A Design Pattern for Swarm-Centric Decision Making

Vito Trianni
ISTC-CNR
Via San Martino della Battaglia 44, 00185 Roma
ITALY

vito.trianni@istc.cnr.it

ABSTRACT

The engineering of large-scale decentralised systems requires sound methodologies to guarantee the attainment of the desired macroscopic system-level behaviour given the microscopic individual-level implementation. While a general-purpose methodology is currently out of reach, specific solutions can be given to broad classes of problems by means of well-conceived design patterns, which recall the well-known concept exploited in software engineering. Design patterns provide formal guidelines to deal with recurring problems in a specific domain. For the particular case of distributed systems, design patterns should prescribe the individual-level microscopic behaviour required to obtain desired system-level macroscopic properties. Here, we present a design pattern for decentralised decision-making—a fundamental ability in several contexts and application domains. The design pattern is based on the nest-site selection behaviour of honeybee swarms. Previous experimental and theoretical studies have demonstrated near-optimal speed-accuracy trade-offs in the selection of the most profitable option among a set of alternative nesting sites by honeybees. 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. The above properties of the nest-site selection process are relevant for many practical decision making scenarios in decentralised systems. Starting from the macroscopic description of the nest-site selection dynamics, we derived the exact relationship between microscopic and macroscopic models—also including finite-size effects—for the general case of a best-of-n decision problem. The inter-related models represent the core of the design pattern, which is completed by formal guidelines for the implementation of collective decisions in multi-agent systems. We report case studies with multi-agent systems, robot swarms and cognitive radio networks to illustrate the application of the design pattern across application domains.

1.0 INTRODUCTION

Decision-making is a crucial ability for autonomous agents immersed in an uncertain, dynamic environment [1][10][13]. In biological systems, the ability to select the most profitable option among the available alternatives is a fundamental adaptive response that can determine the organism survival. This is why decision-making strategies observed in natural systems often present an optimal speed-accuracy trade-off: given a desired accuracy level, a decision is taken in the shortest time possible [2]. Recent studies have identified optimal decision strategies also in decentralised systems, and have recognised interesting parallels between the observed collective behaviours and neural models of optimal decision-making [5]. Speaking in general terms, these studies describe swarm-like systems as information-processing systems capable of some cognitive ability [11]. The observed similarities between swarms and brains reveal the existence of cognitive abilities in distributed systems, which are strongly determined by the interaction patterns among the system components. In other words, cognition in collectives can be recognised as the ability of a system to acquire and process information through the numerous interactions among the system components, therefore limiting the need for explicit representations within the single units.

In this paper, we propose to take a similar perspective in the design of artificial, large-scale distributed systems. By viewing such systems as distributed cognitive systems, it will be possible to maximise their information processing capability while keeping a low complexity of the individual units. That is, individual units would contribute to the overall system behaviour without having the global picture about the cognitive process they are collectively producing. Designing such an information-processing system is clearly a complex endeavour, and it can be successful only if grounded on solid theoretical premises. Therefore, we propose a design methodology that leverages the current understanding of cognitive processing in (natural) distributed systems, and that puts this knowledge in use for the design of artificial ones [8]. Our proposal is based on the concept of *design patterns*, that is, reusable solutions to tackle complex design problems [3]. Similarly to common practice in software engineering, design patterns can be used to guide the design and development of decentralized, swarm-centric systems, independently of the particular implementation technique.

In particular, we present a design pattern for decentralised decision-making based on the nest-site selection behaviour of honeybee swarms [6][9]. Previous experimental and theoretical studies have demonstrated near-optimal speed-accuracy trade-offs in the selection of the most profitable option among a set of alternative nesting sites by honeybees [6]. 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 [9]. The above properties of the nest-site selection process are relevant for many practical decision-making scenarios in decentralised systems. Starting from the macroscopic description of the nest-site selection dynamics, we derived the exact relationship between microscopic and macroscopic models|—also including finite-size effects—for the general case of a best-of-n decision problem [8]. The inter-related models represent the core of the design pattern, which is completed by formal guidelines for the implementation of collective decisions in multi-agent systems. We provided guidelines for implementation by means of either homogenous or heterogeneous agents, as well as guidelines for the inclusion of spatial and topological factors that have a bearing in determining the microscopic interaction patterns [8]. We report here some case studies that illustrate the application of the design pattern, as well as the relevance of the obtained results for better understanding the behaviour of natural systems.

