot (Mondada et al., 2009), equipped with a range-and-bearing board (Gutiérrez et al., 2009) to allow short-range localised communication and an embedded computer running Linux[2]. Experimentation is conducted exploiting the ARGoS simulator (Pinciroli et al., 2012). The robot behaviour is implemented in accordance with the multiagent behaviour described in Section 4.2.2. However, the robot characteristics and the physical embodiment introduce several constraints that require ad-hoc modifications of the behaviour. We group the introduced modifications for (i) sensors and actuators, (ii) random walk, (iii) obstacle avoidance, and (iv) path exploitation.

SENSORS AND ACTUATORS In the multiagent implementation, agents have an abstract perception of the environment as well as abstract interactions with neighbours. In a robotic system, the same functionalities are implemented employing the available sensors and actuators. More specifically, the robotic implementation needs specific solutions for area localisation and robot-robot communication.

2 http://www.gctronic.com/doc/index.php/Overo_Extension

E-pucks detect and discriminate areas according to the floor grey level, as measured through the infrared ground sensor. We paint the floor with three grey levels: one grey level for the target areas, a second grey level for the home area, and a third one for the empty area. In this way, robots can perceive an area when physically over it through their ground sensor, and differentiate between types of area by looking at the different sensor activations resulting from different grey scales.

Communication is implemented through a combination of range-and-bearing infrared communication (IR) and WiFi, in a way similar to the range-and-bearing implementation of Roberts et al. (2009). Even if robots communicate via WiFi, interactions are kept local by delivering messages only to robots within IR range (we limit this range to $R = 30\,\text{cm}$). Robots constantly broadcast via the IR channel their IP address, so that only upon reception of an IR message a robot may deliver a WiFi message to the corresponding IP address. On the receiver side, WiFi messages are filtered out when they do not have an IR counterpart. In this way, WiFi communication gets enhanced by the localised aspect: besides receiving the message content, a robot may localise the sender position using the IR signal strength and angle. In addition, by keeping interactions local, the risk of a communication channel overload is limited even when operating with large groups. However, communication with high robot densities or cluttered environments is limited to the subset of the closest neighbours because IR communication works only in line-of-sight. This may have a bearing on the macroscopic dynamics, as discussed below. Range-and-bearing communication is also used to establish a common frame of reference among two communicating agents, given that robots cannot recognise the relative heading of neighbours. We adopt here the solution described by Gutiérrez et al. (2010), based on sharing the relative bearing among the interacting robots.

In this study we consider noiseless sensors and actuators. In a real robot implementation the odometry error of the physical robot could be compensated for exploiting robot-robot communication as done in (Gutiérrez et al., 2010).

RANDOM WALK Differently from the multiagent system, e-pucks employ a differential-drive motion system that does not allow immediate changes of direction. Therefore, during random walk, if the movement vector $\mathbf{m}$ varies at each control step as described in eq. (5), the robot might not be able to change its position due to the time required to rotate on place for heading towards $\mathbf{m}$. To work out this issue while keeping the implementation as close as possible to the multiagent system, a robot computes $\mathbf{m}$ as in (5) and keeps unaltered the desired motion vector $\mathbf{m}$ for $w = 5$ control steps before the next update. This gives the robot a sufficient time to rotate and move

along the desired direction, while keeping a frequent update of the motion vector $\mathbf{m}$.

OBSTACLE AVOIDANCE Due to their physical embodiment, robots must avoid collisions with each other and consequently alter their ideal motion trajectories. To avoid collisions, a robot computes its motion vector $\mathbf{m}$—either from random walk (5) or from directed movements toward target areas—and sums to this an obstacle avoidance vector $\mathbf{o}$. The latter is computed by taking into account neighbouring robots' locations that are perceived in a range of 20 cm over the IR communication channel, and calculating a sum vector in the opposite direction. Finally, the robot motion vector $\mathbf{m_o}$ is computed as:

$$\mathbf{m_o} = \mathbf{m} + 2\mathbf{o}, \tag{6}$$

where a larger weight is given to the obstacle avoidance component to ensure collision-free motion.

PATH EXPLOITATION Obstacle avoidance alone is not sufficient to allow a smooth navigation back and forth between home and target areas. Indeed, groups of robots moving towards opposite locations interfere with each other. To reduce such interferences, we designed the robot trajectories in a round-trip to create a double-line motion, letting robots keep the right with respect to the robots travelling in the opposite direction. To achieve this organised motion, committed robots going to a location rotate clockwise their motion vector $\mathbf{m}$ by an angle $\theta \in [0^\circ, 30^\circ]$, with θ linearly decreasing as a function of the distance to the target location.

EXPECTED EFFECTS OF EMBODIMENT The physics-based implementation of the robotic simulations allows us to investigate the effects of physical embodiment on the collective dynamics. The solutions described above deal with part of the constraints and interferences caused by the robot embodiment, but do not completely solve them. The most important consideration is that the robotic system cannot work with high robot densities, as motion between and within target areas would be strongly altered or completely disrupted. Therefore, we limit our study to groups of 50 robots. With this group size, we expect only a mild divergence from the ideal motion patterns. Additional constraints are given by the line-of-sight communication, which may prevent well-mixed interactions between the interactive robots. However, given that the average number of interacting robots is limited to 10% of the group size according to the multiagent implementation, minor departures from a well-mixed condition are expected. A further analysis of the effects of embodiment in the robotics implementation is presented in Section 4.3.2.

4.3 RESULTS

We evaluate the correctness of the microscopic implementation in both the multiagent and the swarm robotics simulation. To this end, we compare the dynamics of the macroscopic model with the results of the microscopic simulations. In this study, the parameterisation of the macroscopic model is estimated directly from the simulations and we investigate the decision-making dynamics for varying probabilities of recruitment and of cross-inhibition (P_ρ and P_σ, respectively), and for a set of different decision problems for varying option quality v_i (i.e., varying the target areas distance d_i, $i \in \{A, B\}$).

4.3.1 *Abstract multiagent simulations*

The effects of spatiality can be appreciated by looking at the decision-making dynamics of multiagent simulations by varying the position of target areas. We show results for a set of decision problems in which the distance of target B is fixed ($d_B = 2.5\,\text{m}$), while the distance of target A systematically varies ($d_A \in \{1.5, 2, 2.5, 3, 3.5\}\,\text{m}$). When not stated otherwise, simulations are run for a total population of $N = 500$ agents. We first discuss the relation between estimated transition rates and the target area distances. Then, we evaluate the micro-macro link comparing the final population distributions in multiagent simulations with the macroscopic model attractors, and we verify the presence of the same type of phase transitions at both levels. Finally, we extend the analysis to varying group size N to identify finite-size effects on both microscopic and macroscopic dynamics.

ESTIMATION OF THE MACROSCOPIC TRANSITION RATES To verify the existence of the micro-macro link, it is necessary to relate the microscopic implementation to the macroscopic model. This is possible by estimating the transition rates of the macroscopic model directly from the multiagent simulation. We do so through survival analysis, which provides powerful non-parametric methods (Nelson, 1969) to estimate how the probability of events changes over time directly from the experimental data (for a detailed explanation of the methods, see Appendix A). We designed a set of experiments to estimate all the macroscopic transition rates of the ODE system in (1)—$\gamma_i, \alpha_i, \rho_i, \sigma_i, i \in \{A, B\}$—for each possible distance of the target areas. For the interactive transitions, we also varied the control parameters P_ρ and P_σ in the set $\{0.1, 0.4.0.7, 1\}$.

Figure 6 shows the result of the parameter estimation. In agreement with the implementation choices, the estimated discovery rate $\hat{\gamma}_i$ decreases with the distance d_i of the target area i, while the estimated abandonment rate $\hat{\alpha}_i$ follows the opposite trend (see Figure 6(A)). Note that discovery follows an exponential decay with increasing distance at an estimated decay

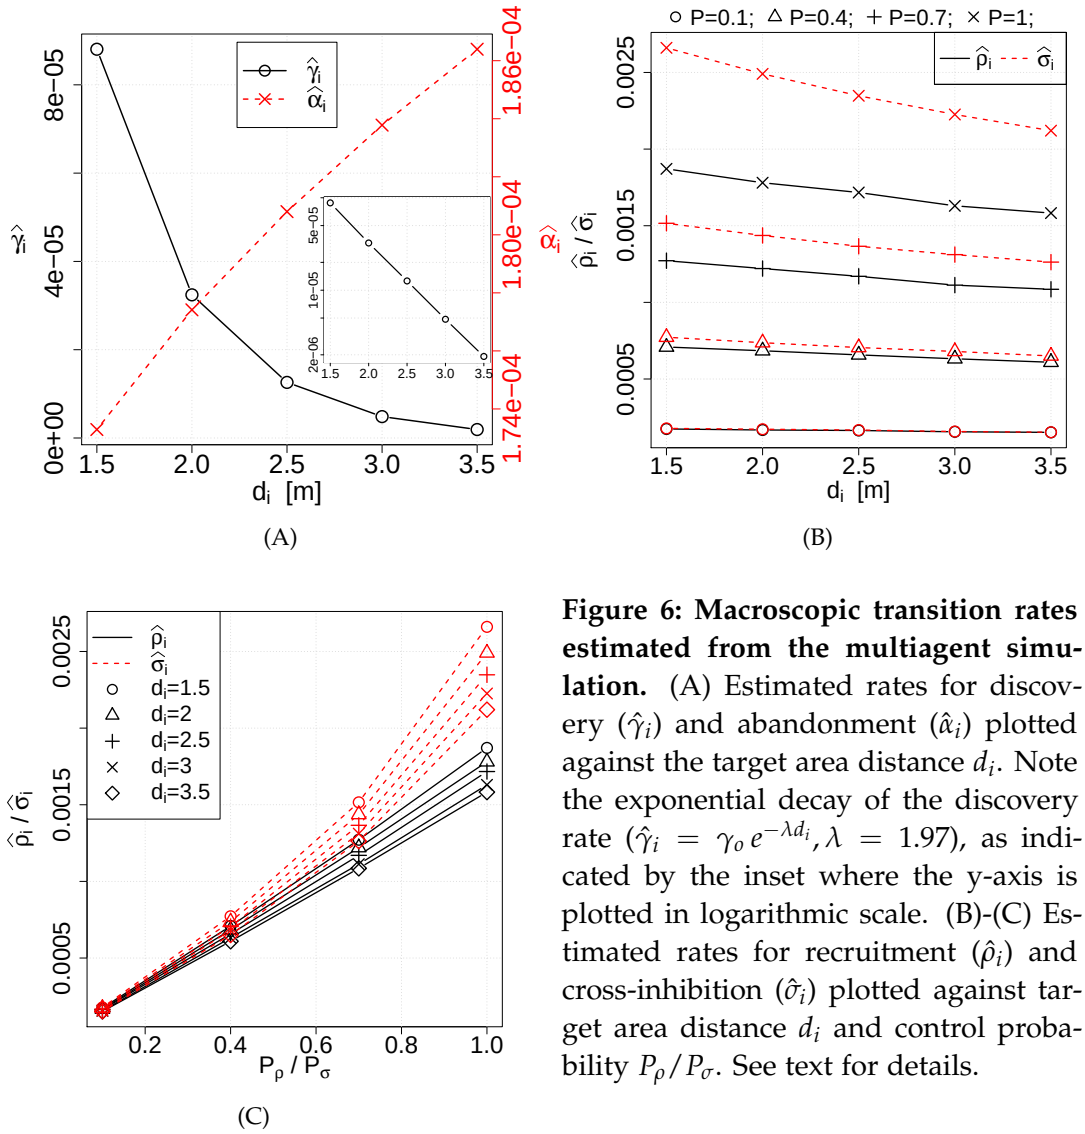

Figure 6: Macroscopic transition rates estimated from the multiagent simulation. (A) Estimated rates for discovery ($\hat{\gamma}_i$) and abandonment ($\hat{\alpha}_i$) plotted against the target area distance d_i. Note the exponential decay of the discovery rate ($\hat{\gamma}_i = \gamma_o\, e^{-\lambda d_i}, \lambda = 1.97$), as indicated by the inset where the y-axis is plotted in logarithmic scale. (B)-(C) Estimated rates for recruitment ($\hat{\rho}_i$) and cross-inhibition ($\hat{\sigma}_i$) plotted against target area distance d_i and control probability P_ρ / P_σ. See text for details.

rate $\lambda = 1.97$, as shown in the inset of Figure 6(A). This is a result of the motion pattern implemented for uncommitted agents, which is governed by exponentially distributed exploratory trips. Such an exponential decay implies that farther targets are more difficult to discriminate than closer ones through discovery, and agent-agent interactions are useful to break possible deadlocks and lead to consistent decisions.

The interaction rates of recruitment $\hat{\rho}$ and cross-inhibition $\hat{\sigma}$ have been estimated for different values of the control probabilities P_ρ and P_σ, respectively. Also in this case, the rates are inversely proportional to the distance of the target area to which the interacting population is committed (see Figure 6(B)), in agreement with the implementation choices. Indeed, such a bias in the transition rates derives from the shorter round-trips performed by agents committed to closer targets. As a result, biased rates favour con-

vergence on shorter paths, confirming the correctness of the multiagent implementation. Note that the effect of distance on the transition rates is linear, and is important for larger values of the control probability but nearly negligible for smaller values, as shown in Figure 6(B). The estimated transition rates vary approximately linearly also with respect to the control probability of the individual agents, P_ρ and P_σ, as shown in Figure 6(C). A slight departure from linearity is visible for cross-inhibition, which is mainly due to errors in the estimation procedure given by the fact that returning home of committed agents is not a purely memoryless process.

FINAL DISTRIBUTION Given the estimated macroscopic parameters for varying distance d_A and varying control parameters P_ρ and P_σ as resulting from the previous analysis, we can evaluate the micro-macro link and compare the dynamics displayed by the two description levels. For each decision problem obtained varying target distances and control probabilities, we perform 500 runs and we compare the population distribution after $5 \cdot 10^3$ s (with an integration time-step of 0.1 s) with the vector field and phase portrait of the ODE system (1) obtained using the estimated macroscopic parameterisation. Figure 7 shows the correspondence between microscopic and macroscopic dynamics for selected decision problems. The complete set of results for every decision problem we investigated is available in Figures S1–S5 in the online supplementary material[3]. As can be observed in Figure 7, we obtained a very good agreement between the macroscopic dynamics and the multiagent simulations, which confirms the existence of a quantitative micro-macro link as a result of the correct implementation of the multiagent behaviour. The agreement is noticeable not only when the microscopic dynamics have stabilised (i.e., population distributions around the macroscopic stable point, see Figures 7(A), (B) and (D)), but also in unstable transitory states which precisely follow the macroscopic vector field, as shown in Figure 7(C). Here, the macroscopic dynamics predict a quick convergence to a one-dimensional manifold, followed by a slower diffusion toward the one or the other attractor (Pais et al., 2013). These dynamics are well reproduced by the multiagent simulations, which show several points scattered around the one-dimensional manifold.

Analysing the system dynamics, it is possible to notice that for low values of P_ρ a decision is not taken and a large majority of the agents remains uncommitted (see Figure 7(A)). This behaviour results from a small positive feedback from recruitment which cannot balance the spontaneous abandonment rate (see also Figure 6). In fact, by increasing the value of P_ρ, the system reliably converges towards the best option (see Figure 7(B)). As shown in the inset, 100% of the runs reach the quorum $Q = 0.75$ for the best option. Cross-inhibition has a different role in the system dynamics: it speeds up

3 http://iridia.ulb.ac.be/supp/IridiaSupp2016-001/index.html

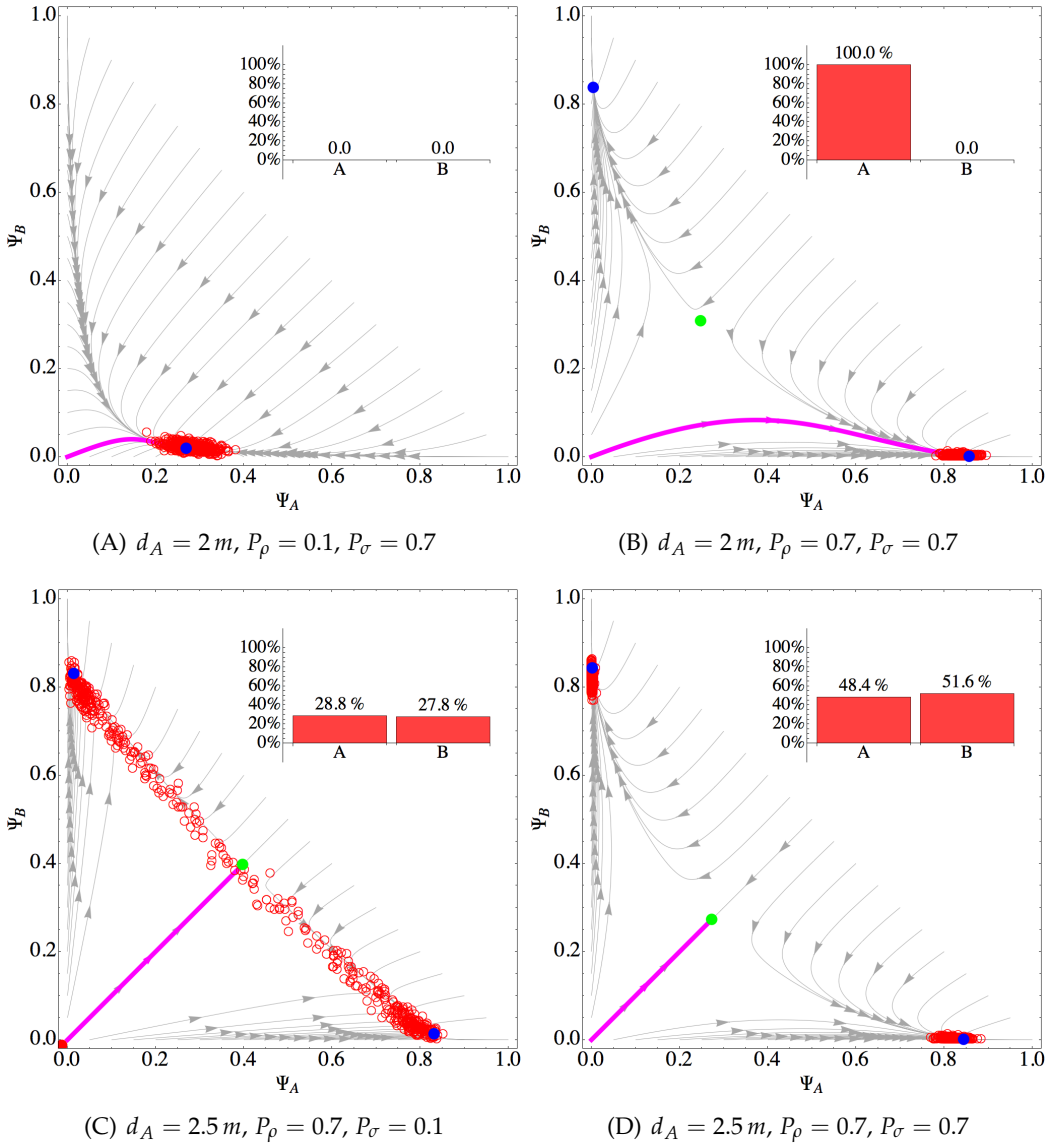

(A) $d_A = 2\,m$, $P_\rho = 0.1$, $P_\sigma = 0.7$

(B) $d_A = 2\,m$, $P_\rho = 0.7$, $P_\sigma = 0.7$

(C) $d_A = 2.5\,m$, $P_\rho = 0.7$, $P_\sigma = 0.1$

(D) $d_A = 2.5\,m$, $P_\rho = 0.7$, $P_\sigma = 0.7$

Figure 7: Comparison of the microscopic and macroscopic dynamics for different decision problems and control probabilities. The macroscopic dynamics are displayed through the phase portrait (grey arrows show the trajectories, filled circles are equilibrium points—dark blue: stable; light green: unstable). The bold magenta arrow represents the trajectory starting from a fully uncommitted population ($\Psi_A = \Psi_B = 0$). The microscopic dynamics are displayed as a scatterplot representing the final distribution of 500 independent runs (red empty circles). The insets show the decision pattern for a quorum $Q = 0.75$, indicating the percentage of runs that resulted in an above-quorum fraction of agents committed to either alternative.

the decision process, and avoids deadlocks at indecision (Pais et al., 2013). In particular, for equally-distant target areas, the dynamics with low values of P_σ are slower and the system is often found to be still at indecision after

$5 \cdot 10^3$ s. As shown by the inset of Figure 7(C), in this time frame only about half of the runs reached the decision quorum for one or the other option. As soon as we increase P_σ, the system always converges to a large majority of agents committed to either one of the two alternatives (see Figure 7(D)).

BIFURCATION To further validate the existence of a precise quantitative micro-macro link, we tested whether the macro and the micro systems undergo the same phase transition at the same predicted value σ^* (see Section 4.1.2). For an unbiased decision problem (i.e., $d_A = d_B$), transition rates for any of the two options are equal, and in such a completely symmetric situation the system risks getting stuck at indecision. As discussed in Section 4.1.2, there exists a value of sigma, σ^*, for which the system breaks the symmetry and converges to either one of the two options. Using the estimated transition rates for $d_A = d_B = 1.5\,\text{m}$ and $P_\rho = 1$, we computed the estimated critical value $\hat{\sigma}^* = 3.96261 \cdot 10^{-5}$. Note that the estimated macroscopic transition rate $\hat{\sigma}$ is linearly related to the individual-level cross-inhibition probability P_σ, as shown in Figure 6(C), especially for small values of P_σ:

$$\hat{\sigma} = K(d_A, d_B) P_\sigma, \tag{7}$$

where $K(d_A, d_B)$ is a constant value that embeds spatial factors related to the distance of both target areas, as well as the constant probability P_L of engaging in an interaction when agents are interactive (see Table 1). Since we have control only on the parameter P_σ, we discount $K(d_A, d_B)$ from $\hat{\sigma}^*$. The pitchfork bifurcation is predicted by the macroscopic model at the value $P_{\sigma^*} = \hat{\sigma}^* / K(d_A, d_B) = 1.4089 \cdot 10^{-2}$ (see Figure 8(A)). We have therefore performed a bifurcation analysis with the multiagent simulation, varying P_σ in the range $[0, 0.1[$ (with a 0.005 step increment). For each condition, we performed 500 runs lasting $2 \cdot 10^4$ s, a sufficient time to ensure convergence. Figure 8(A) illustrates the final distribution of Ψ_A as a density histogram, which results in a very good agreement with the macroscopic bifurcation diagram. Figure 8(B) shows similar results for a different parameterisation: $d_A = d_B = 2.5\,\text{m}$ and $P_\rho = 1$.

