

A Model of Symmetry Breaking in Collective Decision-Making

Heiko Hamann¹, Bernd Meyer², Thomas Schmickl¹, and Karl Crailsheim¹

¹ Artificial Life Lab of the Dep. of Zoology, Karl-Franzens University Graz, Austria
{heiko.hamann, thomas.schmickl, karl.crailsheim}@uni-graz.at

² FIT Centre for Research in Intelligent Systems, Monash University, Melbourne
bernd.meyer@acm.org

Abstract. Symmetry breaking is commonly found in self-organized collective decision making. It serves an important functional role, specifically in biological and bio-inspired systems. The analysis of symmetry breaking is thus an important key to understanding self-organized decision making. However, in many systems of practical importance available analytic methods cannot be applied due to the complexity of the scenario and consequentially the model. This applies specifically to self-organization in bio-inspired engineering. We propose a new modeling approach which allows us to formally analyze important properties of such processes. The core idea of our approach is to infer a compact model based on stochastic processes for a one-dimensional symmetry parameter. This enables us to analyze the fundamental properties of even complex collective decision making processes via Fokker–Planck theory. We are able to quantitatively address the effectiveness of symmetry breaking, the stability, the time taken to reach a consensus, and other parameters. This is demonstrated with two examples from swarm robotics.

1 Introduction

Self-organization is one of the fundamental mechanism used in nature to achieve flexible and adaptive behavior in unpredictable environments [1]. Particularly collective decision making in social groups is often driven by self-organizing processes. Some of the most prominent examples of this are found in social insects, for example the choice of nest sites [2] and food sources [3, 4] by ant colonies and the aggregation behavior of bees [5]. Countless other examples of self-organized collective decision making exist in both biological and human social systems, such as in quorum sensing in bacteria colonies [6] and in trend setting and following in economic markets [7]. Despite this variability the fundamental principles that govern self-organized collective behavior appear to be universal across the range [8].

Understanding the fundamental properties of self-organized collective decision making is thus of central importance in a broad spectrum of disciplines, ranging from biology and social science to bio-inspired engineering.

An important property commonly found in self-organized collective decision making is symmetry breaking. This refers to the fact that a system choosing

between a number of options may make a majority (or even an unanimous) decision for one of these options even if all of them have the same utility. From a biological perspective, symmetry breaking serves an important function. This is immediately obvious when we think about decisions that must be unanimous, such as the choice of a new nest site [9]. In other cases symmetry breaking still serves an important function even though its benefit may be less obvious. Consider food source exploitation by an ant colony. When faced with a choice between two equally good food sources, many species will converge on only one of them instead of exploiting both simultaneously [1]. There are several benefits of concentrating the colony's foragers on a single source. Arguably the most important one is that it is significantly easier to defend a single source against potential competition and aggressors. Further benefit may, for example, arise from the fact that collaborative transport of large prey items can be an efficient form of food retrieval [10] compared with the transport of small items by individual foragers. Similar considerations transfer immediately to technical applications of self-organized collective decision making, for example in swarm robotics.

Because of its central functional role, the analysis of symmetry breaking is one of the keys to a better understanding of collective decision making. This is widely accepted and a substantial body of work on the topics exists, much of which is based on bifurcation analysis. Recruitment in ant colonies is a well-known example where this approach has yielded a significant amount of insight [1].

Symmetry breaking in self-organized collective decision-making usually arises from the interplay of positive and negative feedback loops. For example, in mass recruiting ant colonies pheromone communication is the main mechanism for guiding foragers to food sources. Roughly speaking, pheromones attract potential foragers to a food source, and successful foragers deposit pheromones on paths to good food sources. This positive feedback is counterbalanced by negative feedbacks, namely the evaporation of pheromone and overcrowding. The balance between these two influences leads to a stable yet flexible decision system [3, 4].