2.0 COLLECTIVE DECISIONS THROUGH CROSS-INHIBITION

Models of collective decision-making have been studied in different domains, from ethology to social dynamics. Our proposed design methodology starts from a model of nest site selection in honeybee swarms [9]. In honeybee colonies, 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, get committed to it. Through 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 through this form of positive feedback a greater number of bees get committed to the best quality nests. Eventually, the entire colony reaches consensus on a single site where to swarm. Recently, it has been discovered that bees committed to different alternatives cross-inhibit each other through stop signals. A bee committed for a site that receives several stop signals abandons its commitment and becomes uncommitted. This mechanism allows to break decision deadlocks in case of equal-best alternatives. In this way, the colony minimizes the decision time, thus exposure to dangers. This process, based on peer-to-peer interactions, lets the colony quickly converge towards the highest quality nests without the need of quality comparisons. These advantageous characteristics and low requirements in terms of agent capabilities allow to apply the model to a large number of possible applications.

2.1 From macroscopic to microscopic descriptions

By abstracting from the particular case of nest-site selection, it is possible to introduce an elegant model of the collective decision making among n possible alternatives [8]. The macroscopic dynamics can be modelled at the population level through a system of differential equations that describe the evolution through time of the different sub-populations, each representing the fraction of the group that is uncommitted or committed to any available alternative. Agents become committed to a given alternative either through spontaneous discovery, which can happen at a fixed rate depending on the option quality, or through interactions with other agents that actively recruit them (at a rate proportional to the size of the population already committed to the given alternative). Similarly, committed agents may become uncommitted either through spontaneous abandonment (at a constant rate inversely proportional to the option quality), or through cross-inhibition (at a rate proportional to the fraction of the population committed to a different option). The study of the macroscopic dynamics resulting from such a model reveals that an optimal decision is always taken if the rate of cross-inhibition among agents is sufficiently high, at least for the binary case [6][9]. However, the macroscopic models only suggest a possible implementation for the individual behaviour, in analogy with the behaviour observed in honeybees, but this is normally not sufficient for the implementation of an artificial system.

In [8], we provided a detailed description of the microscopic rules that need to be followed by individual agents so that the desired macroscopic dynamics can be achieved. We have shown that the average agent behaviour should be modelled as a Probabilistic Finite States Machine (PFSM), as depicted in Figure 1. Here, an agent can be in $n+1$ different states: either uncommitted, or committed to any of the n available options. Transitions between states are determined by probabilities which are either constant and depend on the quality of the given alternative, or are modulated by the current fraction of agents in the population committed to a particular alternative. In any case, only transitions between uncommitted and committed states are allowed, forcing every change of commitment to go through the uncommitted state. More in detail, we can identify two different transition types. *Spontaneous* transitions correspond to discovery and abandonment of a given alternative (see bold lines in Figure 1). These transitions happen with a constant probability (in average), which depends on the objective quality of the option (as perceived with noise from the individual agent). *Interactive* transitions instead represent the change in commitment state of an agent on the basis of the current fraction of agents committed to the different available options. An uncommitted agent becomes committed with a higher probability to options that are followed by a large fraction of agents. Conversely, a committed agent becomes uncommitted with a probability that is modulated by the fraction of agents committed to any other option different from the one of the focal agent (see Figure 1). These transitions are referred to as interactive because in a decentralised systems interactions among agents are necessary to estimate the current fraction of agents in each sub-population, which is a pre-requisite to compute the transition probabilities. The way in which such transition probabilities are computed by each agent, as well as the way in which they are linked to the perceived option quality, determine the actual implementation of the design pattern. In the following, we provide guidelines to obtain a working implementation that guarantees to obtain the desired macroscopic dynamics.

2.2 Implementation guidelines

The actual implementation of the agent behaviour requires choosing the way in which transitions are executed in relation to the limited information available to the individual agent. For instance, the estimation of the population-size dependent probability by individual agents requires some sampling of the current population size. The other transition probabilities should instead provide an exact relationship with the macroscopic transition rates to obtain the desired dynamics.

Population-size dependent probabilities can be estimated by looking at the commitment state of neighbours. Through communication, neighbouring agents can share their commitment state so that an estimate of the current population size can be obtained. In a well-mixed system, it is sufficient that each agent interacts with only one agent at the time and makes a choice on the basis of the neighbour's commitment state. Thanks to the well-mixed property of the system, the average agent behaviour tends to the one depicted in Figure 1 (see [8] for details).