FINITE-SIZE EFFECTS A limitation of the macroscopic model (1) is the infinite-size approximation, which turns into the impossibility to precisely analyse the system behaviour for varying group size N. To deal with this, we numerically study the finite-size effects through Monte Carlo simulations of a macroscopic finite-size model exploiting the Gillespie algorithm (Gillespie, 1976), a widely-used method to study the behaviour of continuous-time, well-mixed, memoryless processes. Figure 9 shows that also for small values of N the finite-size dynamics are in agreement with the infinite size ODE system (1). However, for small size N, the macroscopic system is subject to larger random fluctuations which result in more frequent wrong decisions

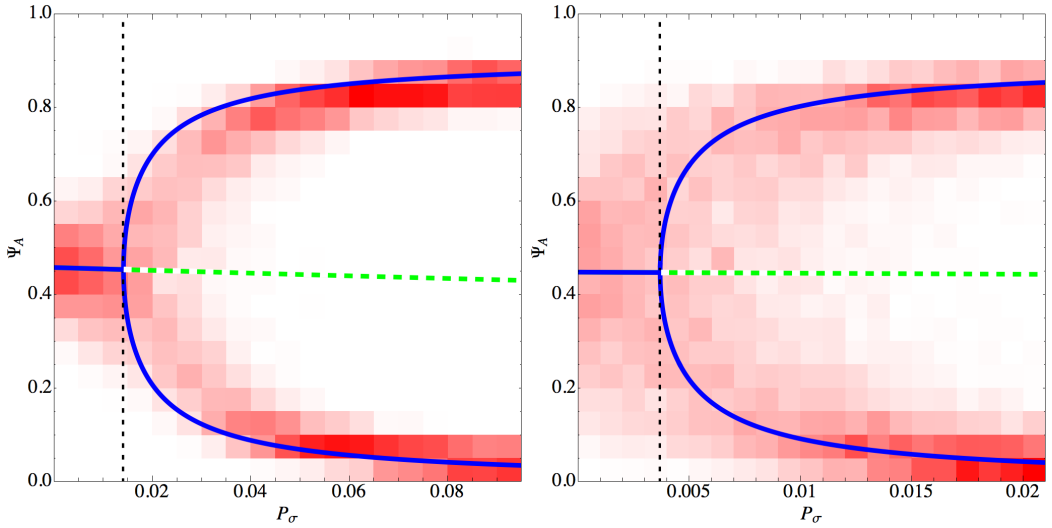

Figure 8: Phase transition predicted by the macroscopic model (lines) and resulting from microscopic multiagent simulations (colour shades) for two unbiased decision problems: (A) $d_A = d_B = 1.5\,\text{m}$ and (B) $d_A = d_B = 2.5\,\text{m}$. We show the proportion of agents committed to option A, Ψ_A, varying the control probability P_σ. The macroscopic bifurcation diagram illustrates stable states as dark (blue) solid lines and unstable state as light (green) dashed lines. The distribution of Ψ_A obtained from multiagent simulations is illustrated through a density histogram in which more frequent values are represented as darker (red) boxes.

especially in decision problems with a small difference between options (see the inset histogram in Figure 9).

To better quantify the accuracy of the micro-macro link including finite-size effects, we introduce the *exit probability*, a measure that indicates the percentage of runs that terminate with a proportion of agents committed for the best (or equally best) option greater than a quorum $Q = 0.75$. Figure 10 shows the estimated exit probability for varying group size for both multiagent and Gillespie simulations. Also in this case, we recognise a good agreement between macroscopic dynamics and multiagent simulation results. Multiagent simulations have a slightly higher exit probability than predicted by the macroscopic model. This is to be accounted to the larger time delay in reporting discoveries of farthest targets, which slightly bias the decision problem towards the closest option. Such effects cannot be grasped by the macroscopic dynamics starting from a fully-uncommitted population, but could be accounted for by an appropriately measured bias in the starting conditions (e.g., starting with a small population fraction committed to the closest target). In future studies, by introducing such a bias in the macroscopic finite-size models, better predictions could be achieved.

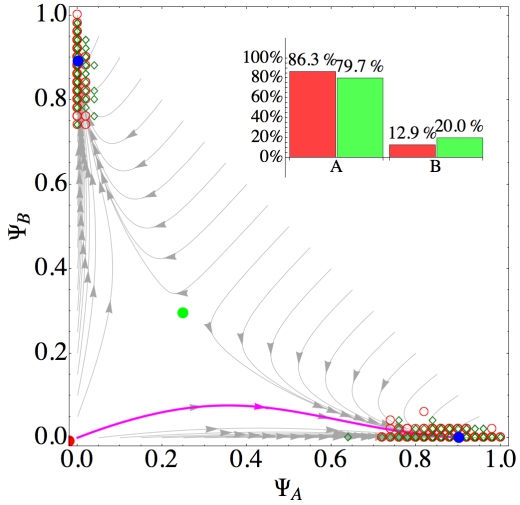

(A) $N = 50$

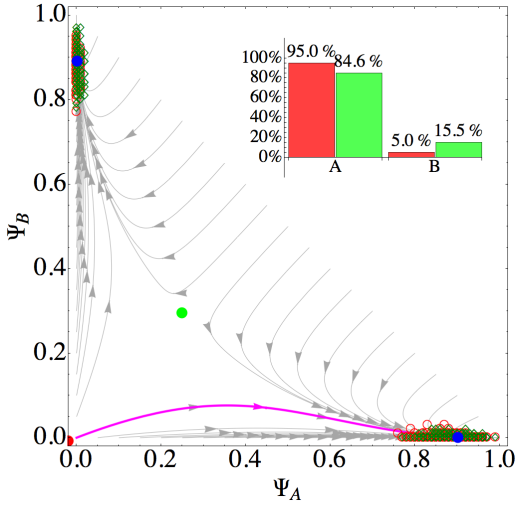

(B) $N = 100$

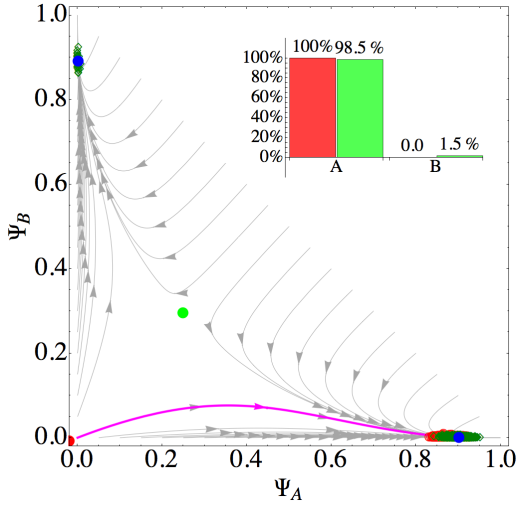

(C) $N = 500$

Figure 9: Comparison of the macroscopic and microscopic dynamics including macroscopic simulations of finite-size effects. The experiment parameterisation is $d_A = 2\,\text{m}$, $d_B = 2.5\,\text{m}$, $P_\rho = P_\sigma = 1$. The populations distribution for 2000 independent runs are displayed as red empty circles for multiagent simulations, and as green diamonds for the Gillespie simulations. The inset shows the percentage of runs in which a decision was made for one or the other option (quorum $Q = 0.75$).

4.3.2 *Physics-based swarm robotics simulations*

The swarm robotics implementation aims to investigate the effect of embodiment on the macroscopic dynamics. In Section 4.2.3 we have discussed several solutions introduced to reduce interferences and deal with physical interactions among robots and between robots and environment. We have therefore run several simulations using the ARGoS framework (Pinciroli et al., 2012) to study the microscopic dynamics. Also in this case, we study several decision problems fixing the distance of target B to $d_B = 2.5\,\text{m}$ and systematically varying the distance of target A, $d_A \in \{1.5, 2, 2.5, 3, 3.5\}\,\text{m}$. Simulations are performed with groups of $N = 50$ robots.

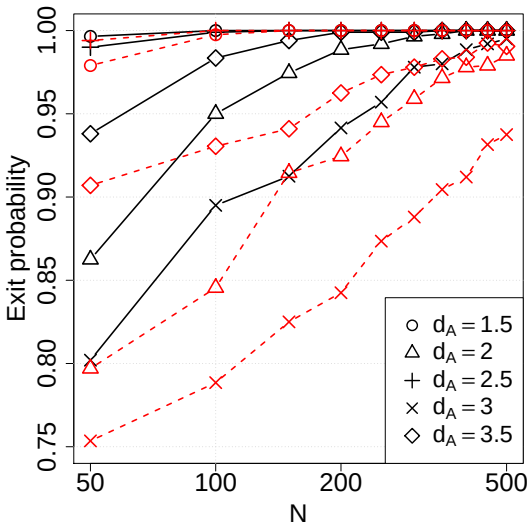

Figure 10: Exit probability as a function of the swarm size N for multiagent (black solid lines) and Gillespie simulations (red dashed line). In these experiments, the distance of target B is fixed, $d_B = 2.5\,\text{m}$, while the distance of target A varies, $d_A \in \{1.5, 2, 2.5, 3, 3.5\}\text{m}$.

ESTIMATION OF THE MACROSCOPIC TRANSITION RATES For each experimental condition, we estimated the macroscopic transition rates through survival analysis, following the same methodology used for multiagent simulations and detailed in Appendix A (see Figure 11). The estimated discovery ($\hat{\gamma}$) and abandonment ($\hat{\alpha}$) rates vary with target distance in a similar way to the multiagent simulations, as shown in Figure 11(A). Abandonment increases with distance while discovery exponentially decays at a decay rate $\lambda = 1.27$ (see also the figure's inset). Abandonment rates are also in a good quantitative agreement with the multiagent simulations, while discovery rates present larger values. This is a result of the difference in the correlated random walk by agents and robots, the former having a smaller correlation in the motion direction due to the ability of instantaneous turning.

The estimated recruitment ($\hat{\rho}_i$) and cross-inhibition ($\hat{\sigma}_i$) rates depend on the distance d_i, as well as on the control probabilities P_ρ and P_σ respectively, as shown by Figures 11(B) and (C). While recruitment correctly decays with distance, similarly to the multiagent implementation, we notice that the cross-inhibition rate slightly increases with distance. This is mainly the result of physical interferences among robots committed to the same option, which are less frequent for longer paths because the density of robots over the path is reduced, therefore leading to higher transition rates. Indeed, notwithstanding the ad-hoc motion pattern for path exploitation discussed in Section 4.2.3, physical interferences cannot be completely neutralised, the stronger their effect the higher the robot density, resulting in a penalty for motion on shorter paths.

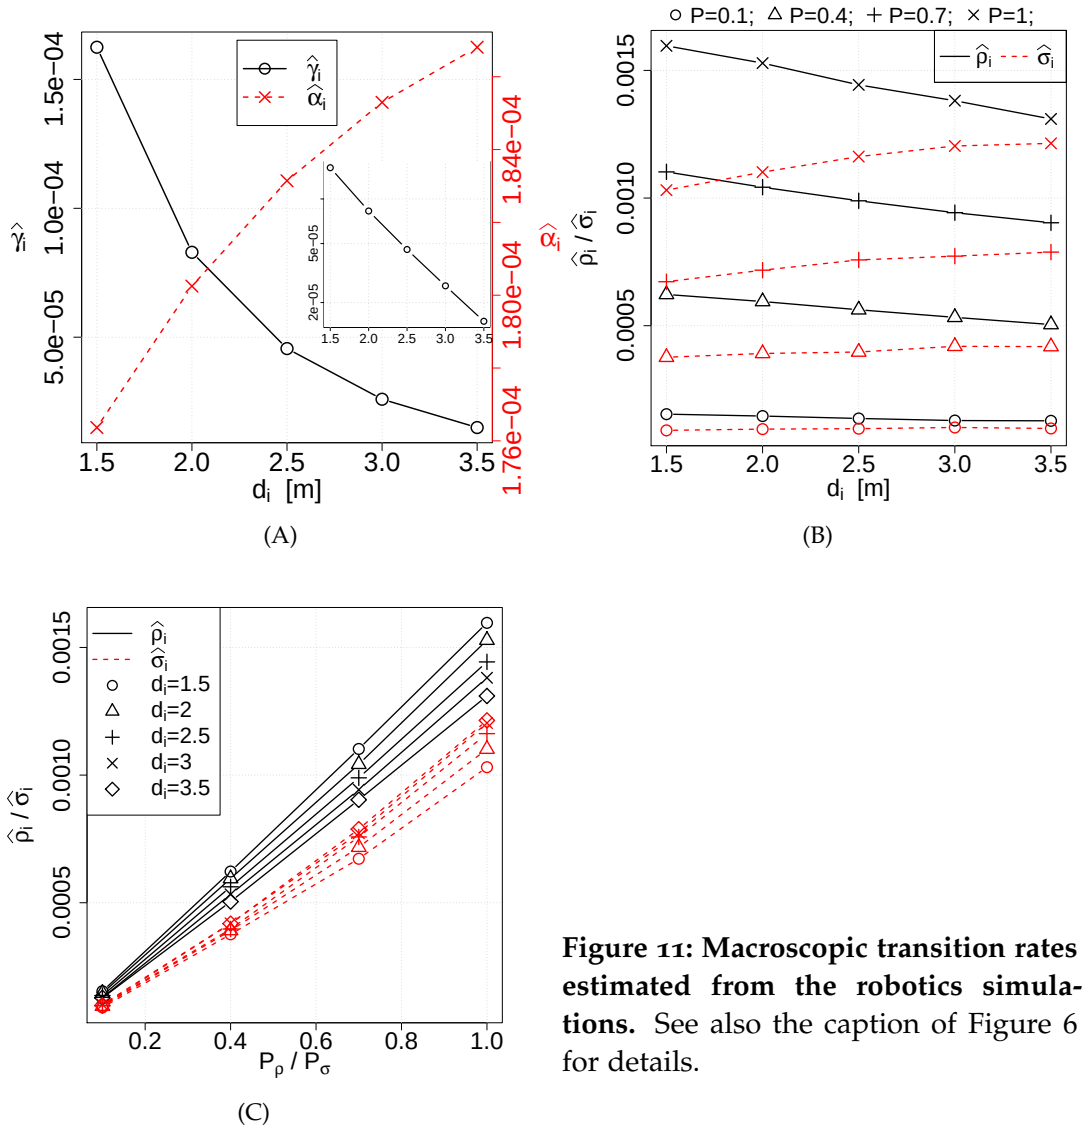

Figure 11: Macroscopic transition rates estimated from the robotics simulations. See also the caption of Figure 6 for details.

FINAL DISTRIBUTION Exploiting the estimated macroscopic transition rates, we can analyse the micro-macro link also for the robotics simulations. We have therefore matched the macroscopic dynamics with the final population distribution from 500 independent runs of the robotics simulations. Figure 12 shows four selected conditions (corresponding to the same decision problems and control parameters used in Figure 7). The complete set of results for every decision problem we investigated is available in Figures S6-S10 in the online supplementary material[4]. Also in this case, the agreement between macroscopic dynamics and microscopic simulations is remarkable. With respect to multiagent simulations, we note a larger scatter of data points, which is due to larger random fluctuations related to

4 http://iridia.ulb.ac.be/supp/IridiaSupp2016-001/index.html

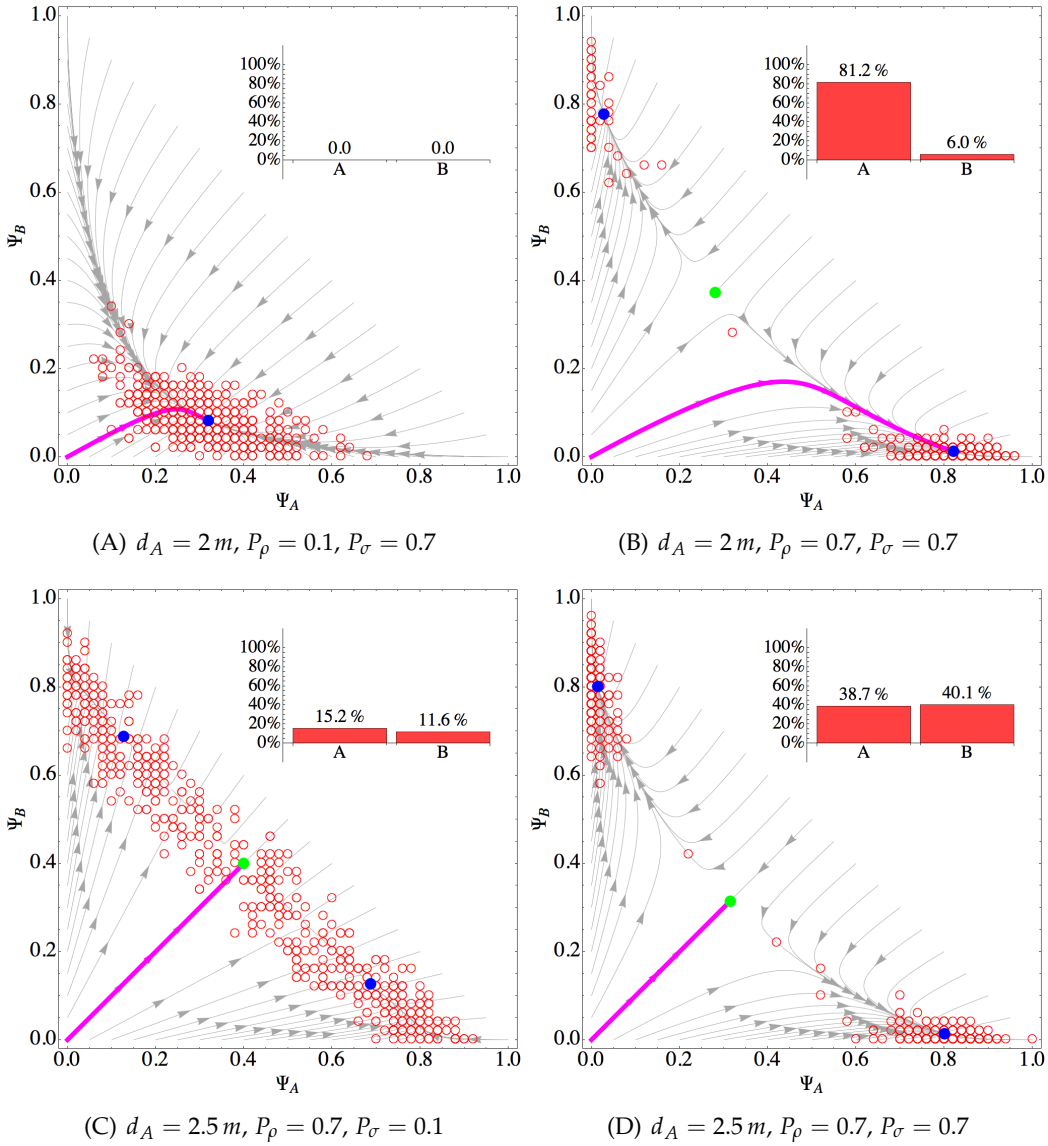

(A) $d_A = 2\,m$, $P_\rho = 0.1$, $P_\sigma = 0.7$

(B) $d_A = 2\,m$, $P_\rho = 0.7$, $P_\sigma = 0.7$

(C) $d_A = 2.5\,m$, $P_\rho = 0.7$, $P_\sigma = 0.1$

(D) $d_A = 2.5\,m$, $P_\rho = 0.7$, $P_\sigma = 0.7$

Figure 12: Comparison of the microscopic and macroscopic dynamics for different decision problems and control probabilities. See the caption of Figure 7 for details.

the physical embodiment of robots, as well as to finite-size effects given the relatively low group size. Apart from this, the microscopic implementation accurately matches the macroscopic dynamics. We also recognise that the decision process is not compromised by a cross-inhibition rate increasing with distance. Indeed, the additional negative feedback is counterbalanced by the positive feedback given by recruitment, which allows to achieve rather accurate decisions.

4.4 DISCUSSION

In this chapter, we have demonstrated how to obtain a quantitative relationship between a desired macroscopic model of decentralised decision making borrowed from studies of nest-site selection in honeybees (Seeley et al., 2012) and a shortest path discovery/selection problem within both multiagent and robotics simulations. This work is particularly relevant in the perspective of formalising a design pattern for decentralised decisions based on the mechanisms observed in honeybees. The high level guidelines presented in Section 4.1 have been verified in an example that is particularly challenging due to the spatial component which is not taken into account in the macroscopic model. The relationship between microscopic parameters and macroscopic transition rates was obtained directly from the experimental data, showing that in many cases an approximately linear relation holds, with few exceptions that can be ascribed to the effects of spatiality and embodiment. Additionally, specific solutions have been proposed to deal with both spatiality and embodiment which will be included in the formalisation of the design pattern for the general case.

For instance, the distinction between interactive and latent states can be generalised to any scenario in which—due to spatiality or other constraints—agent-agent interactions are not always possible, so that a necessity to differentiate between activity states arises. The lesson learned from the present case study is the importance of implementing spontaneous transitions between the activity states with comparable dynamics across different populations of committed and uncommitted agents. This should ensure that the interactive agents are an overall unbiased representation of the entire population. What needs to be further investigated is the amount of bias that can be introduced before disrupting the micro-macro link.

Another important implementation choice consists in linking the possibilities of interaction to the dynamics of activity change. In this respect, the lesson learned is the importance of having similar dynamics between interaction rates and activity changes, in order to maintain unbiased proportions within the interactive population. Also in this case, it would be important to ascertain whether alternative implementation patterns exist in order to speed-up the decision process beyond the constraints imposed by activity dynamics.

Finally, it is worth noticing that spatial factors are well managed also thanks to well designed motion patterns and transitions between states of the PFSM individual-level model. The lesson learned is that a proper design of the individual behaviour should lead to memoryless processes whenever possible. In this way, the microscopic implementation can be easily linked to the macroscopic description. When spatiality provides some form of bias, it is important to design the individual behaviour to neutralize such a bias

(e.g., returning to the home location after an abandonment is mandatory to avoid a bias in the discovery of the abandoned alternative). Neutralising spatiality effects may have costs (e.g., it slows down discovery of potential alternatives), but these are costs that need to be paid to obtain a quantitative micro-macro link.