Arguably the most common approach to the analysis of such coupled feedback systems is via differential equation systems. In sufficiently simple cases, bifurcation analysis can then be used to establish properties of the symmetry breaking process. However, as the complexity of the scenario and the individual behavior increases, an analysis of the full differential equation model soon becomes extremely involved (see e.g. [11]) and often is simply impossible. Moreover, finding such a model can be difficult in the first place, specifically in cases where the behavior of individuals is relatively complex.

Thus, alternative ways to model the symmetry breaking properties in such systems are required. In this paper we propose such a method. The core idea of our approach is to dispense with complex mechanistic models which capture feedback loops or even individual behavior explicitly. We replace these with significantly simpler phenomenological models that focus exclusively on the description of symmetry breaking. The advantage of doing this is that we obtain a compact mathematical model which is amenable to a formal analysis. In this way we can analytically obtain information about the fundamental properties of the process,

such as the stability of decisions. We will illustrate this approach and such an analysis with two examples from swarm robotics. The first one is the aggregation behavior of a swarm whose control algorithms are modeled on honeybee behavior [12–14], the second one is an emergent density classification task [15].

2 Overall approach

To introduce our approach we consider the simplest possible scenario of a binary choice between two options A, B with equal utility. We will say that a system with N individuals (agents) exhibits symmetry breaking, if a significant majority of it decides for either A or B regardless of their equal utilities. We define a symmetry parameter (w.l.o.g. with respect to option A) as $s(t) = L(t)/N$, where $L(t)$ is the number of agents that have chosen option A at time t . A majority decision is any outcome in which the process has converged to a state where at least $L \geq \delta N$ individuals have chosen option A with $0.5 \ll \delta \leq 1.0$. In terms of s this simply means $s \geq \delta$. $s(t)$ essentially captures all information required to analyze the symmetry breaking properties of the process.

If $s(t)$ converges we can in principle obtain the steady state probability density function (PDF) for $s(t)$. This steady state PDF $\rho^*(s)$ tells us how likely it is that a certain proportion of agents decides for option A . In a binary symmetry breaking scenario, $\rho^*(s)$ will be bimodal and (possibly) symmetric (see Fig. 3 for the steady state PDFs of our two example scenarios). Given $\rho^*(s)$ we can thus calculate the expected proportion P of experiments in which a majority decision with at least δ majority occurs as

$$\int_0^\delta \rho^*(s) ds + \int_{1-\delta}^1 \rho^*(s) ds = P. \quad (1)$$

Provided $s(t)$ converges we can in principle always obtain $\rho^*(s)$ at least approximately by statistical evaluation of a large number of sample simulations. The same is true of other properties of $s(t)$, for example the time to convergence. This is indeed the standard approach to the analysis of such systems if no analytic description is available: A symmetry parameter is defined, the full system is simulated, and the statistics of the symmetry parameter is measured from the results of these (typically high-dimensional) simulations. This approach has two drawbacks: Firstly, it can be computationally extremely costly as obtaining numerical values for individual properties requires the simulation of a large number of sample developments. Secondly, this form of statistics only gives us numeric results for individual properties and does not allow us to perform a more general analysis of these properties.

An explicit representation of $s(t)$ would obviously enable us to perform significantly more powerful analysis provided it is simple enough to stay mathematically tractable. This idea is at the very center of our approach.

We simply and somewhat boldly postulate that the development of $s(t)$ can be described by a 1-d Langevin equation, a particular form of stochastic differential

equation:

$$\frac{ds}{dt} = \alpha(s, t) + \beta(s, t)\xi(t), \quad (2)$$

where α describes the deterministic development (so-called drift), ξ a Gaussian noise $|\xi(t)| = 1$, with mean $\langle \xi(t) \rangle = 0$, and uncorrelated in time $\langle \xi(t)\xi(t') \rangle = \delta(t - t')$, and β captures the fluctuation of the noise amplitude. Having a definition of $s(t)$ in this form unlocks a whole repertoire of formal tools for the analysis of stochastic differential equations that allow us to calculate most of the properties that we are interested in [16].