For what concerns the transition probabilities based on the option quality, we propose two strategies based either on a *homogeneous* or on a *heterogeneous* implementation. In the homogeneous case, all agents compute their transition probabilities in the same way as a function of the estimated quality. In this case, it is possible to establish a linear correspondence between microscopic transition probabilities and macroscopic transition rates. Alternatively, agents may compute their own transition probabilities differently from each other. We propose a simple response thresholds scheme, so that agent follows a transition with a fixed probability if the (estimated) option quality exceeds a given threshold. With this implementation, it is possible to establish a relationship between microscopic and macroscopic parameters through the cumulative distribution function of the thresholds within the population of agents. For both homogeneous and heterogeneous strategies, the derivation of the relationship between microscopic and macroscopic description levels passes through the introduction of the master equation of the finite-size macroscopic model [8].

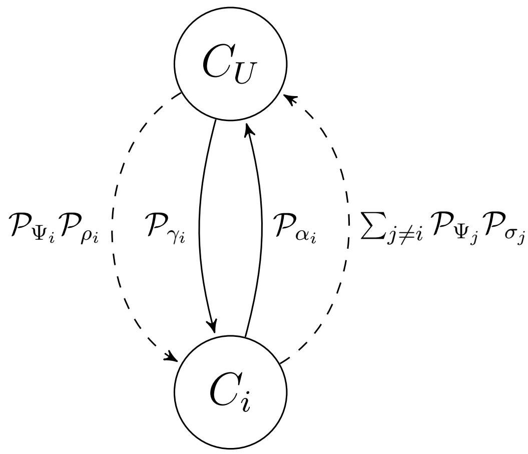

Figure 1: Probabilistic finite state machine (PFSM) representing the average individual behaviour. U represents the uncommitted state, while i represents commitment for the i-th option. Ψ_i represents the fraction of individuals committed to option i. $\alpha_i, \gamma_i, \rho_i, \sigma_i$ represent the processes of abandonment, discovery, recruitment and cross-inhibition, respectively. Spontaneous transitions are represented with bold lines, while dashed lines represent interactive transitions.

3.0 CASE STUDIES

The design pattern methodology introduced above can be exploited in different application domains, with limited adjustments to fulfil the specific requirements. We present here three case studies with different domains of application to exemplify the application of the design pattern for engineering swarm-centric decision making in artificial systems.

3.1 Multi-agent simulations on fully connected networks

As a starting example, we illustrate the implementation of decentralised decision-making for a multi-agent system in which each agent can potentially interact with any other agent. In this simplified scenario, all agents are presented multiple options and can perceive an objective quality, so that the collective choice must be performed by selecting the best or one among the equal-best alternatives. We focus on the general case of value-sensitive decision making as described in [6], hence we adopt the same macroscopic parameterization with a fixed cross-inhibition rate, discovery and recruitment linearly dependent on the option quality, and abandonment inversely proportional to the option quality. We first focus on a binary decision problem, in which the available options are referred to as *A* and *B* and their quality as v_A and v_B, respectively. We simulated the macroscopic dynamics by numerically integrating the ODE system and by means of the Gillespie algorithm [4] to take into account the finite size of the system, and we compared the performance with multi-agent simulations (see Figure 2 left). The correspondence between macroscopic model and microscopic implementation is remarkable. 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 plot in the bottom-right part of Figure 2 left. Conversely, the convergence time is very similar across different system sizes (see the top-left part of Figure 2 left). We also studied the micro-macro link in a best-of-n scenario. 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 results presented in Figure 3 reveal a very good correspondence between ODE, multi-agent and Gillespie simulations, therefore validating the methodology beyond the binary decision problems presented above.