The problem studied in this chapter and the solution proposed have some specificities that do not allow to generalise toward every decentralised decision-making problem. However, this case study can help in understanding the effects of spatiality on the decision dynamics. In this study, the quality of the alternatives is not directly available to the agents and decisions are bound to biases related to spatial factors (i.e., differential latencies related to different options, see also Montes et al., 2010; Scheidler et al., 2015). As a consequence, not all possible macroscopic dynamics can be obtained, because the macroscopic transition rates can only be modulated by the chosen control parameters but not completely controlled. Different dynamics could be achieved in case each agent can individually estimate (with noise) the quality of the available option to contribute to the collective choice (as is shown in the case study of Section 6.3).

In the next chapter, we formalise the suggestions given in this study into a design-pattern for decentralised decision making in the best-of-n problem. The macroscopic and microscopic models for binary decisions are extended to the general case of n alternatives (i.e., the best-of-n problem). More formal guidelines and methods to deal with spatiality issues—or, more generally, with the existence of latent states—are provided, together with a formal relationship between microscopic control parameters and macroscopic transition rates. As a consequence, it becomes possible to select the relevant parameterisation of the system at the macroscopic description level, and immediately derive the corresponding parameterisation of the individual behaviour that leads to the desired collective outcome. This leads to a complete micro-macro link for best-of-n, decentralised decision making allowing both top-down design and bottom-up verification and analysis.

All models are false, but some are useful.

—George Box

5

A DESIGN PATTERN FOR DECENTRALISED DECISION MAKING

Following the definition presented in Chapter 3, we report here a detailed description of five (out of six) attributes of the design pattern for best-of-n, decentralised decision making. The sixth attribute (the set of case studies) is described in Chapter 6. While the case studies are part of the design pattern, we present them in a separate chapter because of their long text. We believe that organising the design pattern in two separate chapters eases the reading. The formalisation of the design pattern's solution, that is presented in Section 5.5, is based on the results presented in Chapter 4.

5.1 NAME

Collective decisions through cross-inhibition (CDCI).

5.2 PROBLEM

The recurring problem tackled by this design pattern is the best-of-n decision problem, that is, the choice of the best option, or any of the equal-best options, among a finite, possibly unknown number n of different alternatives. Each option $i \in \{1, \dots, n\}$ is characterised by the quality $v_i \in [v_m, v_M]$. The decision-making process therefore requires:

- the identification of (possibly all) the available options;
- the estimation of the quality v_i of each identified option i;
- the selection of the best one, or of any of the equal-best options $i^\star \in \arg\max_{i \in \{1,\dots,n\}} v_i$.

We study decision making for a decentralised system composed of N autonomous agents a_g, $g \in \{1, \dots, N\}$. Each agent is either committed to one of the available options ($C(a_g) = C_i$, $i \in \{1, \cdots, n\}$, where C_i indicates commitment to option i), or uncommitted ($C(a_g) = C_U$, where C_U indicates

uncommitment). At the macroscopic level, a decision is taken as soon as the entire population (or a large fraction Ψ_q) becomes committed to one of the options.

The best-of-n decision problem is a challenging problem especially when the number and quality of the available alternatives is not known apriori. Additional complexity might result from uncertain environmental conditions that determine a noisy estimation of the option quality. In other words, the inaccurate quality estimation requires repeated evaluations to increase the decision accuracy. This is a time-consuming process that naturally leads to a speed-accuracy tradeoff (Chittka et al., 2009).

In a decentralised system, each agent may discover and evaluate only a small subset of the available options. However, the system as a whole has to converge on the best option (or any of the equal-best options). Additionally, agents are assumed not to have global knowledge of the system state (e.g., population size, distribution of agents across populations, number of available options).

5.3 CONTEXT

The CDCI supports the implementation of decentralised decision making for a multiagent system in which each agent is autonomous and features the following minimal set of abilities:

- it can individually recognise available options;
- it can individually estimate the options quality;
- it can communicate with peers using small amounts of information;
- it can recognise peers committed to a different option.

Agents neither need to be able to memorise more than one option at a time, nor to explicitly compare different options. The estimated quality $\hat{v}_i$ of a selected option is used only to modulate the individual behaviour (e.g., by altering the probability of performing a certain action). Given these preconditions, a viable solution to achieve a collective decision is the implementation of a truly decentralised algorithm.

5.4 DESIGN RATIONALE

Models of decentralised decision making have been studied in different domains, from ethology to social dynamics. In this design pattern, we propose a methodology starting from a model of nest-site selection in honeybee swarms (Seeley et al., 2012). In honeybee swarms, after spring reproduction, several thousands bees leave their hive and create a cluster in the neighbourhood lasting a few days. During this time, the oldest bees in the swarm

search for new nest sites and, once they *discover* one, they commit to it. On the one hand, committed bees have a tendency to spontaneously *abandon* their commitment. On the other hand, by interacting with other bees through the waggle-dance, committed bees *recruit* uncommitted nest-mates to the site they have discovered. The waggle-dance duration is proportional to the quality of the advertised nest, and this induces a positive feedback that increases the number of bees committed to the best quality nests. Eventually, a quorum is reached for a single site, which is chosen as the new nest site. Recently, it has been discovered that bees committed to different options *cross-inhibit* each other through stop signals (Seeley et al., 2012). A bee committed to a site that receives several stop signals abandons its commitment and becomes uncommitted. This mechanism allows the swarm to break decision deadlocks in case of equal-best options. In this way, the swarm reduces the decision time, thus exposure to dangers such as predation or adverse weather conditions.

The decision-making process is based on individual actions and peer-to-peer interactions (i.e., discovery, abandonment, recruitment and cross-inhibition), and lets the swarm quickly converge towards the highest quality option without the need of quality comparisons. It also allows to break deadlocks between same quality options, as well as to modulate the decision dynamics on the basis of the quality of the discovered options (Pais et al., 2013). These advantageous characteristics and low requirements in terms of agent capabilities allow designers to apply the design pattern in a large number of different application contexts.

5.5 SOLUTION

The decentralised decision-making process of honeybees is modelled as a continuous-time Markov process (Seeley et al., 2012). Starting from this model, through a mean-field approximation, a deterministic macroscopic model is derived as a system of two coupled ODEs for a binary decision problem. Here, we extend the models to the best-of-n problem and complement the multi-level description by introducing the master equation, and the PFSM that describes the individual agent behaviour.

5.5.1 *Macroscopic description: infinite-size, deterministic, time/state-space continuous*

Let us consider a population of N agents (with $N \to \infty$). At the macroscopic level, we model the population fractions of committed agents $\Psi_i = N_i/N$ (with N_i the number of agents committed to option i) and the fraction of uncommitted agents $\Psi_U = N_U/N$ (with N_U the number of uncommitted agents). Agents change their commitment state through four different pro-

cesses: discovery (γ), abandonment (α), recruitment (ρ) and cross-inhibition (σ).

We extend the model for binary decisions proposed in (Seeley et al., 2012) (reported also in Chapter 4 in Equation (1)) to the best-of-n decision problem. The model describes the mean system behaviour as a system of n coupled ODEs and an algebraic equation for mass conservation:

$$\begin{cases} \frac{d\Psi_i}{dt} & = \gamma_i \Psi_U - \alpha_i \Psi_i + \rho_i \Psi_i \Psi_U - \sum_{j \neq i} \sigma_j \Psi_i \Psi_j \\ \Psi_U & = 1 - \sum_i \Psi_i \end{cases} , \quad i \in \{1, \ldots, n\} \tag{8}$$

Each differential equation in (8) describes the variation of the fraction of agents in each population. The fraction of agents committed to option i increases through discovery (at a rate γ_i) and through recruitment proportional to the population committed to i (at a rate $\rho_i \Psi_i$). Conversely, the fraction decreases through abandonment (at a rate α_i) or through cross-inhibition proportional to the contrasting populations (at a rate $\sum_{j \neq i} \sigma_j \Psi_j$,). All model parameters represent rates at which agents change their commitment state. Therefore, we assume all model parameters to be non-negative:

$$\alpha_i, \gamma_i, \rho_i, \sigma_i \geq 0 , \quad i \in \{1, \ldots, n\}. \tag{9}$$

For a decision-making problem based on the quality of the available options, all model parameters could be linked to the option quality v_i:

$$\alpha_i = f_\alpha(v_i),\ \gamma_i = f_\gamma(v_i),\ \rho_i = f_\rho(v_i),\ \sigma_i = f_\sigma(v_i), \tag{10}$$

where each function describes a specific relationship between transition rate and option quality (see (Pais et al., 2013) for an example).

This model describes the average proportion of agents in each population for a system with an infinite number of agents. It is deterministic and continuous in time and in the state space. The model can be exploited to determine the macroscopic behaviour corresponding to a given parameterisation, and to provide constraints to the possible parameterisations in order to obtain a desired system behaviour. This ultimately translates in constraints in the design of the relationship between option quality v_i and transition rates γ_i, α_i, ρ_i, and σ_i.

5.5.2 *Macroscopic description: finite-size, stochastic, time continuous, state-space discrete*

The ODE system introduced in Eq. (8) can be derived as a mean-field approximation for $N \to \infty$ of a finite-size Markov process (Seeley et al., 2012). We can represent this process in the generalised case of best-of-n decisions

through chemical reactions representing agents changing their commitment state, either spontaneously or by interacting with other agents:

$$\begin{aligned} C_U &\xrightarrow{\mathcal{Q}_{\gamma_i}} C_i \\ C_i &\xrightarrow{\mathcal{Q}_{\alpha_i}} C_U \\ C_U + C_i &\xrightarrow{\mathcal{Q}_{\rho_i}} 2C_i \\ C_i + C_j &\xrightarrow{\mathcal{Q}_{\sigma_j}} C_U + C_j, \qquad i \neq j, \end{aligned} \tag{11}$$

where the $\mathcal{Q}_{\lambda_i}, \lambda \in \{\alpha, \gamma, \rho, \sigma\}$ represent reaction constants (Gillespie, 1976). Starting from the above description, it is possible to derive the master equation, which describes the time evolution of the system as a stochastic, discrete-state process. More precisely, the master equation describes the time evolution of the probability mass function related to each possible state in which the system can be found:

$$\frac{d}{dt}P(\mathbf{N}, t) = \sum_{k=1}^{4n} [\beta_k - P(\mathbf{N}, t)\mathcal{Q}_k], \quad \forall \mathbf{N} \tag{12}$$

where $\mathbf{N} = \langle N_U, N_1, \ldots, N_n \rangle$ corresponds to the system state, k is an index for each of the $4n$ possible transitions, and the term β_k is the probability that the system is one transition k "away" from state $\mathbf{N}$ at time t, and undergoes the transition k in $(t, t + dt)$. The quantities $\mathcal{Q}_k$ are defined as follows:

$$\begin{aligned} \mathcal{Q}_1 &= N_U \mathcal{Q}_{\gamma_1} & \mathcal{Q}_2 &= N_1 \mathcal{Q}_{\alpha_1} \\ \mathcal{Q}_3 &= N_U N_1 \mathcal{Q}_{\rho_1} & \mathcal{Q}_4 &= \textstyle\sum_{j \neq 1} N_1 N_j \mathcal{Q}_{\sigma_j} \\ &\vdots & &\vdots \\ \mathcal{Q}_{4n-3} &= N_U \mathcal{Q}_{\gamma_n} & \mathcal{Q}_{4n-2} &= N_n \mathcal{Q}_{\alpha_n} \\ \mathcal{Q}_{4n-1} &= N_U N_n \mathcal{Q}_{\rho_n} & \mathcal{Q}_{4n} &= \textstyle\sum_{j \neq n} N_n N_j \mathcal{Q}_{\sigma_n} \end{aligned} \tag{13}$$

For instance, in the binary case with options A and B, the term $N_U N_A \mathcal{Q}_{\rho_A} dt$ represents the probability that a recruitment transition for option A occurs in the time interval dt, changing the system state from $\langle N_U, N_A, N_B \rangle$ to $\langle N_U - 1, N_A + 1, N_B \rangle$.

The transition rates of the ODE model of equation (8) have direct correspondence with the transition probabilities of the master equation (12). For the generic transition rate λ, the conversion formula is:

$$\mathcal{Q}_{\lambda_i} = \lambda_i N^{1-n_a}, \qquad \begin{matrix} \lambda \in \{\gamma, \alpha, \rho, \sigma\} \\ i \in \{1, \cdots, n\} \end{matrix} \tag{14}$$

where n_a is the number of populations involved in the transition. The factor N^{1-n_a} is consequence of the fact that the transition rates are used in

differential equations that contain population fractions, while the transition probabilities are used in combination with the total number of agents in the population. For transitions involving a single population, i.e., discovery and abandonment, we have $n_a = 1$ and therefore we obtain a direct correspondence between transition rates and probabilities per unit time:

$$\mathcal{Q}_{\gamma_i} = \gamma_i, \qquad \mathcal{Q}_{\alpha_i} = \alpha_i \tag{15}$$

Conversely, for recruitment and cross-inhibition (i.e., transitions that correspond to interactions between populations), $n_a = 2$ and therefore

$$\mathcal{Q}_{\rho_i} = \rho_i N^{-1}, \qquad \mathcal{Q}_{\sigma_i} = \sigma_i N^{-1} \tag{16}$$

These relations provide a mean to link the two macroscopic descriptions of the process, allowing to study the matching of the finite-size system dynamics to the mean-field ones.

At this description level, the model accounts for the stochastic fluctuations of the system due to finite-size effects (i.e., the influence of a finite system size N). Given the complexity of analytically solving the master equation (12), we analyse it through numerical simulation via the Gillespie algorithm (Gillespie, 1976). As we show in the case studies of Chapter 6, the numerical analysis reveals how the system behaviour departs from the predictions of the mean-field approximation. This model allows us to study the effects of the relationship between quality v and transition probabilities $\mathcal{Q}_\lambda$, and therefore to take decisions about the desired macroscopic dynamics at design time.

5.5.3 *Microscopic description: agent-based, stochastic, time/state-space discrete*

The average agent behaviour is modelled as a PFSM (see panel (A) in Figure 13) with $n + 1$ states $\{C_U, C_1, \ldots, C_n\}$ which represent the agent commitment state, and by $4n$ transition probabilities (four transitions for each option i), which determine the state change (either spontaneous or upon interaction with agents of a different population). Differently from the previous models, here the system changes state at discrete time steps of length τ. The $4n$ probabilities determine the individual behaviour and may be modulated according to the option quality v_i.

Upon discovery of option i, agents make a transition from state C_U to state C_i with probability $\mathcal{P}_\gamma(v_i)$ (on average). Similarly, upon abandonment of option i, agents make a transition from state C_i to state C_U with probability $\mathcal{P}_\alpha(v_i)$. Discovery and abandonment are spontaneous transitions, that is, the probability depends solely on the option quality v_i (as estimated by the agent itself). Conversely, the remaining transitions depend also on the size of the different sub-populations, which can be estimated upon interaction with other agents (see below). Recruitment to option i is modelled by a transition from state C_U to state C_i with probability $\mathcal{P}_{\Psi_i}\mathcal{P}_\rho(v_i)$, where the first

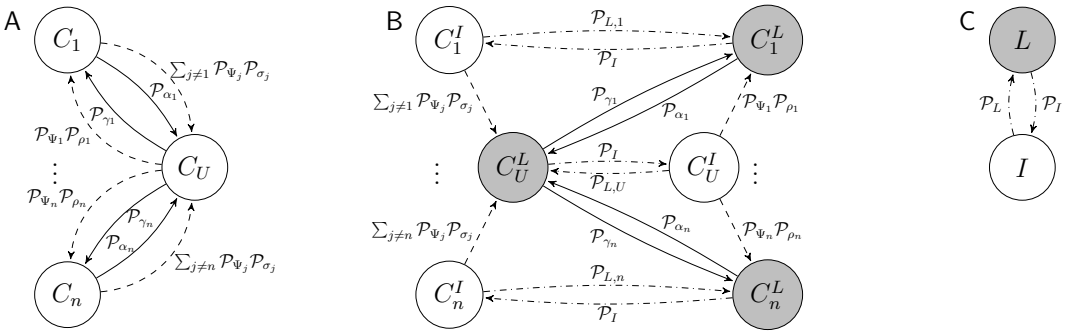

Figure 13: Probabilistic finite state machines (PFSMs) describing the microscopic behaviour of an agent on average. Here, the notation $\mathcal{P}_{\lambda_i}, \lambda \in \{\gamma, \alpha, \rho, \sigma\}$, $i \in \{1, \ldots, n\}$ is a shorthand for $\mathcal{P}_\lambda(v_i)$. (A) PFSM describing the basic commitment dynamics for n possible options. Spontaneous transitions are represented by solid lines, while interactive transitions are represented by dashed lines. (B) PFSM describing the coupled commitment and activity dynamics. Latent states are indicated in grey, and dash-dotted lines represent changes between latent and interactive states.

factor accounts for the probability of interacting with agents already committed to option i given the current population size, and the second factor accounts for a quality-dependent probability of triggering the state change. Similarly, cross-inhibition is modelled by a transition from state C_i to state C_U with probability $\sum_{j \neq i} \mathcal{P}_{\Psi_j} \mathcal{P}_\sigma(v_j)$. Here, the overall transition probability aggregates the probability of interaction with any agent committed to option $j \neq i$. Also in this case, the first factor accounts for the population size and the second factor accounts for a quality-dependent transition probability. The cross-inhibition of an agent committed to option i is influenced by the population size and by the quality of the contrasting options.

As already mentioned, at this level we describe the agent behaviour in average, which might not correspond to the actual implementation. In fact, the transition probabilities between commitment states might be implemented differently for each agent, as we will show in the following sections. The transition probabilities $\mathcal{P}_\lambda(v_i), \lambda \in \{\gamma, \alpha, \rho, \sigma\}$ presented above correspond to the average case, and they can be related to the probabilities per unit time of executing a transition at the macroscopic level, so that the following relations hold:

$$\begin{array}{lll} \mathcal{P}_\gamma(v_i) & = \mathcal{Q}_{\gamma_i}\tau & = \gamma_i\tau \\ \mathcal{P}_\alpha(v_i) & = \mathcal{Q}_{\alpha_i}\tau & = \alpha_i\tau \\ \mathcal{P}_{\Psi_i}\mathcal{P}_\rho(v_i) & = \mathcal{Q}_{\rho_i}N_i\tau & = \rho_i\tau N_i N^{-1} \\ \mathcal{P}_{\Psi_i}\mathcal{P}_\sigma(v_i) & = \mathcal{Q}_{\sigma_i}N_i\tau & = \sigma_i\tau N_i N^{-1} \end{array} \tag{17}$$

where τ is the discrete time step of the PFSM. Under the assumption of a well-mixed system, the probability of interaction with an agent committed to option i corresponds to the fraction Ψ_i, that is:

$$\mathcal{P}_{\Psi_i} = \frac{N_i}{N} \tag{18}$$

Considering also Eq. (10), we can derive a general quality-dependent relationship for the transition probability:

$$\lambda_i = f_\lambda(v_i) \;\rightarrow\; \mathcal{P}_\lambda(v_i) = f_\lambda(v_i)\tau, \qquad \begin{matrix} \lambda \in \{\gamma, \alpha, \rho, \sigma\} \\ i \in \{1, \ldots, n\} \end{matrix} \tag{19}$$

To obtain a desired macroscopic behaviour, one can opportunely define the average transition probabilities as a function of the quality v_i. Conversely, given the relationship between quality and individual probabilities, it is possible to easily derive the macroscopic dynamics.

5.5.4 *Implementation guidelines*

To proceed to the implementation of the agent behaviour, several *design choices* are required to determine how agents change state depending on either the population size or the option quality. The challenge is given by the fact that agents do not have access to global information—e.g., population size, number of available options—as described in the CDCI design pattern context. In such conditions, it is necessary to make design choices about the strategy for executing the state transitions of the PFSM to guarantee a one-to-one correspondence between the microscopic and macroscopic description levels.

5.5.4.1 *Population size dependent probabilities*

The computation of the probability $\mathcal{P}_{\Psi_i}$ requires a decentralised estimation of the population size given that, in the considered decentralised system, neither N_i nor N are available to the individual agent. Each agent compensates this lack of knowledge by estimating a probability P_{Ψ_i} through interactions with neighbours. A possible solution consists in letting each agent take a sample of the total population. This means that before taking action, the agent has to collect enough information about the size of the different populations by sampling the state of neighbour agents:

$$P_{\Psi_i} = \frac{|\tilde{\mathcal{A}}_i|}{|\tilde{\mathcal{A}}|}, \tag{20}$$

where $\tilde{\mathcal{A}}$ is the set of sampled agents, and $\tilde{\mathcal{A}}_i \subseteq \tilde{\mathcal{A}}$ is the set of sampled agents in state C_i. Depending on the pattern of interactions and on the sample size $|\tilde{\mathcal{A}}|$, the quality of the estimation varies. Therefore, it is necessary

to carefully design the sampling size in order to guarantee a certain level of accuracy.

Another possibility is to let each agent draw a random agent $a_{\hat{g}}$ from its local neighbourhood, check its commitment state and compute the probability P_{Ψ_i} as:

$$P_{\Psi_i} = \begin{cases} 1, & \text{if } C(a_{\hat{g}}) = C_i \\ 0, & \text{if } C(a_{\hat{g}}) \neq C_i \end{cases}, \qquad i \in \{1, \cdots, n\} \tag{21}$$

Assuming a well-mixed system, the frequency of picking an agent committed to option i is given exactly by Eq. (18), which corresponds to the desired behaviour on average. This second strategy is more parsimonious, as it requires fewer agent-agent interactions and no additional computations, and is therefore the choice that maximises speed.