Two questions arise immediately: (1) Does this description indeed exist for a given system, i.e. is it possible at all to correctly reduce the behavior of the (potentially high-dimensional) system to a one-dimensional system? (2) Even if it is possible, how can we infer α and β ?

The first question cannot generally be answered, and some systems will not admit such a description. However, our case studies give some indication that it should be possible to approximate the behaviour of many interesting system quite well. Our methodology approaches the question optimistically by assuming the existence of such a description: We first attempt to infer α and β . This can be done via a heuristic argument (as evidenced by the case studies) or with standard numerical fitting techniques. Once candidates for α and β are obtained the crucial step is their verification. This is performed by obtaining a large number of sample developments for $s(t)$ in two complementary ways: One set of samples is generated by simulating the full original (mechanistic) model and measuring $s(t)$ at each simulation step, another set of samples is obtained by numeric forward integration of $s(t)$ according to Eq. 2 for given α and β . If the two sample sets are not in statistical agreement we have to dismiss the candidate functions α, β . However, if the sets agree statistically, we are justified in our choice of α and β and the simplified model Eq. 2 captures the relevant aspects of the process statistically correctly. In this case we may discard the full model and simply proceed with an analysis of the simplified model.

3 Investigated Scenarios

We will now illustrate this approach with two example scenarios from swarm robotics. Both are instances of homogeneous multi-agent systems (MAS), where the agents move in a two dimensional rectangular arena surrounded by walls.

Our first scenario is collision-based adaptive aggregation. The task of the robot swarm is to aggregate at the brightest spot in the arena. This is essentially a physically embodied simulation of the behavior of young honeybees, that typically aggregate at areas of a certain temperature. The system and its distributed control algorithm, termed BEECLUST, have been described in full detail in [13, 17], and we only give a brief summary here.

The robots are equipped with sensors for distance measurements as well as a sensor that allows them to measure a special inhomogeneous property of the arena (e.g., light). In addition, they are able to identify other robots as such. Each

robot moves in a straight line (with initial random heading) until it perceives an obstacle Ω within sensor range. If Ω is a wall it turns away and moves straight again. If Ω is another robot, it counts the number of other robots K in the vicinity. If $K \geq \sigma$ the robot measures the local luminance. The higher the luminance the longer the agent stays stopped. After the waiting period has elapsed, the agent turns away from the other agent and moves straight again.

The collective aggregation at the brightest spot is a consequence of positive feedback. Waiting times are longer at bright spots and the agent density is increased through clustering in these regions. Negative feedback is induced by the saturation of the limited space with high brightness.

Initially, the agents have random headings, are in the state ‘moving’, and are random uniformly distributed in the whole arena (i.e., on average we have initially the same number of robots in the left and in the right half of the arena). The luminance distribution in the test arena is bimodal with maxima of the same value and shape in the left and right half of the arena (for details, see [12, 14]). As a measure of symmetry we use $s_b(t) = L(t)/N$ (‘b’ for BEECLUST) where $L(t)$ is the number of robots in the left half of the arena, and N the swarm size. Our second scenario is an emergent density classification task. Robots in the swarm exist in two states ‘red’ and ‘green’. The task of the robot swarm is to estimate whether there are initially more green or more red swarm members, i.e. to converge on a majority decision. This problem is derived from a well known example of emergent computation in cellular automata [18]. N robots are randomly distributed in the arena and randomly initialized to be red or green in a given proportion. The robots move and perform collision avoidance, i.e. if another robot comes too close they turn away to increase the distance. At each encounter the robots remember the color of each other. After five robot encounters, each robot changes its color to the one it encountered most often. We are interested in the question whether (a majority of) the swarm stably converges on one color and whether this is the original majority color (for details, see [15]). As a measure of symmetry we use $s_d(t) = R(t)/N$ (‘d’ for ‘density classification’) where $R(t)$ is the number of red robots and N the total swarm size.