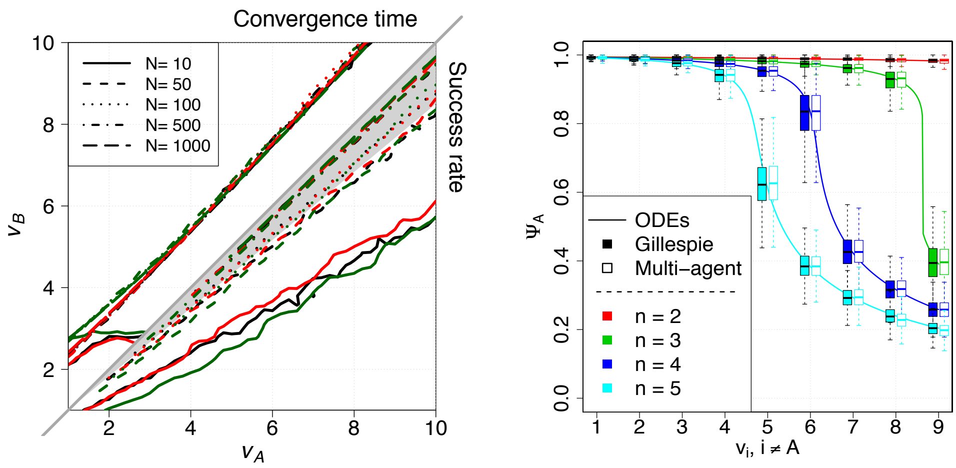

Figure 2: Left: Comparison between the stochastic finite-size macroscopic model (black lines) and the multi-agent implementation with both the homogeneous strategy (red lines) and the heterogeneous strategy (green lines). Results are displayed for varying system size N. For each possible configuration, 500 independent runs are performed. In the bottom-right half of the plot, we show the isolines for the success rate at the value 0.9, and the grey triangle indicates quality value pairs below the target resolution. In the top-left half we show the isolines for convergence time at the value 1s. Right: Micro-macro link with varying number of options. We compare the macroscopic dynamics predicted by the mean-filed ODE model, the finite-size macroscopic dynamics approximated by the Gillespie algorithm and the microscopic dynamics resulting from homogeneous multi-agent simulations ($N=500$ agents). 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.

3.2 Swarm robotics search and exploitation

As a more challenging case study, we deployed a swarm robotics system for search and exploitation of resources [7]. In this system, two or more resources must be identified by a swarm of robots and exploited by navigating back and forth from a home location (e.g., simulating the retrieval of some material from the resource). The quality of a resource is determined by its distance from the home location—closer resources being better than farther ones. We exploited the design pattern to implement the individual behaviour of each robot in the swarm. When uncommitted, robots search for resources using a random walk, and periodically return to the home location should they have not encountered any resource. When stumbling upon a resource, the robot gets committed to it and exploits odometry to navigate back and forth between home and resource locations. When at home, robots committed to a resource interact with other robots to recruit or inhibit them, according to the specifications given by the design pattern. The resource location is communicated too, so that recruited robots can navigate to the target location exploiting odometric information.

To evaluate the quality of the implemented behaviour and the match with macroscopic predictions, we have compared the collective response across several instances of the search and exploitation problem, by varying the distance of the available resource [7]. In the left part of Figure 3, for instance, we show a very good match between the swarm robotics system and the ODE predictions in terms of convergence dynamics. By systematically varying the probability of cross-inhibition in a symmetric decision problem (i.e., two resources placed at the same distance), we could verify also the existence of a bifurcation with increasing cross-inhibition as predicted by the macroscopic model (see the right part of Figure 3). Overall, the results again confirm the suitability of the design pattern for swarm-centric design [7].