5.5.4.2 *Homogeneous versus heterogeneous implementation*

For what concerns the other transition probabilities ($\mathcal{P}_\gamma$, $\mathcal{P}_\alpha$, $\mathcal{P}_\rho$, $\mathcal{P}_\sigma$), we propose two strategies based on either homogeneous or heterogeneous system implementation. In the homogeneous case, all agents share the same transition probability, leading to a direct correspondence between the actual and the average agent behaviour:

$$P_{\lambda,g}(v_i) = \mathcal{P}_\lambda(v_i), \qquad \begin{matrix} g \in \{1, \ldots, N\} \\ \lambda \in \{\gamma, \alpha, \rho, \sigma\} \\ i \in \{1, \ldots, n\} \end{matrix} \tag{22}$$

From (19), we can derive a first constraint to respect for a correct system implementation, following from the need that each probability must be less than 1:

$$P_{\lambda,g}(v_i) \leq 1 \;\rightarrow\; f_\lambda(v_i) \leq \frac{1}{\tau} \tag{23}$$

In the heterogeneous case, each agent a_g computes independently its own transition probability $P_{\lambda,g}(v_i)$, and the system behaviour results from the aggregation of the individual responses. At the macroscopic level, the transition probability per unit time (or conversely the transition rate) depends on the probability that any agent in a given population makes the corresponding transition. This depends on both the way in which the individual agent makes a transition, and on the heterogeneity of the system. In order to relate the macroscopic parameters to the individual probabilities and the option quality v_i, we propose to implement the transition probabilities with a simple response threshold scheme. The agent a_g makes a transition with a fixed probability if the (estimated) option quality exceeds a given response threshold δ_g:

$$P_{\lambda,g}(v_i) = \begin{cases} \mathcal{P}_{\lambda\uparrow} & \text{if } v_i > \delta_g \\ \mathcal{P}_{\lambda\downarrow} & \text{if } v_i \leq \delta_g \end{cases}, \tag{24}$$

where $\mathcal{P}_{\lambda\uparrow}$ and $\mathcal{P}_{\lambda\downarrow}$ are tuneable parameters, and the value δ_g is drawn for each agent a_g from a probability distribution $\mathcal{D}_\lambda$ over the range $[v_m, v_M]$. With this implementation, it is possible to establish a relationship between microscopic and macroscopic parameters through the cumulative distribution function of $\mathcal{D}_\lambda$, $F_{\mathcal{D}_\lambda}$:

$$F_{\mathcal{D}_\lambda} = \frac{\mathcal{P}_\lambda - \mathcal{P}_{\lambda\downarrow}}{\mathcal{P}_{\lambda\uparrow} - \mathcal{P}_{\lambda\downarrow}} \tag{25}$$

For $F_{\mathcal{D}_\lambda}$ to be a cumulative distribution function (CDF), it is required that the relationship between quality and macroscopic transition rate expressed in Eq. (10) be monotonic in v—either increasing or decreasing. As a consequence, the step function (24) can be determined by:

$$\begin{aligned} \mathcal{P}_{\lambda\uparrow} &= \mathcal{P}_\lambda(v_M) = f_\lambda(v_M)\tau \\ \mathcal{P}_{\lambda\downarrow} &= \mathcal{P}_\lambda(v_m) = f_\lambda(v_m)\tau \end{aligned} \quad , \tag{26}$$

which together with Eq. (19) provides the micro-macro link for the heterogeneous case.

Such a micro-macro link holds when each agent a_g re-samples the threshold δ_g from $\mathcal{D}_\lambda$ at every decision step. From an implementation perspective, however, re-sampling is not a parsimonious design choice, neither it is biologically plausible. Instead, fixed thresholds would be a more suitable solution: they would simplify the design and are also biologically plausible (e.g., response thresholds determined genetically or acquired through learning, Jeanson and Weidenmüller, 2013). However, fixed thresholds lead to a quasi-deterministic behaviour of the agents in face of a given option quality, as they are unable to modulate their behaviour according to the perceived quality. The decision problem would therefore lead to "frozen" sub-populations, and the microscopic dynamics would diverge from the macroscopic predictions. By studying the behaviour of different parameterisation, we recognised that an approximation with fixed thresholds is still valid for recruitment and cross-inhibition, because re-sampling is ensured by changing partner in each different interaction (as shown in case study I-A). Instead, the micro-macro link is hampered by the usage of fixed thresholds for spontaneous transitions, unless the macroscopic dynamics are dominated by recruitment and cross-inhibition (as shown in case study 1-B). Therefore, a principled choice about the usage of fixed thresholds can be made on the basis of the desired macroscopic parameterisation. Should the system be governed principally by spontaneous transitions, the fixed threshold scheme would not be suitable. Otherwise, it represents a viable solution, which is also biologically plausible (Robinson et al., 2011). Finally, note that homogeneous and heterogeneous strategies can be mixed together, so that the agent behaviour can be homogeneous with respect to some transition probabilities, and heterogeneous with respect to others.

5.5.4.3 *Latent and interactive agents*

As a further implementation guideline, we discuss here the case in which agents cannot interact every τ seconds. This is a very common condition in practical application scenarios, because of spatial/topological factors that determine the interaction pattern, or because interactions are constrained by limitations of the computing power or by the communication channel. To model such conditions, we introduce the possibility for agents to be either latent or interactive. When an agent is latent, it cannot communicate or receive messages from neighbours, but is still capable of changing its commitment state following spontaneous transitions. In the interactive state, agents are capable of communicating with other agents, and therefore can change commitment state accordingly. We refer to changes in the latent/interactive state as *activity dynamics*, as opposed to the commitment dynamics resulting into changes of the commitment state. We model the activity dynamics by considering that an agent becomes latent with probability $\mathcal{P}_L$ and returns interactive with probability $\mathcal{P}_I$ (see the PFSM in Figure 13(C)). Equivalently, agents may remain in the interactive or latent state for exponentially distributed time intervals, respectively with mean time $\tau_I = 1/\mathcal{P}_L$ and $\tau_L = 1/\mathcal{P}_I$. Under these conditions, the distribution of agents between interactive and latent states reaches asymptotically the fractions $\eta_I = \mathcal{P}_I/(\mathcal{P}_I + \mathcal{P}_L)$ and $\eta_L = \mathcal{P}_L/(\mathcal{P}_I + \mathcal{P}_L)$.

By coupling together activity and commitment dynamics, we obtain a microscopic description with $2(n+1)$ states. Here, agents can be uncommitted and latent (state C_U^L), uncommitted and interactive (state C_U^I), committed to option $i \in \{1, \ldots, n\}$ and latent (state C_i^L), or committed to i and interactive (state C_i^I). Transitions between these states can be arranged in different ways, constrained by the need to correctly represent both the commitment and the activity dynamics. Recruitment and cross-inhibition are available only when agents are interactive, and the final state can be either interactive or latent, depending on the application. Conversely, discovery and abandonment may be available in any state, and the actual choice depends on the application needs. An example PFSM is provided in panel (B) of Figure 13, and corresponds to the microscopic description for the search and exploitation task described in the case study II of Section 6.3. Here, the transition probabilities between different activity states must be appropriately tuned to reduce to the overall activity dynamics described by the PFSM of Figure 13(C). In the given example, the following relations hold:

$$\begin{matrix} \mathcal{P}_{L,i} = \mathcal{P}_L - \sum_{j \neq i} \mathcal{P}_{\Psi_j} \mathcal{P}_{\sigma_j} \\ \mathcal{P}_{L,U} = \mathcal{P}_L - \sum_i \mathcal{P}_{\Psi_i} \mathcal{P}_{\rho_i} \end{matrix} , \qquad i \in \{1, \ldots, n\}, \tag{27}$$

which ensure that the overall transition probability from interactive to latent states sums up to $\mathcal{P}_L$. Note that in the PFSM of Figure 13(B), there is al-

ways just one possible transition from a latent to an interactive state, with probability $\mathcal{P}_I$.

In order to maintain a micro-macro link despite the existence of latent agents, the population of interactive agents must always be an unbiased sample of the entire population. More precisely, given the fraction Ψ^I (Ψ^L) of agents in the interactive state I (latent state L), we require that:

$$\frac{\Psi_i^I}{\Psi^I} \approx \frac{\Psi_i^L}{\Psi^L} \approx \Psi_i, \qquad i \in \{1, \ldots, n\}, \tag{28}$$

where Ψ_i^I and Ψ_i^L represent the fractions of agents that are found in state C_i^I and C_i^L within the entire population. In fact, if changes in the commitment state within the interactive sub-population (fraction Ψ^I) are much faster than changes in the activity state (i.e., agents switching between states I and L), the distribution of commitment states among interactive agents would misrepresent the global population distribution, and therefore the microscopic and macroscopic dynamics would diverge. As a consequence, we require that the transitions in the activity state must be faster than transitions in the commitment state, for instance by constraining the decision to change commitment state to each transition from latent to interactive, as done in the search and exploitation task presented in Section 6.3.

The correspondence between microscopic and macroscopic parameters depends on the way in which microscopic transitions are implemented. Because some transitions are available only to interactive (latent) agents, the corresponding macroscopic rate must be reduced by η_I (η_L), which represent the fraction of the population that can actually change commitment state. Conversely, given a desired macroscopic transition rate, the average probabilities per agent must be increased by $1/\eta_I$ ($1/\eta_L$). For the example of Figure 13(B), recruitment and cross-inhibition transitions are available only when agents are interactive; therefore, Eq. (19) should be written as follows:

$$\lambda_i = f_\lambda(v_i) \;\rightarrow\; \mathcal{P}_\lambda(v_i) = \frac{f_\lambda(v_i)\tau}{\eta_I}, \qquad \begin{matrix} \lambda \in \{\rho, \sigma\} \\ i \in \{1, \ldots, n\} \end{matrix} \tag{29}$$

On the other hand, discovery and abandonment are only available to latent agents, and therefore the average probability per agent must take into account the proportion of agents in the latent state:

$$\lambda_i = f_\lambda(v_i) \;\rightarrow\; \mathcal{P}_\lambda(v_i) = \frac{f_\lambda(v_i)\tau}{\eta_L}, \qquad \begin{matrix} \lambda \in \{\gamma, \alpha\} \\ i \in \{1, \ldots, n\} \end{matrix} \tag{30}$$

5.5.4.4 *Minimum speed of the process*

The time step τ at which the agent updates its commitment state determines the process speed. In order to obtain a precise correspondence between macroscopic transition rates and microscopic transition probabilities,

the value of τ must be conveniently sized. To calculate an upper bound for τ, we consider the coexistence of transitions exiting from a same state of the PFSM, and we require that the total probability of leaving the state be lower than one. For instance, an agent that directly follows the behaviour described by the PFSM of Figure 13(A) has n recruitment and n discovery transitions all exiting from state C_U. This means that the overall probability from all outgoing transitions must not exceed one:

$$\sum_{i=1}^{n} \mathcal{P}_{\gamma}(v_i) + \mathcal{P}_{\Psi_i}\mathcal{P}_{\rho}(v_i) \leq 1 \tag{31}$$

Similarly, a constraint is given by the overall outgoing probability from the commitment state C_i:

$$\mathcal{P}_{\alpha}(v_i) + \sum_{j \neq i}^{n} \mathcal{P}_{\Psi_j}\mathcal{P}_{\sigma}(v_j) \leq 1 \tag{32}$$

Recall however that the PFSM of Figure 13(A) is the representation of the average agent, which differs from the actual implementation. The implementation guidelines described above prescribe that at most one interactive transition is available at a time. Additionally, we can assume that at most one spontaneous transition may become available at a time, given that in most application scenarios the evaluation of available alternatives is performed sequentially by individual agents. Overall, to compute the upper bound, we consider only one interactive and one spontaneous transitions at a time. A safe upper bound of τ is guaranteed by considering the extreme case in which $\mathcal{P}_{\Psi_i} = 1$ and v_i maximizes $\mathcal{P}_{\lambda}(v_i)$ (with $\lambda \in \{\gamma, \alpha, \rho, \sigma\}$):

$$\begin{cases} \max_{v_i} \mathcal{P}_{\rho}(v_i) + \max_{v_i} \mathcal{P}_{\gamma}(v_i) \leq 1 \\ \max_{v_i} \mathcal{P}_{\alpha}(v_i) + \max_{v_i} \mathcal{P}_{\sigma}(v_i) \leq 1 \end{cases}, \tag{33}$$

which, together with (19), can be rewritten as follows:

$$\begin{cases} \tau \leq \left(\max_{v_i} f_{\rho}(v_i) + \max_{v_i} f_{\gamma}(v_i)\right)^{-1} \\ \tau \leq \left(\max_{v_i} f_{\alpha}(v_i) + \max_{v_i} f_{\sigma}(v_i)\right)^{-1} \end{cases} \tag{34}$$

The constraint for the upper limit of the agent's time-step τ reduces to the minimum value of Eq. (34). Note that in Eq. (34), we specify a constraint on the multiagent system speed as a function of only macroscopic transition rates independently of the implementation strategy of the individual agent behaviour (whether homogenous, heterogenous or mixed). This upper-bound can be further refined considering the actual implementation, and the possible existence of latent and interactive states.

5.5.4.5 *Dealing with episodic discovery*

Discovery is the process that allows to report the existence of an option to the swarm. In many practical application scenarios, discovery is an episodic

event, that is, a given option is recognised by the agents only occasionally due to the limited individual agent capabilities. This can be related either to temporal or spatial constraints (e.g., discovery of options is correlated with the position in space of the agent, see the search and exploitation task presented in the case study II of Section 6.3). When no option is available, the agent cannot make a discovery transition. This has a bearing on the macroscopic dynamics resulting from the agent behaviour, because the macroscopic rate depends on the probability that an agent actually encounters the option i. If we refer to this probability as E_i, then we can rewrite Eq. (19) as follows:

$$\gamma_i = f_\gamma(v_i) \;\rightarrow\; \mathcal{P}_\gamma(v_i) = \frac{f_\gamma(v_i)\tau}{E_i} \tag{35}$$

Knowing how agents discover potential options is necessary to correctly link the microscopic and the macroscopic dynamics. In many practical scenarios, E_i can be estimated a priori to support the choice of the microscopic parameterisation.

5.5.4.6 *Skeleton of the agent behaviour*

Here, we present a generic structure of the individual agent behaviour, following the design strategies discussed above. The algorithm, reported in Figure 14, is rather general and, in a real application, each function needs to be implemented to meet the constraints and requirement of the specific application. Despite its simplicity, this algorithm contains all the instructions required by an agent to implement the CDCI. We assume that the main loop is repeated every τ seconds, where τ satisfies the requirements given above in Eq. (33). We also assume that each agent a_g, $g \in \{1, \ldots, N\}$, is characterised by its commitment state $C \in \{C_U, C_i\}$, $i \in \{1, \ldots, n\}$, which identifies the option i to which it is committed to.

When an agent is uncommitted, $C = C_U$, it explores the solution space to collect information about the available alternatives, and the discovered options are stored in the set $\mathcal{I}$ (see line 3). In real application scenarios, the cardinality of $\mathcal{I}$ may vary in $[0, n]$, because not every option can be discovered at a time. Upon discovery, the quality of the available options is estimated and stored in the set $\mathcal{V}$ (see line 4). However, agents need only to memorize the option they are committed to. Therefore, when the agent is committed for a specific option i, $\mathcal{I}$ contains only that option and $\mathcal{V}$ the corresponding quality. No other information needs to be stored by the agent (see lines 6-7).

Additionally, agents interact with each other in order to implement recruitment and cross-inhibition. In this respect, an agent samples the population to obtain the set $\mathcal{N}$ of neighbours which allows to compute the probability P_{Ψ_i} (see line 9). As mentioned in Section 5.5.4.1, two strategies

```
1: loop
2:    if (C == C_U) then /* the agent is uncommitted */
3:       I ← GetAvailableOptions()/* discovery of available options */
4:       V ← GetQualityEstimates(I)/* estimate option quality */
5:    else /* the agent is committed to some option */
6:       I ← {C}
7:       V ← {V}
8:    end if
9:    N ← InteractWithNeighbours()/* get neighbours' state for recruitment/inhibition */
10:   C ← UpdateCommitmentState(C, I, V, N)/* change commitment state */
11:   V ← GetQuality(C)/* store option quality */
12: end loop
```

Figure 14: Pseudo-code of the individual agent behaviour for the CDCI design pattern. We assume that the loop presented in this algorithm is executed every τ seconds.

are possible: (i) sampling for the current size of the different populations of committed agents, which implies to get information from a relatively large number of neighbours, or (ii) choosing a single neighbour to compute P_Ψ. In the latter case, the probability P_Ψ assumes a binary value, either 0 or 1, which is equivalent to activate or deactivate the respective transition. In this case, the agent can follow only a single recruitment/cross-inhibition transition per control loop. However, at the macroscopic level no difference is observable.

Finally, the new commitment state is computed at line 10 on the basis of the current commitment state C. The sets $\mathcal{I}$, $\mathcal{V}$ and $\mathcal{N}$ are exploited to activate potential transitions between the commitment states. When an agent gets committed to a given option, it stores the corresponding quality V for subsequent iterations (see line 11).

In theory, theory and practice are the same.
In practice, they are not.

—Albert Einstein

6

CASE STUDIES

Following the solution described in Chapter 5, we present here two case studies that showcase the usage of the proposed design pattern. Case study I, in Section 6.2, concerns decentralised decisions making by static agents interacting on a fully-connected network, and is divided in two parts, I-A and I-B. Case study II, presented in Section 6.3, concerns decentralised decisions by mobile agents in the context of a search and exploitation task. To ease the discussion and simplify the visualisation, we present here a binary decision problem (options A and B with quality v_A and v_B), and we report additional results for the best-of-n scenario only for the case study I-B.

To quantify the agreement between macroscopic models and microscopic implementation, we look at the system performance through a set of metrics detailed in Section 6.1, and we compare the process dynamics at different abstraction levels.

6.1 METRICS

Different metrics are used in the literature to evaluate the results of a collective decision-making process, which are linked to the correctness of the response, the coherence of the decisions within the group, as well as the speed of the process. Whenever time is required to gather sufficient information, decision making gives rise to speed-accuracy tradeoffs—a very common phenomenon in biological systems (Chittka et al., 2009). Accuracy measures the correctness of the decision and, in decentralised systems, can be defined as the proportion of the group that is committed to the best option, or to any of the equal-best options. Conversely, coherence measures the ability of the group to be committed to the same option, notwithstanding its quality (Franks et al., 2013). Therefore, one can simultaneously have low accuracy and high coherence, for instance if 10% of the group choses a high-quality option and the remaining 90% goes for a low-quality one. A high coherence of the group is important, as it can minimise the costs for conflict-

ing choices by individuals. As maintaining coherence is a time-consuming process that requires to spread information widely within the group, speed-coherence tradeoffs (also named speed-cohesion tradeoffs) may also appear (Franks et al., 2013). In the context of engineering artificial systems, it is important to quantify both the aspects of decision accuracy and group coherence, and constrast them with the time required to arrive at a decision.

To this purpose, we first introduce the **resolution** $\mathcal{R}$, which refers to the ability to discriminate between different-quality options, and is measured by the normalised quality difference between two options A and B: $\mathcal{R} = |v_A - v_B| / \max(v_A, v_B)$. Fixing a target resolution determines in which portion of the problem space the system accuracy is a relevant metric. In fact, any chosen solution that is below the given resolution threshold is correct, instead, when the difference between option's quality is above the resolution threshold, the designer requires that the superior option is selected. Note that resolution is normalised so that the minimum quality difference that can be detected is proportional to the quality magnitude, in analogy to many biological decision-making processes following Weber's law.

Then, we consider the **effectivity** $\mathcal{E}$ as the ability of the group to take a decision within the maximum execution time T. Effectivity is measured as the fraction of runs that reach the quorum Ψ_q within the given time limit. Effectivity is related to the coherence of decision making, as it measures the ability to take a decision (i.e., reach the predefined quorum) within the maximum allotted time, notwithstanding the quality of the chosen option. By requiring a minimum effectivity threshold, the designer can require that the system reaches a coherent state within a maximum time T.

Having defined resolution and effectivity, we introduce the main performance metrics we take into account. The **success rate** $\mathcal{S}$ corresponds to the fraction of effective runs resulting in a correct decision—i.e., the quorum is reached for the best option, or any of the equal-best options—when starting from a fully-uncommitted population. The success rate is similar to the exit probability in stochastic processes, and is related to the accuracy of decision making, because the quality of the chosen option is taken into account. Note that, by looking at the effective runs only—e.g., runs with effectivity larger than the given threshold—we limit ourselves only to high-coherence results.

The **convergence time** $\mathcal{C}$ is the average time required to reach the quorum Ψ_q computed over all effective runs, and is similar to the exit time of stochastic processes. This metric actually corresponds to the speed of the decision-making process, and can be exploited together with the success rate to select the most convenient solution that optimises the speed-accuracy/coherence tradeoff.

6.2 CASE STUDY I: COLLECTIVE DECISIONS IN A FULLY-CONNECTED MULTIAGENT SYSTEM

The first case study illustrates the implementation of decentralised decision making for a multiagent system in which each agent can potentially interact with any other agent. We implemented a synchronous simulation for a multiagent system with a fully-connected communication network that directly derives from the design pattern implementation guidelines. At simulation start, each agent a_g estimates the quality $\hat{v}_i$ of all available options $i \in \{1, \dots, n\}$, and on that basis computes its own transition probabilities $P_{\lambda,g}(\hat{v}_i)$, with $\lambda \in \{\gamma, \alpha, \rho, \sigma\}$. In the homogeneous case, these are computed in the same way for each agent according to the desired parameterisation, as prescribed by Eq. (19). In the heterogeneous case, each agent a_g draws a random threshold δ_g from the random distribution D_λ and computes the transition probabilities as prescribed by Eq. (24).

The simulation proceeds in discrete time steps of length τ. At each time step t, every agent updates its state following the PFSM of Figure 13(A). All spontaneous transitions are always available. Conversely, interactive transitions depend on the interaction with a randomly selected partner, who shares its own commitment state and probabilities of recruitment and cross-inhibition. Given the well-mixed property ensured by the fully-connected topology, the population-dependent probabilities $\mathcal{P}_{\Psi_i}$ are estimated by randomly choosing a different agent $a_{\hat{g}}$ as partner at each time step and checking its state: transitions are activated if the selected partner is committed to some option i (see Section 5.5.4.1 for details). In this case, the probability of recruitment $P_{\rho,\hat{g}}$ and of cross-inhibition $P_{\sigma,\hat{g}}$ are received from the selected partner, otherwise they are null. In this way, the agent a_g has complete information to update its commitment state.

We present two parameterisations as case study I-A and I-B.

6.2.1 *Case study I-A*

In case study I-A, we study consensus decisions, that is, we design a system in which the desired outcome is complete convergence of the group towards the choice of one or the other option. To this end, we set the decision quorum $\Psi_q = 1$ and we require that a decision is taken within $T = 400\,\text{s}$. Here, we also assume that option quality varies in $v \in [0, 1]$.

The first step towards implementation is the definition of the macroscopic parameterisation and its relationship with the option quality. The analysis of the macroscopic dynamics from Eq. (8) reveals that consensus can be achieved only when abandonment is null ($\alpha_A = \alpha_B = 0$). We arbitrarily choose a constant cross-inhibition rate $\sigma_A = \sigma_B = \bar{\sigma}$, which is sufficient for determining a collective decision (Pais et al., 2013). The value $\bar{\sigma}$ can be tuned

to determine the time scale of the process: the higher the rate, the quicker the convergence dynamics. Here, we choose $\bar{\sigma} = 1$. In these conditions, the model predicts two equilibrium points corresponding to consensus decision for either of the two options, but their stability may vary depending on the relative strength of discovery and recruitment. Assuming $v_A \geq v_B$, the model (8) predicts that the equilibrium at consensus for A is always stable, while consensus for B is stable only when $\gamma_A < \gamma_B + \rho_B$ (see Appendix B for a the detailed stability analysis). Thanks to this result, an informed choice can be made about the macroscopic parameterisation and the relation with option quality: $\gamma_i = f_\gamma(v_i)$, $\rho_i = f_\rho(v_i)$, $i \in \{\mathrm{A}, \mathrm{B}\}$. In particular, assuming a target resolution $R = 0.15$, we can minimise the chances of a wrong decision by designing the system to have a single stable equilibrium for the best option in any decision problem characterised by above-resolution quality differences. We select linear functions that link macroscopic transition rates to the quality:

$$f_\gamma(v_i) = k\, v_i, \qquad f_\rho(v_i) = h\, v_i \tag{36}$$

where k and h are tuneable parameters. Next, we compute the constraint on the above functions to satisfy our design choice: $k > h\,(1 - R)/R$ (see Appendix B for the details about the parameterisation choice). Finally, we choose values that comply with the prescribed bounds: $h = 0.1$ and $k = 0.6$.

The second step towards implementation is the analysis of the system performance in the complete decision space for varying system size N. This can be studied numerically by approximating the finite-size macroscopic dynamics using the Gillespie algorithm (Gillespie, 1976). Finally, the multiagent system can be deployed following the prescriptions of the design pattern and choosing a convenient implementation strategy. In the homogeneous case, all agents determine P_γ and P_ρ in the same way according to Eq. (19) and (36). Conversely, in the heterogeneous case transition probabilities are determined by the step function of Eq. (24), and vary from agent to agent with thresholds randomly sampled from the distribution determined by Eq. (25), (26) and (36). In this case study, we use fixed thresholds and we therefore limit the heterogeneous implementation to the recruitment probability P_ρ, while we keep the discovery probability P_γ homogeneous across agents. Finally, we let agents update their state every $\tau = 0.2\,\mathrm{s}$ following the design pattern guidelines (see Section 5.5.4.4).

The performance of the multiagent system for both homogenous and heterogeneous implementations is compared to the macroscopic Gillespie simulations for varying system size N (see Figure 15(A)). An excellent match between microscopic and macroscopic dynamics can be observed for every system size, for both the success rate $\mathcal{S}$ and the convergence time $\mathcal{C}$. When the difference in quality between the two options is above the resolution $R = 0.15$, the correct decision is taken in at least 90% of the cases (i.e., $\mathcal{S} = 0.9$, as evidenced by the isolines in the bottom-right part of Figure 15(A)

laying within the grey shaded area) for every system size but $N = 10$. Indeed, small groups suffer from stochastic fluctuations, reflected by a substantially lower success rate with respect to larger groups. Conversely, the speed of the process is lower for larger groups, as indicated by the isolines for $\hat{\mathcal{C}} = 50\,\mathrm{s}$ in the top-left part of Figure 15(A). To quantify the scaling properties with respect to the system size, we analysed the convergence time for each decision problem as a function of N. We found a generalised adherence with a power law behaviour $\mathcal{C} = b\,N^a$, with exponent $a \approx 0.2$ as shown in Figure 15(B) for macroscopic Gillespie simulations, and Figure 16 for multiagent simulations. The coefficient b also varies with the decision problem: the lower the option quality difference, the higher the coefficient. Looking at Figure 15(B) and 16, we observe that $\mathcal{C}$ scales similarly across different decision problems, with the exception of problems characterised by similar qualities (i.e., $v_A \approx v_B$) that require in general more time for convergence. Finally, we show in Figure 17 an example of the convergence dynamics for a specific decision problem (e.g., $v_A = 0.9$, $v_B = 0.6$), which highlights the close correspondence between ODEs, Gillespie and multiagent simulations.

6.2.2 *Case study I-B*

Case study I-B is concerned with the general case of value-sensitive decision making (Pais et al., 2013), and discusses the implementation in case of less restricting conditions with respect to the previous case study. We consider a quality range $v \in [1, 10]$, we fix the quorum for the collective decision to $\Psi_q = 0.8$ and we limit the total execution time to $T = 40\,\mathrm{s}$. Here, we also demonstrate a fully heterogeneous implementation of the multiagent system.

To obtain value-sensitive decision making (Pais et al., 2013), discovery and recruitment rates are assumed to be linearly proportional to the option quality v_i (i.e., $\gamma_i = \rho_i = v_i, i \in \{\mathrm{A}, \mathrm{B}\}$), the abandonment rate is inversely proportional to quality (i.e., $\alpha_i = 1/v_i$), while the cross-inhibition rate is constant ($\sigma_i = \bar{\sigma}$), which we fix to $\bar{\sigma} = 10$. Given the macroscopic parameterisation, we follow the same methodology described above to analyse the finite-size effects produced by the system size N, and we design the multiagent system following both the homogeneous and the heterogeneous strategies (see Section 5.5.4.2). In the latter case, we use fixed thresholds for all transition probabilities and discuss the error introduced with respect to the macroscopic dynamics. According to the design pattern prescriptions, agents are updated every $\tau = 20\,\mathrm{ms}$ (see Section 5.5.4.4).

Figure 15(C) shows the match between the macroscopic Gillespie simulations and the multiagent implementation with both the homogeneous and the heterogeneous strategy, for varying system size N (see also Figure 17 for an example of convergence dynamics). The correspondence between macro-

scopic model and microscopic implementation is remarkable also in this case. Even with the rough approximation of fixed thresholds for the heterogeneous case, we note a good match of the micro-macro dynamics, although not perfect especially for the convergence time $\mathcal{C}$ at low qualities. The results show that the studied parameterisation allows to reliably take decisions for above-resolution decision problems already with $N = 100$, as indicated by the success rate $\mathcal{S}$ in the bottom-right part of Figure 15(C). Conversely, the convergence time $\mathcal{C}$ is very similar across different system sizes. Also in this case, we analysed the scaling behaviour of the convergence time, and found adherence with a power law behaviour $\mathcal{C} = b\,N^a$, but with a very low exponent a (see Figure 15(D) for Gillespie simulations and Figure 16 for multiagent simulations). With the proposed parameterisation, $\mathcal{C}$ becomes nearly independent of the system size N in large parts of the problem space. The coefficient b is rather low too, indicating fast decisions especially for large differences in quality. This is a result of the higher transition rates chosen for the macroscopic model, which are reflected by a smaller timestep τ in the multiagent implementation as prescribed by the design pattern. Finally, we studied the micro-macro link in a best-of-n scenario. The results presented in Figure 18 reveal a very good correspondence between multiagent and Gillespie simulations, therefore validating the methodology beyond the binary decision problems presented above.

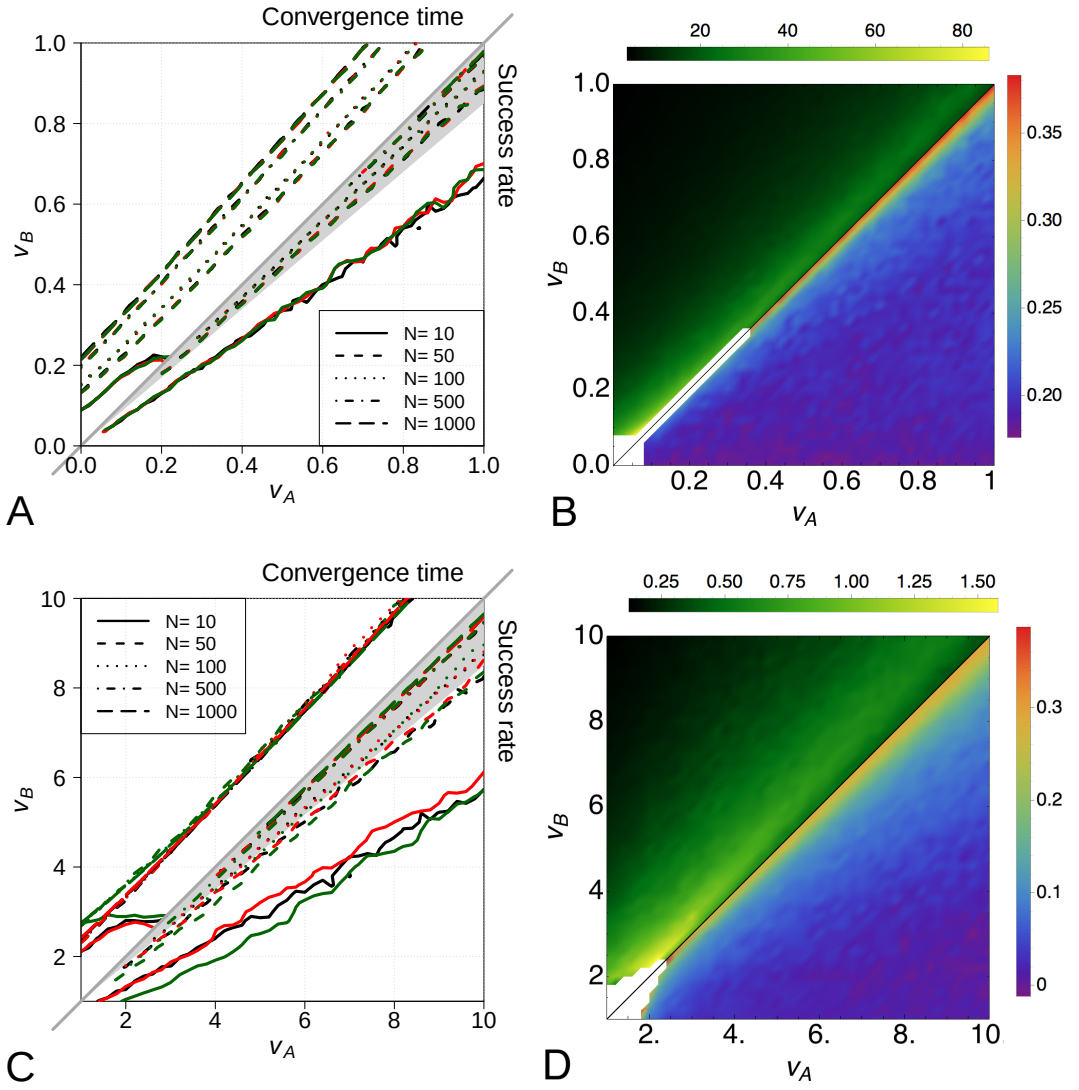

Figure 15: Collective decisions in a fully-connected multiagent system: comparison between the micro and the macro dynamics and scaling properties. (A),(C) Comparison between the stochastic finite-size macroscopic model (black lines) and the multiagent implementation with both the homogeneous strategy (red lines) and the heterogenous strategy (green lines). Results are displayed for varying system size N. For each possible configuration (v_A, v_B), 500 independent runs are performed. We show results for configurations with effectivity $\mathcal{E} > 0.7$. The plot is divided in two parts: in the bottom-right triangle, we consider the success rate $\mathcal{S}$ for each configuration (v_A, v_B), where $v_A \geq v_B$. For each group size N, we show the isolines at $\mathcal{S} = 0.9$. The gray triangle indicates quality value pairs below the target resolution $R = 0.15$ (i.e., configurations in which the two options are considered equivalent). In the top-left half of the plot, we consider the convergence time $\mathcal{C}$ for each configuration (v_A, v_B), where $v_B \geq v_A$ (i.e., the symmetric problems with respect to the bottom-right plot). For each group size N, we show the isolines at $\mathcal{C} = \hat{\mathcal{C}}$. (B),(D) Scaling of the convergence time $\mathcal{C}$ with the system size N. For each configuration (v_A, v_B), we fit the curve $\mathcal{C} = b\,N^a$ and we show the heat-map for the fitted coefficient a (see the bottom-right triangle showing the coefficient value for each configuration $(v_A, v_B), v_A \geq v_B$) and b (see the top-left triangle showing the coefficient value for symmetric configurations $(v_A, v_B), v_B \geq v_A$) across the decision space. Also in this case we show only configurations where $\mathcal{E} > 0.7$ for all N, and the white space indicates configurations with low effectivity. (A),(B) Results for case-study 1A with $v_i \in [0, 1]$, $\gamma_i = 0.6 v_i$, $\alpha_i = 0$, $\rho_i = 0.1 v_i$, $\sigma_i = 1$ and $i \in \{\mathrm{A}, \mathrm{B}\}$, $\hat{\mathcal{C}} = 50\,\mathrm{s}$. (C),(D) Results for case-study 1B with $v_i \in [1, 10]$, $\gamma_i = \rho_i = v_i$, $\alpha_i = 1/v_i$, $\sigma_i = 10$ and $i \in \{\mathrm{A}, \mathrm{B}\}$, $\hat{\mathcal{C}} = 1\,\mathrm{s}$.

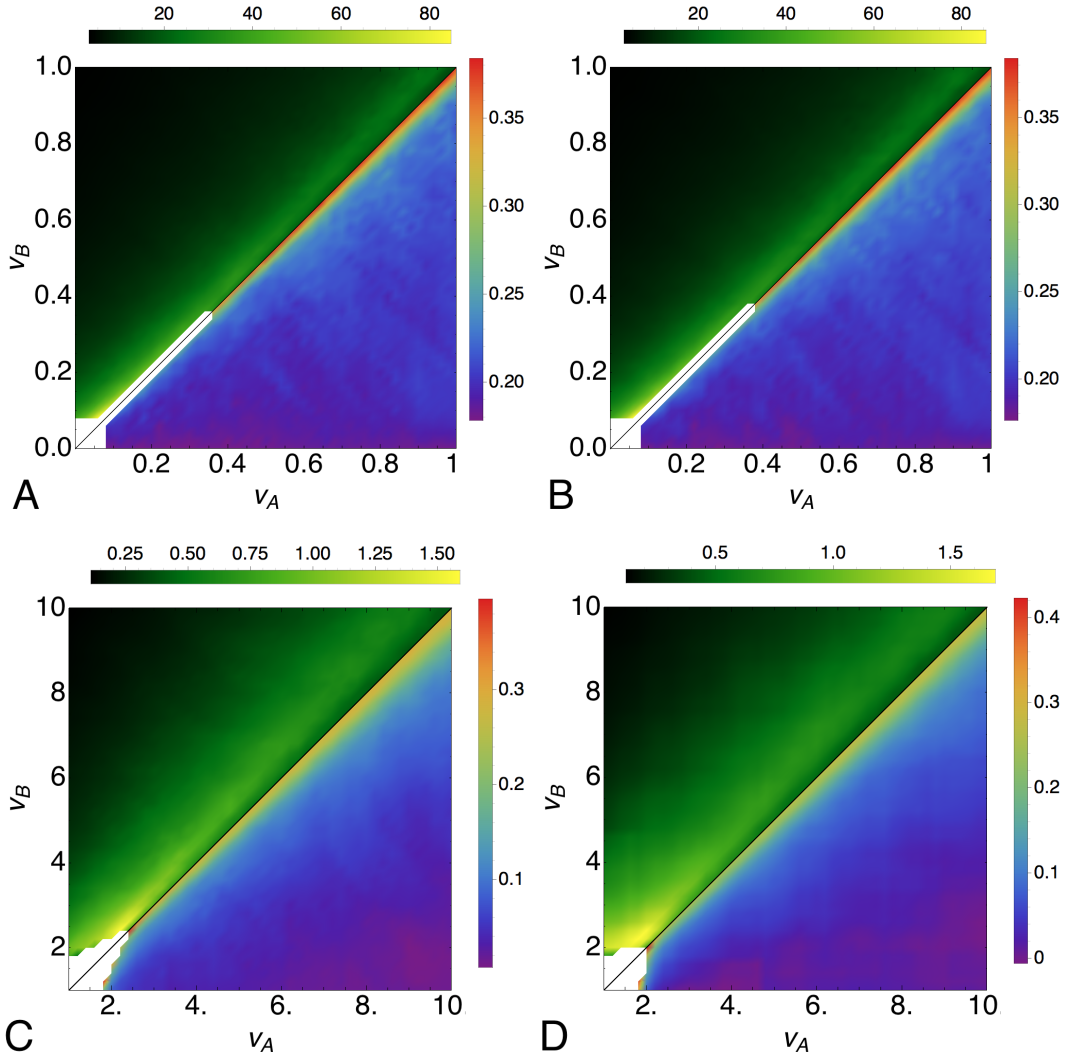

Figure 16: Scaling of the convergence time $\mathcal{C}$ with the system size N. For each configuration $(v_A, v_B), v_A > v_B$—and the symmetric case $(v_A, v_B), v_B > v_A$—we fit the curve $\mathcal{C} = b\, N^a$ and we show the heat-map for the fitted coefficient a (bottom-right) and b (top-left) across the decision space (see Figure 15 for details). Also in this case we show only configurations where $\mathcal{E} > 0.7$. (A,B) Results for case study I-A with $v_i \in [0,1]$, $\gamma_i = 0.6 v_i$, $\alpha_i = 0$, $\rho_i = 0.1 v_i$, $\sigma_i = 1$ and $i \in \{\mathrm{A}, \mathrm{B}\}$ for the homogenous (A) and the heterogeneous (B) implementation. (C,D) Results for case study I-B with $v_i \in [1, 10]$, $\gamma_i = \rho_i = v_i$, $\alpha_i = 1/v_i$, $\sigma_i = 10$ and $i \in \{\mathrm{A}, \mathrm{B}\}$ for the homogenous (C) and the heterogeneous (D) implementation.

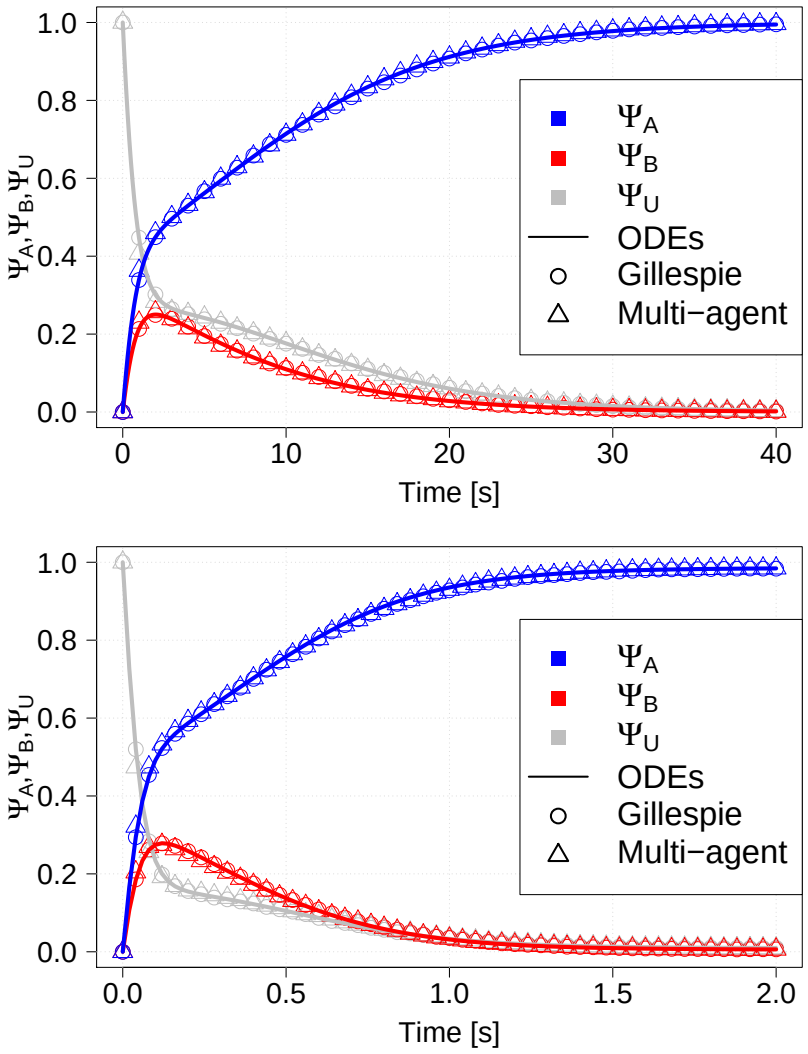

Figure 17: Mean trajectory of the population fractions Ψ_A, Ψ_B, Ψ_U over time from the initial condition $(\Psi_A = 0, \Psi_B = 0, \Psi_U = 1)$**.** Comparison at various levels of abstractions: mean field model (solid lines), macroscopic, finite-size master equation (circles), and multiagent simulations (triangles). Simulations results are averaged over 100 independent runs. Errorbars are smaller than the symbols size, and are not displayed. (Top) Parameterisation of case study I-A with homogenous multiagent implementation and $v_A = 0.9$, $v_B = 0.6$. (Bottom) Parameterisation of case study I-B with heterogenous multiagent implementation and $v_A = 9$, $v_B = 6$.