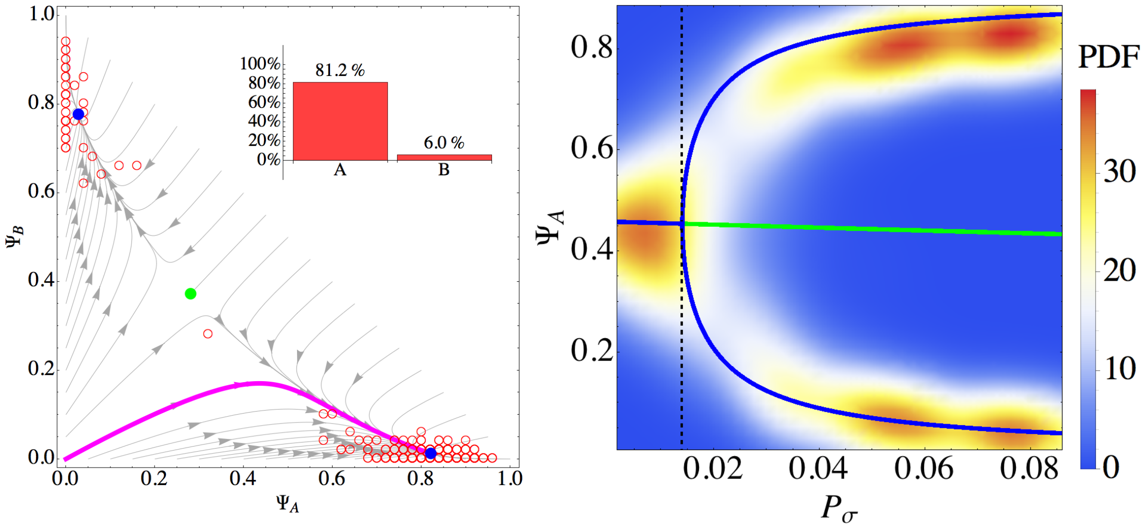

Figure 3: Left: Comparison of the swarm robotics system with the macroscopic ODE model for a binary choice scenario in which resource A is closer than resource B. The macroscopic dynamics are represented by trajectories and fixed points (blue: stable; green: unstable). The red dots represent the final repartition between resrouce A and B within the robot swarm, as resulting from simulations with 50 robots. Right: Bifurcation diagram comparing the dynamics predicted by the macroscopic model and the observation from multi-agent simulations of the swarm robotics search and exploitation scenario.

3.3 Coexistence in Cognitive Radio Networks

A different application scenario is provided by cognitive radio networks [1]. We consider the opportunistic access of unlicensed networks to the TV white space (TVWS, i.e., the channels within the TV spectrum that are not used for licensed broadcast). In such scenario, multiple cognitive radio networks may access in parallel the same channels, leading to possible interferences. We have exploited the design pattern for swarm-centric decision making to deploy a decentralised coexistence strategy that leads a cognitive radio network to (i) decide on which channel to allocate all its users and (ii) minimise interferences with other (heterogeneous) networks [12]. In this scenario, each channel represents an alternative for the decision process, with a quality dependent on the channel capacity q_i relative to the network demand d_j. The individual commitment dynamics follow the usual stochastic processes of discovery and abandonment, and interaction with other users of the same network lead to recruitment and cross-inhibition. In this way, each cognitive radio network can selfishly allocate itself to the best channel to fit its communication demand. At the same time, however, multiple networks compete to occupy the available channels, which results in communication interferences. By linking the individual abandonment rate to the perceived interference on the occupied channel, it is possible to deploy a decentralised strategy that provides the following attractive features: i) fully distributed, i.e., it avoids the need of centralized interference management; ii) over-the-air communications free, i.e., it avoids the need of direct communications among the heterogeneous networks; iii) adaptive to the time- and space-dynamics of the coexistence interference; iv) selfless, i.e., it allows a fair TVWS spectrum sharing by accounting for the communication demands of each unlicensed network. In Figure 4, we present results for numerical simulations with two available channels (A and B, with normalised capacity $q_A = 1$ and varying q_B) and three competing networks (with normalized demand $d_1=0.8$, and varying aggregated demand $d_{2,3}= d_2+d_3$, having $d_2/d_3=2/3$). The results are shown in Figure 4. Whenever an interference-free allocation is possible (e.g. $d_{2,3} \leq q_B$), the proposed strategy quickly provides a feasible solution, otherwise the best possible allocation is achieved by selfish choices starting from the distribution of users over the channels achieved at the end of the allotted decision time (i.e., 18 seconds, see Figure 4 left). In any case, the normalized achievable demand $d^\star$ (i.e., the actual allocated demand normalized on the maximum demand that can be allocated on the available channels) is always close to the optimal value of 1, as shown in the central panel of Figure 4. Additionally, the ratio between $d^\star$ and the allocated demand d^g using a random allocation proportional to channel qualities q_A and q_B suggests that the proposed strategy always outperforms a basic allocation with only a little overhead in terms of decision time (see Figure 4 right). Overall, the obtained results demonstrate the achievement of near-optimal performance through a completely decentralized decision process [12].