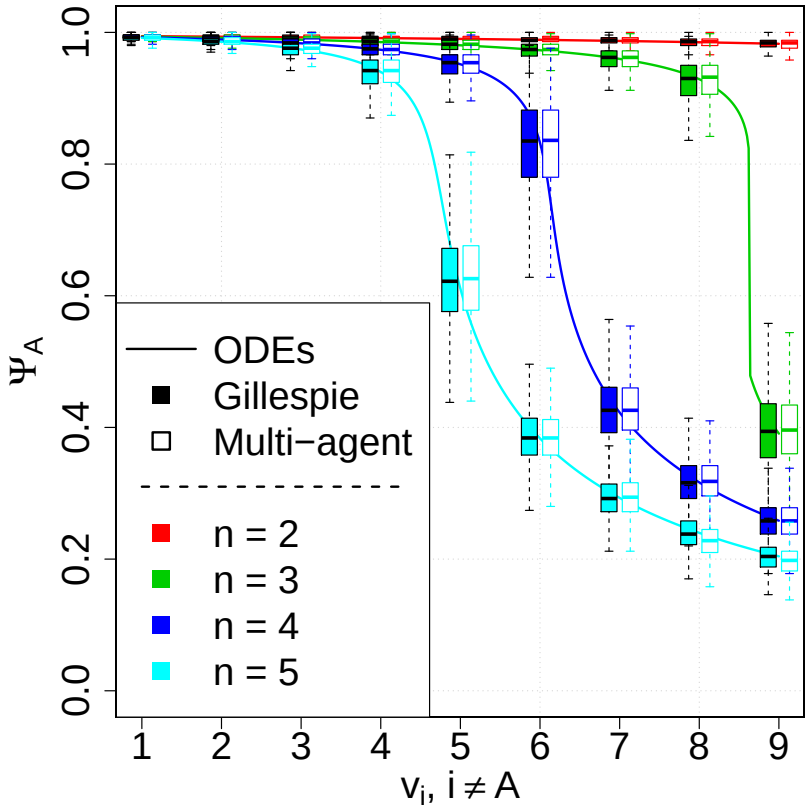

Figure 18: Micro-macro link with varying number of options for case study I-B. We compare the macroscopic dynamics predicted by the mean-filed model, the finite-size macroscopic dynamics simulated by the Gillespie algorithm and the microscopic dynamics resulting from homogeneous multiagent simulations ($N = 500$ agents). We fix the best option (A) to the maximum quality $v_A = 1$, and all other options to the same, lower quality v_i. The plot shows the fraction of the population committed to option A at the end of the simulation, plotted against the lower option quality v_i. Solid lines show the macroscopic prediction of the ODE system. The box-and-whiskers plots represent the statistics from Gillespie and multiagent simulations. Boxes represent the inter-quartile range of the data (2000 runs), while the horizontal lines inside the boxes mark the median values. The whiskers extend to the most extreme data points within 1.5 times the inter-quartile range. Outliers are not shown. A very good match can be appreciated between microscopic and macroscopic dynamics, therefore validating the design pattern for best-of-n scenarios.

6.3 CASE STUDY II: COLLECTIVE DECISIONS IN A SEARCH AND EXPLOITATION PROBLEM

Here, we present a case study that includes as main features spatiality and local interactions between agents. The experimental scenario is similar to the swarm robotics study of Chapter 4. Point-size agents move in a 2D environment characterised by three areas—a *home* and two *target* areas—and must select the best quality target. Uncommitted agents explore the environment to discover new options. Committed agents recruit and cross-inhibit other agents and periodically re-estimate the quality of the option they are committed to. Differently from Chapter 4, here quality is independent of distance. However, spatiality may influence the decision dynamics (e.g., the rate of discovery is higher for closer targets), and only an accurate design of the agent behaviour can lead to the systematic choice of the best available option.

PROBLEM DEFINITION Agents are point-size particles capable to move in a 2D environment. Movement is simulated through kinematic equations on the basis of the current agent speed v and orientation θ. The environment is an infinite plane and does not contain any wall or obstacle. No collision or physical interference among agents is taken into account, and agents are free to move in any location of the 2D plane. The environment contains three circular areas with radius $r = 0.3\,\text{m}$: home, target A and target B. Target areas are located at a variable distance $d_i \in [1.5\,\text{m}, 3.5\,\text{m}], i \in \{\text{A}, \text{B}\}$, from the home area. Each target area is further characterised by a quality $v_i \in [0.1, 1]$; each agent can individually estimate the target area quality when is inside the area.

Agents move at a constant speed $v = 0.01\,\text{m/s}$ and communicate with their neighbours within an interaction range $d_I = 0.6\,\text{m}$. Every $\tau = 1s$, agents update both their state and their motion direction θ as detailed below.

INTERACTIVE AND LATENT AGENTS As a consequence of local communications and of the distance between target areas, agents committed to different targets and uncommitted agents cannot always interact with each other. To ensure a well-mixed system, we limit interactions only when agents are within the home area. Agents remain in the home area for a time interval τ_I exponentially distributed with average $1/\mathcal{P}_L$. When a timeout expires, agents become latent and leave the home area. Agents remain latent for a time interval τ_L exponentially distributed with average $1/\mathcal{P}_I$. The motion pattern is conceived to ensure that agents are within the home area once they become interactive again (see below). To ensure that on average 10% of the agents are interactive within the home area, we set $\mathcal{P}_L = 9\mathcal{P}_I$.

MOTION PATTERN The agent motion direction θ is determined by the current agent state. While moving, odometry sensors are exploited to track the position of known areas. In this case study, we model noiseless sensors. In a more realistic implementation, agent-to-agent communication can be exploited to compensate for odometry noise (Gutiérrez et al., 2010).

An uncommitted agent a_g ($C(a_g) = C_U$) explores the environment in search of target areas: when latent, it chooses a random direction and leaves home moving in a straight line. Upon finding target area i, the agent makes an estimate of the quality $\hat{v}_i$ and gets committed with probability $P_{\gamma,g}(\hat{v}_i)$ (which is computed according to the homogeneous or heterogeneous implementation strategy). Then, it stores the target area estimated quality and location—which is kept updated through odometry—and returns back home. If an uncommitted agent has explored the environment for $\tau_L/2\,\mathrm{s}$ without encountering any target area, it returns home to ensure that it becomes interactive when already within the home area.

An agent a_g committed to option i ($C(a_g) = C_i$), while interactive, randomly moves within the home area to communicate with neighbours. When latent, it leaves home to return to the chosen target area i and re-estimate its quality $\hat{v}_i$. The agent returns home after $\tau_L/2\,\mathrm{s}$ to become interactive within the home area. While latent, the committed agent abandons commitment with probability $P_{\alpha,g}(\hat{v}_i)$ and returns home. This probability is also computed according to either the homogeneous or the heterogeneous implementation strategy.

INTERACTION PATTERN When interactive, all agents located in the home area can exchange short communication messages with a randomly chosen neighbour. These messages contain information on the agent state, which is used to estimate the population-dependent probabilities (see Section 5.5.4.1 for more details).

When committed to option i, agent a_g also communicates the stored location of the target area and its own probability of recruitment $P_{\rho,g}(\hat{v}_i)$ and cross-inhibition $P_{\sigma,g}(\hat{v}_i)$, computed either with the homogeneous or the heterogeneous implementation strategy. Recruitment takes place if the partner is uncommitted, otherwise cross-inhibition takes place only if the distance between the target area locations internally stored by the interacting agents is larger than the target area radius r. In this way, cross-inhibition takes place only between agents committed to different options, as prescribed by the design pattern.

When the population is divided between interactive and latent agents, the design pattern prescribes that the dynamics of activity change must be faster than changes in the commitment state. To achieve this, we let agents interact only upon becoming latent. By doing so, we guarantee that the

fraction of interactive agents committed to option i is always an unbiased representation of the global fraction Ψ_i (see Section 5.5.4.3).

INVESTIGATED PARAMETERISATION We choose a macroscopic parameterisation similar to the one of case study I-B, but option quality v_i varies in $[0.1, 1]$ and also the cross-inhibition rate varies linearly with quality (i.e., $\sigma_i = v_i$). Besides depending on option quality, discovery is episodic (see Section 5.5.4.5), being determined by the diffusive motion pattern of uncommitted agents that start searching from home. We model the macroscopic discovery rate to be proportional to the quality v_i and to decay with the target distance d_i as follows:

$$\gamma_i = \frac{v_i \, \mu e^{-\xi d_i}}{d_i} \tag{37}$$

where ξ and μ are parameters estimated from preliminary experiments ($\mu \approx 0.12$ and $\xi \approx 0.24$), although geometrical approximations could be used as well.

Besides discovery, spatiality influences also the interaction patterns among agents, given that interactions are possible only with agents in the local neighbourhood. To ensure a well-mixed system and comply with the design pattern requirements, we limit interactions within the home area and we force agents to periodically return home. As described above, the design pattern prescribes to have fixed probabilities to become interactive (e.g., return home with probability $\mathcal{P}_I$) or latent (e.g., leave home with probability $\mathcal{P}_L$). Therefore, we designed the individual behaviour after the microscopic description of Figure 13(B) following both the homogeneous and the heterogeneous strategy, and we set $\mathcal{P}_I = 0.001$ and $\mathcal{P}_L = 9\,P_I$ to ensure a fraction $\eta_I = 0.1$ of interactive agents on average.

The design choices detailed above allow us to determine the microscopic parameterisation starting from the desired macroscopic transition rates (see Figure 19 for a comparison between the rates determined by design and those estimated from multiagent simulations). We have tested the micro-macro link varying both option quality and target area distance, to observe how the implementation deals with the inclusion of spatial factors and interactive-latent dynamics. Indeed, the macroscopic model does not consider such factors, exception made for the model of discovery of Eq. (37). In particular, we are interested in making consistent choices notwithstanding the target area distance. Figure 20*A* shows one such case for a homogeneous system in which the better option is also the farthest. The macroscopic model predicts convergence on the best-quality option (B in this case, see the trajectory starting from $\Psi_U = 1$), and the simulations are centred at the predicted stable point. Good agreement between macroscopic Gillespie simulations and multiagent implementation is observable also for the success rate $\mathcal{S}$ (see Figure 20*A* inset). For same quality options, the target area dis-

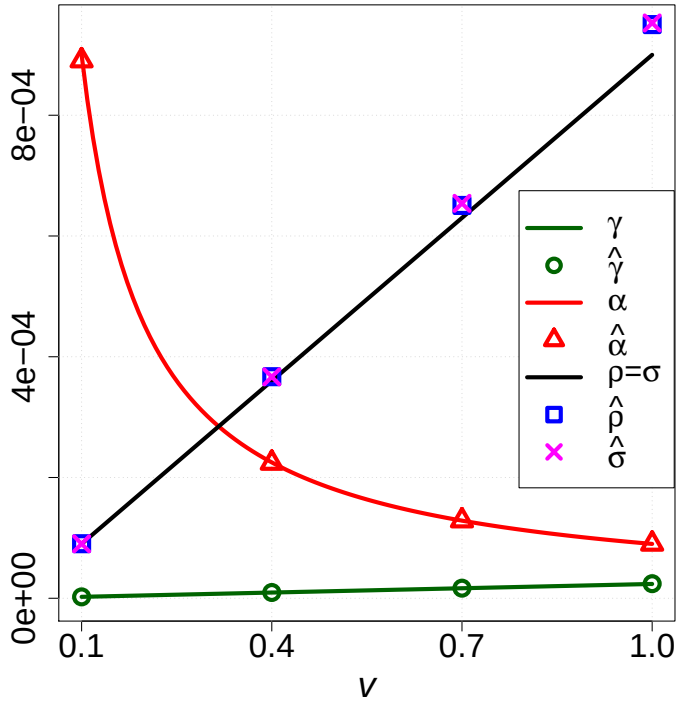

Figure 19: Comparison of the macroscopic transition rates resulting from design choices (solid lines) and estimated from the implemented multiagent system (points). The case for distance $d_A = d_B = 2.5\,\text{m}$ is shown. Estimates have been obtained through survival analysis computing the Nelson-Haelen estimator for the permanence time of agents in each state (Nelson, 1969). Survival analysis provides powerful non-parametric methods to estimate how the probability of events changes over time directly from the experimental data. See also Appendix A for details.

tance biases the choice towards the closer target area (see Figure S11 for the homogeneous case and Figure S12 for the heterogeneous case in the online supplementary material[1]). When both distance and option quality are equal (i.e., a completely symmetric condition), the system converges toward the one or the other option with equal probability, as shown in Figure 20*B* for the homogeneous case. Here too, the adherence between microscopic and macroscopic dynamics is remarkable. All these tests have been performed with $N = 500$ agents. Good agreement is observed also for different system sizes, as shown in Figure S13 in the online supplementary material[1].

1 http://iridia.ulb.ac.be/supp/IridiaSupp2016-001/index.html

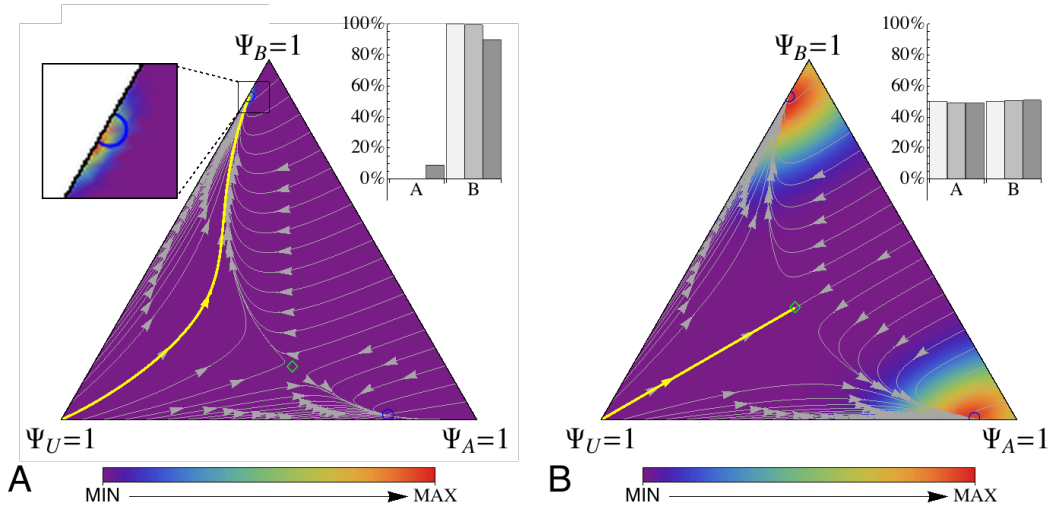

Figure 20: Collective decisions in a search and exploration problem: comparison between the micro and the macro dynamics. The state space of the system is presented as a ternary plot characterised by $\Psi_U + \Psi_A + \Psi_B = 1$, so that vertices correspond to fully-uncommitted or fully-committed populations. Macroscopic dynamics are indicated by trajectories and equilibrium points from the ODE model of Eq. (8), parameterised according to the specific configuration. The bold yellow trajectory indicates the behaviour starting from a fully-uncommitted population ($\Psi_U = 1$). Stable equilibrium points are indicated as blue empty circles, while unstable points are indicated as green empty diamonds. The density map in the background represents the results of homogeneous multiagent simulations (1000 runs). The inset shows the success rate $\mathcal{S}$ for macroscopic Gillespie simulations (white bars) and multiagent simulations (homogeneous in light gray and heterogeneous in dark grey). (*A*) Micro-macro link for a decision problem in which the best option is also the farthest one ($v_A = 0.7 < v_B = 1$ and $d_A = 1.5\,\mathrm{m} < d_B = 2.5\,\mathrm{m}$). The magnify-glass effect allows to appreciate the close correspondence between the stable point predicted by the macroscopic model and the results from the multiagent simulations. (*B*) Micro-macro link for a completely symmetric decision problem ($v_A = v_B = 1$ and $d_A = d_B = 2.5\,\mathrm{m}$).

An expert is a person who has made all the mistakes that can be made in a very narrow field.

—Niels Bohr

7

CONCLUSIONS

The design pattern methodology we propose provides a complete framework that allows moving from the choice of the macroscopic parameterisation down to the implementation of the individual behaviour. Each step is supported by the principled understanding of the causal relationship between microscopic choices and macroscopic effects. We have substantiated the methodology with case studies that, despite being idealised, contain all the ingredients to be taken as reference for practical applications. In this respect, the inclusion of latent states for individual agents is particularly important, as it allows to preserve the micro-macro link also when interactions are sporadic or when spatiality interferes with the well-mixed assumption. Indeed, departures from the macroscopic predictions are expected in case of heterogeneous interaction topologies, as it happens in other ordering processes such as the naming game (Baronchelli et al., 2006; Trianni et al., 2016b) or the voter model (Sood and Redner, 2005). In this case, the micro-macro link could be preserved through the inclusion of heterogeneous mean-field approximations, which can correct the departure from the assumed well-mixed condition at the macroscopic level (Moretti et al., 2012). Future work should also take into account the macroscopic effects of interactions over adaptive and multi-layer networks, in order to (i) take into consideration the variability of the topology of interactions with time (Groß and Blasius, 2008) and (ii) allow for the existence of different layers of connectivity among agents, each pertaining to specific context-related properties (Boccaletti et al., 2014).

In the near future, we plan to implement the proposed design pattern in a collective decision-making scenario involving a swarm of hundreds of robots (in particular, we plan to use the kilobot platform, Rubenstein et al., 2014). Preliminary results confirm the viability of these experiments. We believe that the significance of the proposed design pattern is not limited to swarm robotics but extends to other swarm systems, such as cognitive radios or cyber-physical systems. In fact, in (Trianni et al., 2016a), the CDCI

design pattern has been implemented to allow fair bandwidth allocation of TV White Space (TVWS) spectrum in decentralised networks. While this study has been conducted only in simulation, it has the potential to become a real-world application of the design pattern proposed in this thesis. Despite the generality of the proposed method, we do not think that every swarm system may benefit from it. For instance, blockchain systems aim to collective consensus (thus, collective decision making) and may appear as a potential application domain of the CDCI design pattern. However, a blockchain works with the very strong assumption that any agent may act maliciously to tamper with the correct system functioning (Nakamoto, 2008). Our design pattern cannot be employed in such systems because does not take into consideration the case of malicious agents but assumes that all the agents are cooperative.

Besides engineering, our results can be relevant for better understanding the behaviour of natural systems. With respect to honeybee nest-site selection, our results provide testable hypotheses about the algorithm employed by individual bees in relation to the proposed macroscopic model (Seeley et al., 2012; Pais et al., 2013). Our algorithm is simpler than other individual-based approaches (Passino and Seeley, 2006; Janson et al., 2007), as it abstracts several details that require assumptions difficult to be verified experimentally. Similarly to previous studies, (Janson et al., 2007; Laomettachit et al., 2015), we have shown in case study II, in Section 6.3, that spatiality affects the outcome of the collective decision biasing it towards closer sites. In our implementation, this is mainly the result of the quicker discovery of closer sites, while different latencies (e.g., shorter travel times) play a negligible role, thanks to the fixed probability of becoming interactive required to preserve the micro-macro link (as discussed in Section 5.5.4.3). Field experiments should be made to verify the existence of a tradeoff between distance and quality. Furthermore, the effects of finite-size groups in the decision dynamics predicted by the stochastic macroscopic simulations adhere with studies about group-size effects in natural conditions (Schaerf et al., 2013): the larger the swarm the more accurate its decision. The scaling of decision time with group size that we have highlighted here represents another interesting aspect to investigate with field experiments. The distinction between latent and interactive states is also biologically relevant and it has been included in some honeybee nest-site selection models (Britton et al., 2002). We show how this microscopic state distinction may impact on the macroscopic dynamics. Finally, it would be interesting to study the extent to which behavioural heterogeneities influence honeybee nest-site selection, as genetic and molecular determinants of honeybee behaviour seem to play an important role (Mattila and Seeley, 2007; Liang et al., 2012).

Behaviour heterogeneity in social systems is an important aspect not to be overlooked, as it can lead to interesting collective dynamics that are

not attainable in fully homogeneous systems (Martino and Marsili, 2006; Helbing, 2012; Huang et al., 2015). In our study, the choice of response thresholds for the heterogeneous implementation strategy is supported by the large literature on inter-individual variability in social insects (Burns, 2005; Mailleux et al., 2006; Dussutour et al., 2008; Robinson et al., 2011; Jeanson and Weidenmüller, 2013). Recent studies have recognised the importance of including individual differences in behaviour—often referred to as personality or behavioural syndrome (Jandt et al., 2014)—to better understand the collective dynamics (Wray and Seeley, 2011; Planas-Sitjà et al., 2015). Here, we have highlighted the relationship between the distribution of individual thresholds and the collective response function, so that macroscopic predictions could be matched against estimates of the real threshold distribution (Weidenmüller, 2004). We have also shown that fixed response thresholds well approximate the macroscopic dynamics especially for interactive processes like recruitment (Robinson et al., 2011). Associating fixed response thresholds with variable probability and intensity of responses may lead to more flexible and robust behaviour at the colony level (Jeanson and Weidenmüller, 2013). Response thresholds are linked with adaptive mechanisms for threshold adaptation (Theraulaz et al., 1998; Weidenmüller, 2004), allowing to fine tune the macroscopic response to match the statistical regularities that characterise the task. This adaptivity can result from evolutionary factors (Duarte et al., 2012) as well as from development and learning (Jeanson and Weidenmüller, 2013). Integrating adaptive mechanisms in the microscopic implementation could lead to improved performance (Franklin et al., 2012; Westhus et al., 2013), and represents a natural extension of the proposed design pattern.

BIBLIOGRAPHY

W. Agassounon, A. Martinoli, and K. Easton. Macroscopic modeling of aggregation experiments using embodied agents in teams of constant and time-varying sizes. *Autonomous Robots*, 17(2-3):163–192, 2004.

I. F. Akyildiz, W.-Y. Lee, M. C. Vuran, and S. Mohanty. Next generation/dynamic spectrum access/cognitive radio wireless networks: a survey. *Computer networks*, 50(13):2127–2159, 2006.

I. F. Akyildiz, B. F. Lo, and R. Balakrishnan. Cooperative spectrum sensing in cognitive radio networks: A survey. *Physical Communication*, 4(1):40–62, 2011.

C. Alexander, S. Ishikawa, and M. Silverstein. *A Pattern Language: Towns, Buildings, Construction*. Oxford University Press, 1977.

A. P. Athreya and P. Tague. Survivable smart grid communication: Smart-meters meshes to the rescue. In *Computing, Networking and Communications (ICNC), 2012 International Conference on*, pages 104–110, 2012.

O. Babaoğlu, G. Canright, A. Deutsch, G. A. Di Caro, F. Ducatelle, L. M. Gambardella, N. Ganguly, M. Jelasity, R. Montemanni, A. Montresor, and T. Urnes. Design patterns from biology for distributed computing. *Transactions on Adaptive and Autonomous Systems*, 1(1):26–66, 2006.

R. Bachmayer and N. E. Leonard. Vehicle networks for gradient descent in a sampled environment. In *Proceedings of the 41st IEEE Conference on Decision and Control*, volume 1, pages 112–117. IEEE Press, 2002.

A. Baronchelli, L. Dall'Asta, A. Barrat, and V. Loreto. Topology-induced coarsening in language games. *Physical review E, Statistical, nonlinear, and soft matter physics*, 73(1):015102, 2006.

F. Bartumeus, M. G. E. Da Luz, G. M. Viswanathan, and J. Catalan. Animal search strategies: A quantitative random-walk analysis. *Ecology*, 86(11): 3078–3087, 2005.

J. Beal. Programming an amorphous computational medium. In *Proceedings of the 2004 International Conference on Unconventional Programming Paradigms*, UPP'04, pages 121–136, Berlin, Germany, 2005. Springer.

J. Beal and J. Bachrach. Infrastructure for engineered emergence in sensor/actuator networks. *IEEE Intelligent Systems*, 21:10–19, 2006.

J. Beal and M. Viroli. Space–time programming. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 373(2046), 2015.

S. Berman, V. Kumar, and R. Nagpal. Design of control policies for spatially inhomogeneous robot swarms with application to commercial pollination. In *Proceedings of the 2011 IEEE International Conference on Robotics and Automation (ICRA)*, pages 378–385. IEEE Press, 2011a.

S. Berman, R. Nagpal, and A. Halasz. Optimization of stochastic strategies for spatially inhomogeneous robot swarms: A case study in commercial pollination. In *Proceedings of the IEEE/RSJ 2011 International Conference on Intelligent Robots and Systems (IROS)*, pages 3923–3930. IEEE Press, 2011b.

S. Berman, A. Halász, V. Kumar, and S. Pratt. Bio-inspired group behaviors for the deployment of a swarm of robots to multiple destinations. In *Proceedings of the 2007 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2318–2323. IEEE Press, 2007a.

S. Berman, A. Halász, V. Kumar, and S. Pratt. Algorithms for the analysis and synthesis of a bio-inspired swarm robotic system. In E. Şahin, W. M. Spears, and A. F. Winfield, editors, *Swarm Robotics*, volume 4433 of *LNCS*, pages 56–70. Springer, Berlin, Germany, 2007b.

S. Berman, A. Halász, M. A. Hsieh, and V. Kumar. Optimized Stochastic Policies for Task Allocation in Swarms of Robots. *IEEE Transactions on Robotics*, 25(4):927–937, 2009.

S. Bhattacharya, R. Ghrist, and V. Kumar. Multi-robot coverage and exploration on riemannian manifolds with boundaries. *The International Journal of Robotics Research*, 33(1):113–137, 2014.

S. Boccaletti, G. Bianconi, R. Criado, C. del Genio, J. Gómez-Gardeñes, M. Romance, I. Sendiña-Nadal, Z. Wang, and M. Zanin. The structure and dynamics of multilayer networks. *Physics Reports*, 544(1):1–122, 2014.

E. Bonabeau, M. Dorigo, and G. Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, New York, 1999.

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

M. Brambilla, A. Brutschy, M. Dorigo, and M. Birattari. Property-driven design for robot swarms. *ACM Transactions on Autonomous and Adaptive Systems*, 9(4):17:1–17:28, 2015.

N. F. Britton, N. R. Franks, S. C. Pratt, and T. D. Seeley. Deciding on a new home: how do honeybees agree? *Proceedings of the Royal Society of London B: Biological Sciences*, 269(1498):1383–1388, 2002.

J. Burns. Impulsive bees forage better: the advantage of quick, sometimes inaccurate foraging decisions. *Animal Behaviour*, 70(6):e1–e5, 2005.

C. Castellano, S. Fortunato, and V. Loreto. Statistical physics of social dynamics. *Reviews of Modern Physics*, 81(2):591–646, 2009.

L. Chittka, P. Skorupski, and N. E. Raine. Speed-accuracy tradeoffs in animal decision making. *Trends in Ecology & Evolution*, 24(7):400–407, 2009.

A. L. Christensen, R. O'Grady, and M. Dorigo. From fireflies to fault-tolerant swarms of robots. *IEEE Transactions on Evolutionary Computation*, 13(4):754–766, 2009.

E. A. Codling, M. J. Plank, and S. Benhamou. Random walk models in biology. *Journal of The Royal Society, Interface*, 5(25):813–34, 2008.

N. Correll and A. Martinoli. System identification of self-organizing robotic swarms. In M. Gini and R. Voyles, editors, *Distributed Autonomous Robotic Systems 7 (DARS)*, pages 31–40. Springer Japan, 2006.

N. Correll and A. Martinoli. Modeling and designing self-organized aggregation in a swarm of miniature robots. *The International Journal of Robotics Research*, 30(5):615–626, 2011.

J. Cortes, S. Martinez, T. Karatas, and F. Bullo. Coverage control for mobile sensing networks. *IEEE Transactions on Robotics and Automation*, 20(2):243–255, 2004.

P. Dames, M. Schwager, V. Kumar, and D. Rus. A decentralized control policy for adaptive information gathering in hazardous environments. In *Proceedings of the 51st IEEE Conference on Decision and Control (CDC)*, pages 2807–2813, 2012.

K. Dantu, S. Berman, B. Kate, and R. Nagpal. A comparison of deterministic and stochastic approaches for allocating spatially dependent tasks in micro-aerial vehicle collectives. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 793–800, 2012.

T. De Wolf and T. Holvoet. Design patterns for decentralised coordination in self-organising emergent systems. In S. Brueckner, S. Hassas, M. Jelasity, and D. Yamins, editors, *Engineering Self-Organising Systems*, volume 4335 of *Lecture Notes in Computer Science*, pages 28–49. Springer Berlin Heidelberg, 2007.

P. Derler, E. Lee, and A. Sangiovanni Vincentelli. Modeling cyber-physical systems. *Proceedings of the IEEE*, 100(1):13–28, 2012.

J. P. Desai, J. P. Ostrowski, and V. Kumar. Modeling and control of formations of nonholonomic mobile robots. *IEEE Transactions on Robotics and Automation*, 17(6):905–908, 2001.

M. Dorigo and M. Birattari. Swarm intelligence. *Scholarpedia*, 2(9):1462, 2007.

M. Dorigo, M. Birattari, and M. Brambilla. Swarm robotics. *Scholarpedia*, 9 (1):1463, 2014.

F. Dressler. *Self-organization in sensor and actor networks*. John Wiley & Sons, 2008.

J. Driesen and F. Katiraei. Design for distributed energy resources. *IEEE Power and Energy Magazine*, 6(3):30–40, 2008.

A. Duarte, I. Pen, L. Keller, and F. Weissing. Evolution of self-organized division of labor in a response threshold model. *Behavioral Ecology and Sociobiology*, 66(6):947–957, 2012.

A. Dussutour, S. Nicolis, E. Despland, and S. Simpson. Individual differences influence collective behaviour in social caterpillars. *Animal Behaviour*, 76(1):5–16, 2008.

J. L. Fernandez-Marquez, G. Di Marzo Serugendo, S. Montagna, M. Viroli, and J. L. Arcos. Description and composition of bio-inspired design patterns: a complete overview. *Natural Computing*, 12(1):43–67, 2013.

E. Ferrante, A. E. Turgut, M. Dorigo, and C. Huepe. Elasticity-based mechanism for the collective motion of self-propelled particles with spring-like interactions: A model system for natural and artificial swarms. *Physical Review Letters*, 111(268302):1–5, 2013.

E. Fiorelli, N. E. Leonard, P. Bhatta, D. A. Paley, R. Bachmayer, and D. M. Fratantoni. Multi-AUV control and adaptive sampling in Monterey Bay. *IEEE Journal of Oceanic Engineering*, 31(4):935–948, 2006.

G. Francesca, M. Brambilla, A. Brutschy, L. Garattoni, R. Miletitch, G. Podevijn, A. Reina, T. Soleymani, M. Salvaro, C. Pinciroli, F. Mascia, V. Trianni, and M. Birattari. AutoMoDe-Chocolate: automatic design of control software for robot swarms. *Swarm Intelligence*, 9(2-3):125–152, 2015.

G. Francesca and M. Birattari. Automatic design of robot swarms: achievements and challenges. *Frontiers in Robotics and AI*, 3(29), 2016.

G. Francesca, M. Brambilla, A. Brutschy, L. Garattoni, R. Miletitch, G. Podevijn, A. Reina, T. Soleymani, M. Salvaro, C. Pinciroli, V. Trianni, and M. Birattari. An experiment in automatic design of robot swarms: AutoMoDe-Vanilla, EvoStick, and human experts. In M. Dorigo et al., editor, *Swarm Intelligence (ANTS 2014)*, volume 8667 of *LNCS*, pages 25–37. Springer, Berlin, Germany, 2014a.

G. Francesca, M. Brambilla, A. Brutschy, V. Trianni, and M. Birattari. AutoMoDe: A novel approach to the automatic design of control software for robot swarms. *Swarm Intelligence*, 8(2):89–112, 2014b.

E. Franklin, E. Robinson, J. Marshall, A. Sendova-Franks, and N. Franks. Do ants need to be old and experienced to teach? *The Journal of Experimental Biology*, 215(8):1287–1292, 2012.

N. Franks, T. Richardson, N. Stroeymeyt, R. Kirby, W. Amos, P. Hogan, J. Marshall, and T. Schlegel. Speed-cohesion trade-offs in collective decision making in ants and the concept of precision in animal behaviour. *Animal Behaviour*, 85(6):1233–1244, 2013.

E. Gamma, R. Helm, R. Johnson, and J. Vlissides. *Design Patterns: Elements of Reusable Object-oriented Software*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1995.

L. Gardelli, M. Viroli, and A. Omicini. Design patterns for self-organising systems. In *Multi-Agent Systems and Applications V*, volume 4696 of *LNCS*, pages 123–132. Springer, Berlin, Germany, 2007.

S. Garnier, C. Jost, J. Gautrais, M. Asadpour, G. Caprari, R. Jeanson, A. Grimal, and G. Theraulaz. The embodiment of cockroach aggregation behavior in a group of micro-robots. *Artificial life*, 14(4):387–408, 2008.

V. Gazi and K. M. Passino. Stability analysis of swarms. *IEEE Transactions on Automatic Control*, 48(4):692–697, 2003.

V. Gazi. Swarm aggregations using artificial potentials and sliding-mode control. *IEEE Transactions on Robotics*, 21(6):1208–1214, 2005.

V. Gazi and B. Fidan. Coordination and control of multi-agent dynamic systems: Models and approaches. In E. Şahin, W. M. Spears, and A. F. T. Winfield, editors, *Swarm Robotics*, volume 4433 of *LNCS*, pages 71–102. Springer, Berlin, Germany, 2007.

V. Gazi and K. M. Passino. A class of attraction/repulsion functions for stable swarm aggregations. In *Proceedings of the 41st IEEE Conference on Decision and Control*, volume 3, pages 2842–2847, 2002.

V. Gazi and K. M. Passino. Stability analysis of social foraging swarms. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, 34 (1):539–557, 2004.

A. Ghassemi, S. Bavarian, and L. Lampe. Cognitive radio for smart grid communications. In *Smart Grid Communications (SmartGridComm)*, pages 297–302. IEEE Press, 2010.

D. T. Gillespie. A General Method for Numerically Simulating Stochastic Time Evolution of Coupled Chemical Reactions. *Journal of Computational Physics*, 22:403–434, 1976.

T. Groß and B. Blasius. Adaptive coevolutionary networks: a review. *Journal of The Royal Society, Interface*, 5(20):259–271, 2008.

A. Gutiérrez, A. Campo, M. Dorigo, J. Donate, F. Monasterio-Huelin, and L. Magdalena. Open e-puck range and bearing miniaturized board for local communication in swarm robotics. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 3111–3116, 2009.

A. Gutiérrez, A. Campo, F. Monasterio-Huelin, L. Magdalena, and M. Dorigo. Collective decision-making based on social odometry. *Neural computing & applications*, 19(6):807–823, 2010.

J. Guzzi, A. Giusti, L. M. Gambardella, and G. A. Di Caro. Human-friendly robot navigation in dynamic environments. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 423–430, Karlsruhe, Germany, 2013.

A. Halász, M. Hsieh, S. Berman, and V. Kumar. Dynamic redistribution of a swarm of robots among multiple sites. In *Proceedings of the IEEE/RSJ 2007 International Conference on Intelligent Robots and Systems (IROS)*, pages 2320–2325. IEEE Press, 2007.

J. Halloy, G. Sempo, G. Caprari, C. Rivault, M. Asadpour, F. Tache, I. Said, V. Durier, S. Canonge, and M. Amé, J. Social integration of robots into groups of cockroaches to control self-organized choices. *Science*, 318(5853): 1155, 2007.

H. Hamann and H. Wörn. An analytical and spatial model of foraging in a swarm of robots. In E. Şahin, W. M. Spears, and A. F. T. Winfield, editors, *Swarm Robotics - Second SAB 2006 International Workshop*, volume 4433 of *LNCS*, pages 43–55. Springer, Berlin, Germany, 2007.

H. Hamann and H. Wörn. Aggregating robots compute: An adaptive heuristic for the euclidean steiner tree problem. In M. Asada, J. Hallam, J.-A. Meyer, and J. Tani, editors, *From Animals to Animats 10 (SAB 2008)*, volume 5040 of *LNCS*, pages 447–456. Springer, Berlin, Germany, 2008a.

H. Hamann and H. Wörn. A framework of spacetime continuous models foralgorithm design in swarm robotics. *Swarm Intelligence*, 2(2-4):209–239, 2008b.

H. Hamann, H. Wörn, K. Crailsheim, and T. Schmickl. Spatial macroscopic models of a bio-inspired robotic swarm algorithm. In *Proceedings of the IEEE/RSJ 2008 International Conference on Intelligent Robots and Systems (IROS)*, pages 1415–1420. IEEE Press, Los Alamitos, CA, USA, 2008.

H. Hamann, G. Valentini, Y. Khaluf, and M. Dorigo. Derivation of a micro-macro link for collective decision-making systems: Uncover network features based on drift measurements. In T. Bartz-Beielstein, J. Branke, B. Filipič, and J. Smith, editors, *Parallel Problem Solving from Nature – PPSN XIII*, volume 8672 of *LNCS*, pages 181–190. Springer, 2014.

N. Hatziargyriou, H. Asano, R. Iravani, and C. Marnay. Microgrids. *IEEE Power and Energy Magazine*, 5(4):78–94, 2007.

S. Hauert and S. N. Bhatia. Mechanisms of cooperation in cancer nanomedicine: towards systems nanotechnology. *Trends in Biotechnology*, 32(9):448–455, 2014.

S. Hauert, S. Berman, R. Nagpal, and S. N. Bhatia. A computational framework for identifying design guidelines to increase the penetration of targeted nanoparticles into tumors. *Nano Today*, 8(6):566 – 576, 2013.

S. Haykin. Cognitive radio: brain-empowered wireless communications. *IEEE Journal on Selected Areas in Communications*, 23(2):201–220, 2005.

J. P. Hecker and M. E. Moses. Beyond pheromones: evolving error-tolerant, flexible, and scalable ant-inspired robot swarms. *Swarm Intelligence*, 9(1): 43–70, 2015.

D. Helbing, editor. *Social Self-Organization*. Understanding Complex Systems. Springer Verlag, Berlin, Germany, 2012.

D. Helbing. Globally networked risks and how to respond. *Nature*, 497 (7447):51–59, 2014.

D. Helbing and E. Pournaras. Build digital democracy. *Nature*, 527:33–34, 2015.

M. A. Hsieh, A. Halász, S. Berman, and V. Kumar. Biologically inspired redistribution of a swarm of robots among multiple sites. *Swarm Intelligence*, 2(2-4):121–141, 2008.

K. Huang, T. Wang, Y. Cheng, and X. Zheng. Effect of heterogeneous investments on the evolution of cooperation in spatial public goods game. *PLoS ONE*, 10(3):e0120317–10, 2015.

M. D. Ilic, L. Xie, U. A. Khan, and J. M. F. Moura. Modeling of future cyber-physical energy systems for distributed sensing and control. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, 40(4):825–838, 2010.

J. Jandt, S. Bengston, N. Pinter-Wollman, J. Pruitt, N. Raine, A. Dornhaus, and A. Sih. Behavioural syndromes and social insects: personality at multiple levels. *Biological Reviews*, 89(1):48–67, 2014.

S. Janson, M. Middendorf, and M. Beekman. Searching for a new home–scouting behavior of honeybee swarms. *Behavioral Ecology*, 18(2):384–392, 2007.

R. Jeanson, S. Blanco, R. Fournier, J.-L. Deneubourg, V. Fourcassié, and G. Theraulaz. A model of animal movements in a bounded space. *Journal of Theoretical Biology*, 225(4):443–451, 2003.

R. Jeanson and A. Weidenmüller. Interindividual variability in social insects - proximate causes and ultimate consequences. *Biological Reviews*, 89(3):671–687, 2013.

M. Ji and M. Egerstedt. Distributed coordination control of multiagent systems while preserving connectedness. *IEEE Transactions on Robotics*, 23(4): 693–703, 2007.

D. Jones, S. Stewart, and L. Power. Patterns: using proven experience to develop online learning. In *Proceedings of ASCILITE*, volume 99, pages 155–162, 1999.

S. Kalantar and U. R. Zimmer. Distributed shape control of homogeneous swarms of autonomous underwater vehicles. *Autonomous Robots*, 22(1): 37–53, 2006.

A. B. Kao, N. Miller, C. Torney, A. Hartnett, and I. D. Couzin. Collective Learning and Optimal Consensus Decisions in Social Animal Groups. *PLoS Computational Biology*, 10(8):e1003762, 2014.

M. D. Kennedy, L. Guerrero, and V. Kumar. Decentralized algorithm for force distribution with applications to cooperative transport. In *the 39th Mechanisms and Robotics Conference (ASME 2015)*, volume 5C, page V05CT08A013, New York, NY, USA, 2015. American Society of Mechanical Engineers.

Y. Khaluf and M. Dorigo. Modeling robot swarms using integrals of birth-death processes. *ACM Transactions on Autonomous and Adaptive Systems*, In press, 2016.

T. Laomettachit, T. Termsaithong, A. Sae-Tang, and O. Duangphakdee. Decision-making in honeybee swarms based on quality and distance information of candidate nest sites. *Journal of Theoretical Biology*, 364:21–30, 2015.

D. Latella, M. Loreti, and M. Massink. On-the-fly PCTL fast mean-field approximated model-checking for self-organising coordination. *Science of Computer Programming*, 110:23 – 50, 2015.

J. R. T. Lawton, R. W. Beard, and B. J. Young. A decentralized approach to formation maneuvers. *IEEE Transactions on Robotics and Automation*, 19(6): 933–941, 2003.

E. Lee, B. Hartmann, J. Kubiatowicz, T. Simunic Rosing, J. Wawrzynek, D. Wessel, J. Rabaey, K. Pister, A. Sangiovanni Vincentelli, S. Seshia, D. Blaauw, P. Dutta, K. Fu, C. Guestrin, B. Taskar, R. Jafari, D. Jones, V. Kumar, R. Mangharam, G. Pappas, R. Murray, and A. Rowe. The swarm at the edge of the cloud. *IEEE Design & Test*, 31(3):8–20, 2014.

E. A. Lee. The past, present and future of cyber-physical systems: A focus on models. *Sensors*, 15(3):4837, 2015.

N. E. Leonard, D. A. Paley, F. Lekien, R. Sepulchre, D. M. Fratantoni, and R. E. Davis. Collective motion, sensor networks, and ocean sampling. *Proceedings of the IEEE*, 95(1):48–74, 2007.

N. E. Leonard and E. Fiorelli. Virtual leaders, artificial potentials and coordinated control of groups. In *Proceedings of the 40th IEEE Conference on Decision and Control*, volume 3, pages 2968–2973, 2001.

K. Lerman, A. Martinoli, and A. Galstyan. A review of probabilistic macroscopic models for swarm robotic systems. In E. Şahin and W. Spears, editors, *Swarm Robotics Workshop: State-of-the-art Survey*, number 3342 in LNCS, pages 143–152. Springer, Berlin, Germany, 2005.

K. Lerman, C. V. Jones, A. Galstyan, and M. J. Matarić. Analysis of dynamic task allocation in multi-robot systems. *International Journal of Robotics Research*, 25(3):225–242, 2006.

Z. Liang, T. Nguyen, H. Mattila, S. Rodriguez-Zas, T. Seeley, and G. Robinson. Molecular determinants of scouting behavior in honey bees. *Science*, 335(6073):1225–1228, 2012.

Y.-Y. Liu, J.-J. Slotine, and A.-L. Barabási. Controllability of complex networks. *Nature*, 473(7346):167–173, 2011.

R. Ma, H. H. Chen, Y. R. Huang, and W. Meng. Smart grid communication: Its challenges and opportunities. *IEEE Transactions on Smart Grid*, 4(1): 36–46, 2013.

A. Mailleux, C. Detrain, and J.-L. Deneubourg. Starvation drives a threshold triggering communication. *The Journal of Experimental Biology*, 209(21): 4224–4229, 2006.

J. Marshall, R. Bogacz, A. Dornhaus, R. Planqué, T. Kovacs, and N. Franks. On optimal decision-making in brains and social insect colonies. *Journal of the Royal Society, Interface*, 6(40):1065–1074, 2009.

A. D. Martino and M. Marsili. Statistical mechanics of socio-economic systems with heterogeneous agents. *Journal of Physics A: Mathematical and General*, 39(43):R465–R540, 2006.

A. Martinoli, K. Easton, and W. Agassounon. Modeling Swarm Robotic Systems: A Case Study in Collaborative Distributed Manipulation. *International Journal of Robotics Research*, 23(4):415–436, 2004. Special Issue on Experimental Robotics, B. Siciliano, editor.

M. Massink, M. Brambilla, D. Latella, M. Dorigo, and M. Birattari. On the use of Bio-PEPA for modelling and analysing collective behaviours in swarm robotics. *Swarm Intelligence*, 7(2–3):201–228, 2013.

H. Mattila and T. Seeley. Genetic diversity in honey bee colonies enhances productivity and fitness. *Science*, 317(5836):362–364, 2007.

J.-M. McNew, E. Klavins, and M. Egerstedt. Solving coverage problems with embedded graph grammars. In A. Bemporad, A. Bicchi, and G. Buttazzo, editors, *Hybrid Systems: Computation and Control*, volume 4416 of *LNCS*, pages 413–427. Springer, Berlin, Germany, 2007.

G. Mermoud, U. Upadhyay, W. C. Evans, and A. Martinoli. Top-down vs. bottom-up model-based methodologies for distributed control: Acomparative experimental study. In O. Khatib, V. Kumar, and G. Sukhatme, editors, *Experimental Robotics*, volume 79 of *STAR*, pages 615–629. Springer, Berlin, Germany, 2014.

N. Michael and V. Kumar. Planning and control of ensembles of robots with non-holonomic constraints. *International Journal of Robotics Research*, 28(8): 962–975, 2009.