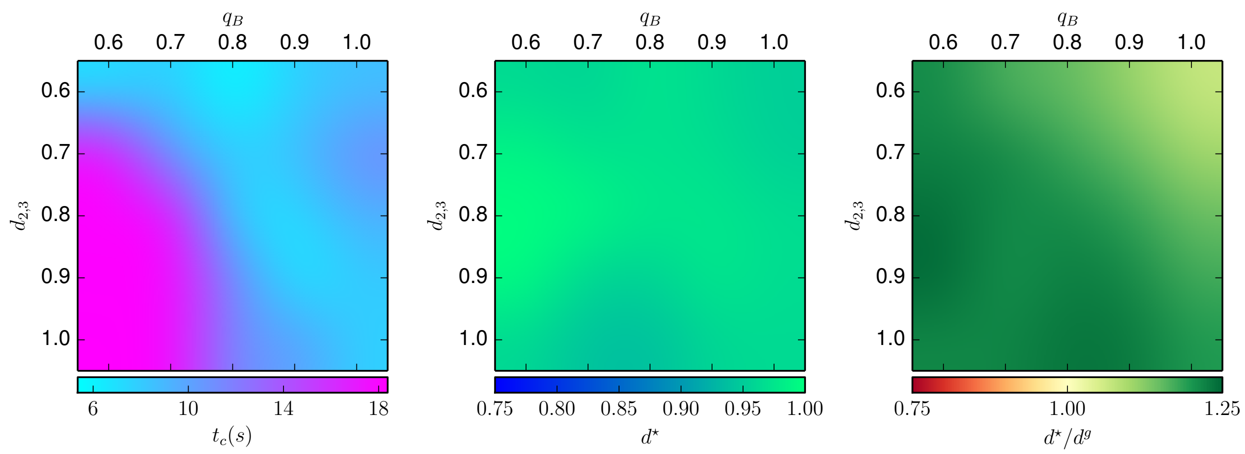

Figure 4: Results of the numerical simulations by varying the capacity q_B and the aggregated demand $d_{2,3}$. Left: average convergence time t_c required to achieve a stable allocation of three networks over the available channels. Centre: normalised allocated demand. Right: ratio between the normalised allocated demand and a random allocation proportional to the channel qualities.

3.0 CONLCUSIONS

The design pattern methodology we propose provides a complete framework that allows to move 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 [8]. We have substantiated the methodology with several case studies which confirm the suitability of the proposed strategy across application domains [7][8][12]. Besides engineering, our results can be relevant for better understanding the behaviour of natural systems, by providing testable hypotheses to be verified by field experiments.

REFERENCES

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

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

[3] Gamma, E., Helm, R., Johnson, R. & Vlissides, J. (1995). Design Patterns: Elements of Reusable Object-Oriented Software. Addison-Wesley Professional, Boston, MA.

[4] Gillespie, D. T. (1976). A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. Journal of Computational Physics, 22(4):403–434.

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

[6] Pais, D., Hogan, P. M., Schlegel, T., Franks, N. R., Leonard, N. E., & Marshall, J. A. R. (2013). A Mechanism for Value-Sensitive Decision-Making. PLoS ONE, 8(9):e73216.

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

[8] Reina, A., Valentini, G., Fernández-Oto, C., Dorigo, M., & Trianni, V. (2015b). A Design Pattern for Decentralised Decision Making. PLoS ONE, 10(10):e0140950–18.

[9] Seeley, T. D., Visscher, P. K., Schlegel, T., Hogan, P. M., Franks, N. R., & Marshall, J. A. R. (2012). Stop Signals Provide Cross Inhibition in Collective Decision-Making by Honeybee Swarms. Science, 335(6064):108–111.

[10] Srivastava, V., & Leonard, N. E. (2014). Collective Decision-Making in Ideal Networks: The Speed-Accuracy Tradeoff. IEEE Transactions on Control of Network Systems, 1(1):121–132.

[11] Trianni, V., Tuci, E., Passino, K. M., & Marshall, J. A. R. (2011). Swarm Cognition: an interdisciplinary approach to the study of self-organising biological collectives. Swarm Intelligence, 5(1):3–18.

[12] Trianni, V., Cacciapuoti, A. S., Caleffi, M., (2016) Distributed Design for Fair Coexistence in TVWS. Proceedings of the 2016 IEEE International Conference on Communciations, 23-27 May 2016, Kuala Lampur, Malaysia

[13] Vigelius, M., Meyer, B., & Pascoe, G. (2014). Multiscale Modelling and Analysis of Collective Decision Making in Swarm Robotics. PLoS ONE, 9(11):e111542–19.