F. Mondada, M. Bonani, X. Raemy, J. Pugh, C. Cianci, A. Klaptocz, S. Magnenat, J.-C. Zufferey, D. Floreano, and A. Martinoli. The e-puck, a robot designed for education in engineering. In *Proceedings of the 9th conference on autonomous robot systems and competitions*, volume 1(1), pages 59–65. IPCB, Castelo Branco, Portugal, 2009.

M. Montes, 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, 2010.

P. Moretti, S. Liu, A. Baronchelli, and R. Pastor-Satorras. Heterogenous mean-field analysis of a generalized voter-like model on networks. *The European Physical Journal B*, 85(3):1–6, 2012.

R. Nagpal. A catalog of biologically-inspired primitives for engineering self-organization. In G. Di Marzo Serugendo, A. Karageorgos, O. Rana, and F. Zambonelli, editors, *Engineering Self-Organising Systems*, volume 2977 of *LNCS*, pages 53–62. Springer, Berlin, Germany, 2004.

S. Nakamoto. Bitcoin: A peer-to-peer electronic cash system, 2008.

W. Nelson. Hazard plotting for incomplete failure data. *Journal of Quality Technology*, 1:27–52, 1969.

P. Ögren, M. Egerstedt, and X. Hu. A control lyapunov function approach to multi-agent coordination. In *Proceedings of the 40th IEEE Conference on Decision and Control*, volume 2, pages 1150–1155. IEEE Press, 2001.

P. Ögren, E. Fiorelli, and N. E. Leonard. Cooperative control of mobile sensor networks:adaptive gradient climbing in a distributed environment. *IEEE Transactions on Automatic Control*, 49(8):1292–1302, 2004.

D. E. Olivares, A. Mehrizi-Sani, A. H. Etemadi, C. A. C. nizares, R. Iravani, M. Kazerani, A. H. Hajimiragha, O. Gomis-Bellmunt, M. Saeedifard, R. Palma-Behnke, G. A. Jiménez-Estévez, and N. D. Hatziargyriou. Trends in microgrid control. *IEEE Transactions on Smart Grid*, 5(4):1905–1919, 2014.

D. Pais, P. Hogan, T. Schlegel, N. Franks, N. Leonard, and J. Marshall. A mechanism for value-sensitive decision-making. *PLoS ONE*, 8(9):e73216, 2013.

L. Panait and S. Luke. Cooperative multi-agent learning: The state of the art. *Autonomous Agents and Multi-Agent Systems*, 11(3):387–434, 2005.

C. A. C. Parker and H. Zhang. Cooperative decision-making in decentralized multiple-robot systems: the best-of-N problem. *IEEE Transactions on Mechatronics*, 14(2):240–251, 2009.

C. A. C. Parker and H. Zhang. Collective unary decision-making by decentralized multiple-robot systems applied to the task-sequencing problem. *Swarm Intelligence*, 4(3):199–220, 2010.

K. Passino and T. Seeley. Modeling and analysis of nest-site selection by honeybee swarms: the speed and accuracy trade-off. *Behavioral Ecology and Sociobiology*, 59(3):427–442, 2006.

C. Pinciroli, V. Trianni, R. O'Grady, G. Pini, A. Brutschy, M. Brambilla, N. Mathews, E. Ferrante, G. A. Di Caro, F. Ducatelle, M. Birattari, L. M. Gambardella, and M. Dorigo. ARGoS: A modular, parallel, multi-engine simulator for multi-robot systems. *Swarm Intelligence*, 6(4):271–295, 2012.

I. Planas-Sitjà, J.-L. Deneubourg, C. Gibon, and G. Sempo. Group personality during collective decision-making: a multi-level approach. *Proceedings of the Royal Society of London B: Biological Sciences*, 282(1802):20142515, 2015.

A. Prorok, N. Corell, and A. Martinoli. Multi-level spatial modeling for stochastic distributed robotic systems. *The International Journal of Robotics Research*, 30(5):574–589, 2011.

R. Rajkumar, I. Lee, L. Sha, and J. Stankovic. Cyber-physical systems: The next computing revolution. In *Proceedings of the 47th Design Automation Conference*, DAC '10, pages 731–736, New York, NY, USA, 2010. ACM.

A. Reina and V. Trianni. Deployment and redeployment of wireless sensor networks: a swarm robotics perspective. In N. Mitton and D. Simplot-Ryl, editors, *Wireless Sensor and Robot Networks*, chapter 7, pages 143–162. World Scientific, Singapore, 2014.

A. Reina, M. Dorigo, and V. Trianni. Collective decision making in distributed systems inspired by honeybees behaviour. In *Proceedings of 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1421–1422. IFAAMAS, New York, 2014a.

A. Reina, M. Dorigo, and V. Trianni. Towards a cognitive design pattern for collective decision-making. In M. Dorigo et al., editor, *Swarm Intelligence (ANTS 2014)*, volume 8667 of *LNCS*, pages 194–205. Springer, Berlin, Germany, 2014b.

A. Reina, L. M. Gambardella, M. Dorigo, and G. A. Di Caro. zePPeLIN: Distributed path planning using an overhead camera network. *International Journal of Advanced Robotic Systems*, 11(119), 2014c.

A. Reina, R. Miletitch, M. Dorigo, and V. Trianni. A quantitative micro-macro link for collective decisions: The shortest path discovery/selection example. *Swarm Intelligence*, 9(2–3):75–102, 2015a.

A. Reina, M. Salvaro, G. Francesca, L. Garattoni, C. Pinciroli, M. Dorigo, and M. Birattari. Augmented reality for robots: virtual sensing technology applied to a swarm of e-pucks. In *Proceedings of the IEEE 2015 NASA/ESA Conference of Adaptive Hardware and Systems (AHS)*, pages 1–6, Los Alamitos, CA, 2015b. IEEE Computer Society Press. Paper ID sB_p3.

A. Reina, G. Valentini, C. Fernández Oto, M. Dorigo, and V. Trianni. A design pattern for best-of-n collective decisions. Extended abstract presented at the Biologically Distributed Algorithms Workshop (BDA 2015), 2015c.

A. Reina, G. Valentini, C. Fernández-Oto, M. Dorigo, and V. Trianni. A design pattern for decentralised decision making. *PLoS ONE*, 10(10): e0140950, 2015d.

J. F. Roberts, T. S. Stirling, J.-C. Zufferey, and D. Floreano. 2.5D infrared range and bearing system for collective robotics. In *Proceeding of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3659–3664. IEEE Press, 2009.

E. Robinson, N. Franks, S. Ellis, S. Okuda, and J. Marshall. A simple threshold rule is sufficient to explain sophisticated collective decision-making. *PLoS ONE*, 6(5):e19981, 2011.

M. Rohden, A. Sorge, M. Timme, and D. Witthaut. Self-organized synchronization in decentralized power grids. *Physical Review Letters*, 109:064101, Aug 2012.

M. Rubenstein, C. Ahler, N. Hoff, A. Cabrera, and R. Nagpal. Kilobot: A low cost robot with scalable operations designed for collective behaviors. *Robotics and Autonomous Systems*, 62(7):966–975, 2014.

G. Sartoretti, M.-O. Hongler, M. E. de Oliveira, and F. Mondada. Decentralized self-selection of swarm trajectories: from dynamical systems theory to robotic implementation. *Swarm Intelligence*, 8(4):329–351, 2014.

T. Schaerf, J. Makinson, M. Myerscough, and M. Beekman. Do small swarms have an advantage when house hunting? The effect of swarm size on nest-site selection by Apis mellifera. *Journal of The Royal Society, Interface*, 10 (87):20130533, 2013.

A. Scheidler, A. Brutschy, E. Ferrante, and M. Dorigo. The k-Unanimity rule for self-organized decision making in swarms of robots. *IEEE Transactions on Cybernetics*, PP(99):1, 2015.

M. Schillo, K. Fischer, and C. T. Klein. The Micro-Macro Link in DAI and Sociology. In S. Moss and P. Davidsson, editors, *Multi-Agent-Based Simulation*, volume 1979 of *LNCS*, pages 133–148. Springer, Berlin, Germany, 2001.

T. Schmickl, H. Hamann, H. Wörn, and K. Crailsheim. Two different approaches to a macroscopic model of a bio-inspired robotic swarm. *Robotics and Autonomous Systems*, 57(9):913–921, 2009.

T. Seeley, P. Visscher, T. Schlegel, P. Hogan, N. Franks, and J. Marshall. Stop signals provide cross inhibition in collective decision-making by honeybee swarms. *Science*, 335(6064):108–11, 2012.

R. Sepulchre, D. A. Paley, and N. E. Leonard. Stabilization of planar collective motion: All-to-all communication. *IEEE Transactions on Automatic Control*, 52(5):811–824, 2007.

R. Sepulchre, D. A. Paley, and N. E. Leonard. Stabilization of planar collective motion with limited communication. *IEEE Transactions on Automatic Control*, 53(3):706–719, 2008.

B. Smith, A. Howard, J.-M. McNew, J. Wang, and M. Egerstedt. Multi-robot deployment and coordination with embedded graph grammars. *Autonomous Robots*, 26(1):79–98, 2009.

T. Soleymani, V. Trianni, M. Bonani, F. Mondada, and M. Dorigo. Bio-inspired construction with mobile robots and compliant pockets. *Robotics and Autonomous Systems*, pages 1–24, 2015.

V. Sood and S. Redner. Voter model on heterogeneous graphs. *Physical review letters*, 94(17):178701–4, 2005.

V. Srivastava and N. E. Leonard. On the speed-accuracy tradeoff in collective decision making. In *Proceedings of the IEEE 52nd Annual Conference on Decision and Control (CDC)*, pages 1880–1885, 2013.

V. Srivastava and N. E. Leonard. Collective decision-making in ideal networks: The speed-accuracy tradeoff. *IEEE Transactions on Control of Network Systems*, 1(1):121–132, 2014.

V. Srivastava and N. E. Leonard. On first passage time problems in collective decision-making with heterogeneous agents. In *American Control Conference (ACC)*, pages 2113–2118, 2015.

A. Stranieri, A. E. Turgut, M. Salvaro, G. Francesca, A. Reina, M. Dorigo, and M. Birattari. Iridia's arena tracking system. Technical Report TR/IRIDIA/2013-013, IRIDIA, Université Libre de Bruxelles, Brussels, Belgium, July 2014. revision r003.

G. Theraulaz, E. Bonabeau, and J.-L. Denuebourg. Response threshold reinforcements and division of labour in insect societies. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 265(1393):327–332, 1998.

V. Trianni. *Evolutionary Swarm Robotics - Evolving Self-Organising Behaviours in Groups of Autonomous Robots*. Studies in Computational Intelligence. Springer, Berlin, Germany, 2008.

V. Trianni, A. Cacciapuoti, and M. Caleffi. Distributed design for fair coexistence in TVWS. In *Proceedings of the IEEE International Conference on Communications (ICC 2016)*. IEEE Press, 2016a.

V. Trianni, D. De Simone, A. Reina, and A. Baronchelli. Emergence of consensus in a multi-robot network: from abstract models to empirical validation. *Robotics and Automation Letters, IEEE*, PP(99):1–1, 2016b.

W. F. Truszkowski, M. G. Hinchey, J. L. Rash, and C. A. Rouff. Autonomous and autonomic systems: a paradigm for future space exploration missions. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 36(3):279–291, 2006.

G. Valentini, H. Hamann, and M. Dorigo. Self-organized collective decision making: The weighted voter model. In A. Lomuscio, P. Scerri, A. Bazzan, and M. Huhns, editors, *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems*, AAMAS 2014, pages 45–52. IFAAMAS, 2014.

G. Valentini, E. Ferrante, H. Hamann, and M. Dorigo. Collective decision with 100 kilobots: Speed versus accuracy in binary discrimination problems. *Autonomous Agents and Multi-Agent Systems*, 30(3):553–580, 2016.

T. Vicsek and A. Zafeiris. Collective motion. *Physics Reports*, 517(3–4):71–140, 2012.

M. Vigelius, B. Meyer, and G. Pascoe. Multiscale modelling and analysis of collective decision making in swarm robotics. *PLoS ONE*, 9(11):e111542, 2014.

M. Viroli and F. Damiani. A calculus of self-stabilising computational fields. In E. Kühn and R. Pugliese, editors, *Coordination Models and Languages*, volume 8459 of *LNCS*, pages 163–178. Springer, Berlin, Germany, 2014.

A. Weidenmüller. The control of nest climate in bumblebee (bombus terrestris) colonies: interindividual variability and self reinforcement in fanning response. *Behavioral Ecology*, 15(1):120–128, 2004.

C. Westhus, C. Kleineidam, F. Roces, and A. Weidenmueller. Behavioural plasticity in the fanning response of bumblebee workers: impact of experience and rate of temperature change. *Animal Behaviour*, 85(1):27–34, 2013.

S. Wilson, T. P. Pavlic, G. P. Kumar, A. Buffin, S. C. Pratt, and S. Berman. Design of ant-inspired stochastic control policies for collective transport by robotic swarms. *Swarm Intelligence*, 8(4):303–327, 2014.

A. F. Winfield, W. Liu, J. Nembrini, and A. Martinoli. Modelling a wireless connected swarm of mobile robots. *Swarm Intelligence*, 2(2-4):241–266, 2008.

M. Wray and T. Seeley. Consistent personality differences in house-hunting behavior but not decision speed in swarms of honey bees (*Apis mellifera*). *Behavioral Ecology and Sociobiology*, 65(11):2061–2070, 2011.

APPENDICES

A

SURVIVAL ANALYSIS

In Chapter 4, we estimate the transition rates of the macroscopic model from the multiagent/swarm robotics simulations through survival analysis. In this appendix, we shortly introduce this methodology and we present the Nelson-Aalen estimator that we employ for our analysis. Finally, we describe how to estimate the transition rates directly from experimental data. We believe that showing the application of this statistical tool for the analysis of a swarm robotics experiment may be of interest to the community.

SURVIVAL ANALYSIS Survival analysis is a branch of statistics that offers tools to estimate the change over time of the probability of an event from experimental data. Survival analysis has been initially introduced in medicine to estimate the probability of survival (or death) of an organism under some treatment. Subsequently, these tools generalised to the estimate of any transition probability between populations. Nowadays, survival analysis is employed in several fields, such as economics—e.g., to estimate the probability of a stock market crash—or mechanical engineering—e.g., to estimate the probability of engine failures. In this work, we apply survival analysis to estimate the transition rates of the macroscopic model of decentralised decision making.

Other works have used this type of analysis to estimate the parameters of multiagent systems behaviour. Jeanson et al. (2003) use survival analysis to estimate the probability with which cockroaches change their behaviour. Garnier et al. (2008) and Reina et al. (2014b) employed survival analysis to compute transition rates of artificial agents behaviour.

HAZARD CURVE We consider the three populations $\{A, B, U\}$ in accordance with the agent's commitment state described in Section 4.1. To compute the rate at which agents switch (*transit* between) their commitment state, we log the number of timesteps t interlaying between two commitment switches (transitions) and the relative type of event causing the switch

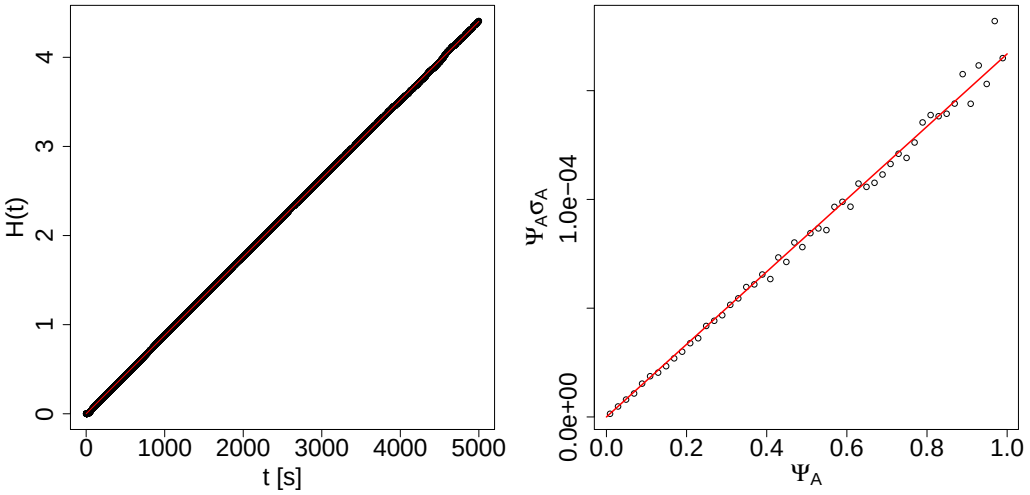

(A) (B)

Figure 21: Estimation of the macroscopic transition rates from experimental data through survival analysis in two exemplifying cases taken from the shortest path discovery/selection case study of Chapter 4. (A) Cumulative discovery probability over time for target area at distance $d = 1.5\,\text{m}$. The slope of the resulting line corresponds to the discovery rate estimate. (B) Cross-inhibition rate σ_A estimates over population fraction Ψ_A for target area distances $d_A = d_B = 2.5\,\text{m}$ and probability $P_\sigma = 0.1$.

(e.g., discovery or recruitment). At the end of an experiment, we log the timesteps t from the last commitment switch as censored event, which indicates that after t timesteps no transition happened. We use the Nelson-Aalen estimator (Nelson, 1969) to compute the hazard curve $H(t)$ from the collected experimental data. The hazard curve $H(t)$ shows the cumulative probability of events occurring until time t, and is computed as follows:

$$H(t) = \sum_{t_i \leq t} d_i / n_i, \tag{38}$$

where d_i is the number of events recorded at t_i, and n_i is the number of events occurring (or censored) at time $t \geq t_i$. In a memory-less system, the probability of an event does not change over time, therefore the curve of the cumulative probability as function of time corresponds to a line with a slope equal to the constant event rate (i.e., the transition rate). Assuming our system as memory-less, we compute the transition rate by linear fitting the hazard curve with a line passing through the origin. Additionally, the quality of the fitting can be used to verify the correctness of a memoryless implementation.

RATE ESTIMATION The rate of spontaneous transitions, in this study discovery and abandonment, can be directly estimated by computing the slope of the hazard curve, as detailed above. For instance, Figure 21(A) shows the hazard curve of the discovery rate computed from the multiagent experiments of Chapter 4 with target area at distance $d = 1.5\,\mathrm{m}$. Differently, rate estimates of transitions consequent to an interaction—in this work, recruitment and cross-inhibition— include the probability of interaction with an agent of the other population, which we call hereafter the *interacting population*. This probability changes during the process as the interacting population size changes, and must be taken into account to estimate a constant transition rate independently from the interacting population size. Therefore, we first compute one aggregate transition rate for every interacting population size, and then we normalise the rates for the interacting population fraction. For instance, from Equation (1), the cross-inhibition rate for the population B is $(-\sigma_A \Psi_A)$, which includes the size of the interacting population A that delivers the inhibition signal. Through survival analysis, we compute the aggregate rate $(\sigma_A \Psi_A)$ for varying values of Ψ_A in the range $]0, 1[$. Then, by linear fitting, we discount from the aggregate rates the interacting population fraction to obtain σ_A. Figure 21(B) shows a set of 50 aggregate rate estimates plotted as function of Ψ_A, the slope of the fitted line corresponds to the estimate of σ_A.

To estimate the transition rates with constant population sizes, we run adhoc experiments where the population sizes are fixed. In these experiments, agents follow the normal behaviour and, in case of commitment state transitions, they only log the event but do not change commitment state. In this way, we can quickly gather a large amount of data for every population size and parameterisation.

B

STABILITY ANALYSIS AND PARAMETERISATION CHOICE FOR CASE STUDY I-A

To determine the parameters of case study I-A, we analyse the macroscopic ODE model. To meet the requirement of consensus decision, the model must display equilibria for the points $S_A = \{\Psi_A = 1, \Psi_B = 0\}$ and $S_B = \{\Psi_A = 0, \Psi_B = 1\}$. The system has equilibria S_A and S_B for $\alpha_A = \alpha_B = 0$. We choose $\sigma_A = \sigma_B = 1$, and we perform a stability analysis of the system as a function of the parameters γ_i, ρ_i, $i \in \{A, B\}$. Given the domain space:

$$\begin{cases} 0 \leq \Psi_i \leq 1, \quad i \in \{A, B\} \\ 0 \leq \Psi_U \leq 1 \\ \Psi_A + \Psi_B + \Psi_U = 1 \end{cases}, \tag{39}$$

the analysis of the macroscopic ODE model reveals three possible equilibria: S_A, S_B and S_X. Equilibria S_A and S_B do not change position and are always present for any parameterisation but change stability. In contrast, equilibrium S_X is always unstable and appears only when S_A and S_B are both stable, as a function of the parameters γ_i and ρ_i. The stability analysis gives:

$$\begin{array}{rcl} S_A \text{ is stable} & \Leftrightarrow & \gamma_B < \gamma_A + \rho_A \\ S_B \text{ is stable} & \Leftrightarrow & \gamma_A < \gamma_B + \rho_B \end{array} \tag{40}$$

The equilibria S_A and S_B change stability through transcritical bifurcations, and when one of the consensus solutions is unstable, the other one is the unique stable solution.

Given Eq. (40), we can parameterise the system to minimise the chance of wrong decisions. We require that there exists only one stable solution for a quality difference above a target resolution R:

$$|v_A - v_B| / \max(v_A, v_B) > R \tag{41}$$

Let us assume $v_A > v_B$. On the one hand S_A must be the only stable solution, and therefore S_B should be unstable, hence $\gamma_A > \gamma_B + \rho_B$. On the other

hand, the constraint on the target resolution implies that $v_A - v_B/v_A > R$. We select linear functions that link macroscopic transition rates to quality:

$$\gamma_i = f_\gamma(v_i) = k\, v_i \quad \rho_i = f_\rho(v_i) = h\, v_i, \quad i \in \{\mathrm{A}, \mathrm{B}\} \tag{42}$$

where k and h are tuneable parameters. To satisfy our design choices, we combine the above equations and solve the system:

$$\begin{cases} k\, v_A > k\, v_B + h\, v_B \\ v_A - v_B > v_A R \end{cases} \quad \rightarrow \quad k > h\,(1-R)/R \tag{43}$$

An identical relation can be obtained assuming $v_B > v_A$. We arbitrarily select the target resolution $R = 0.15$, and we finally select a parameterisation that complies with the prescribed bounds: $h = 0.1$ and $k = 0.6$.