Both scenarios include a stochastic component. The density classification scenario includes a stochastic component as we explicitly account for errors in the color recognition. We assume that a robot recognizes the color of the encountered robot correctly only with a given probability $\gamma = 0.8$. While noise is explicitly incorporated in the case of the density classification scenario, it implicitly enters the BEECLUST scenario. This is because the basic movement and collision-avoidance mechanism is a billiard-like system that introduces pseudo-randomness through deterministic chaos.

4 Dynamics of the Symmetry Parameter

As reported in [12], most macroscopic characteristics of the collective decision processes of these systems can approximately be captured by two features of the

symmetry parameter s . First, the mean of the absolute changes

$$\Delta s^{\text{abs}}(s, t) = \frac{1}{K} \sum_i |s_i(t) - s_i(t-1)|, \quad (3)$$

averaged over K samples $s_i(t)$ obtained from many independent simulation runs. Second, the mean of the relative changes

$$\Delta s^{\text{rel}}(s, t) = \frac{1}{K} \sum_i s_i(t) - s_i(t-1), \quad (4)$$

which is an approximation of the derivative. In Fig. 1, we show the results of the measurements for both scenarios. Note that Δs^{rel} at $s = 0.5$ averages to $\Delta s^{\text{rel}}(s = 0.5, t) \approx 0$ as expected (Fig. 1(b), 1(d)), while the converged absolute changes keep a maximum at $s = 0.5$ (Fig. 1(a), 1(c)). Fig. 1 shows that Δs^{abs} and Δs^{rel} are indeed time-variant. Δs^{abs} basically keeps its shape and is simply scaled down with time, whereas Δs^{rel} even changes its shape. It only converges after some time to a function that can cause multiple stable (or meta-stable) decisions (as it exhibits multiple zero-crossings corresponding to fixpoints).

For numeric calculations we will subsequently simply use the (time-invariant) values for Δs^{abs} and Δs^{rel} which are reached with an error exponentially decreasing in time as shown in Fig. 1. While this is clearly a rather drastic simplification, our choice will be justified by the verification step of our basic methodology.

Based on Δs^{rel} and Δs^{abs} we suggest the following heuristics to infer candidates for α and β . Consider the development of s in the time-discrete simulation. We suggest to approximate this as

$$s_{t+1} = s_t + \Delta s^{\text{rel}}(s_t) + (\Delta s^{\text{abs}}(s_t) - |\Delta s^{\text{rel}}(s_t)|)\xi_t, \quad (5)$$

for Δs^{rel} and Δs^{abs} as defined above, and for a Gaussian white noise ξ_n . The assumption of Gaussian white noise is of course only an approximation which could be improved by measuring the distributions of the Δs for each time step. To justify Eq. 5, we consider the three points $s \in \{0, 0.5, 1\}$ in Fig. 1. At $s = 0$ and $s = 1$ we have $|\Delta s^{\text{rel}}| = \Delta s^{\text{abs}}$ because changes are only possible in one direction at the borders. Hence, the influence of the noise term should be zero there. At $s = 0.5$ we have, as expected, $\Delta s^{\text{rel}} = 0$. Hence, the only influence is due to Δs^{abs} , which is fulfilled by Eq. 5.

We make the transition to a continuous time version by using Eq. 2 with drift $\alpha(s) = \Delta s^{\text{rel}}(s)$ and diffusion coefficient $\beta(s) = \Delta s^{\text{abs}}(s) - |\Delta s^{\text{rel}}(s)|$. For this continuous model we can now use the Fokker–Planck equation

$$\frac{\partial \rho_s}{\partial t} = \frac{\partial}{\partial s}(\alpha(s, t)\rho_s) + \frac{1}{2} \frac{\partial^2}{\partial s^2}(\beta^2(s, t)\rho_s) \quad (6)$$

to obtain the time development of the probability density function for s , and specifically its steady-state PDF. As outlined earlier for our methodology, we must verify whether this candidate model is valid by comparing two versions of

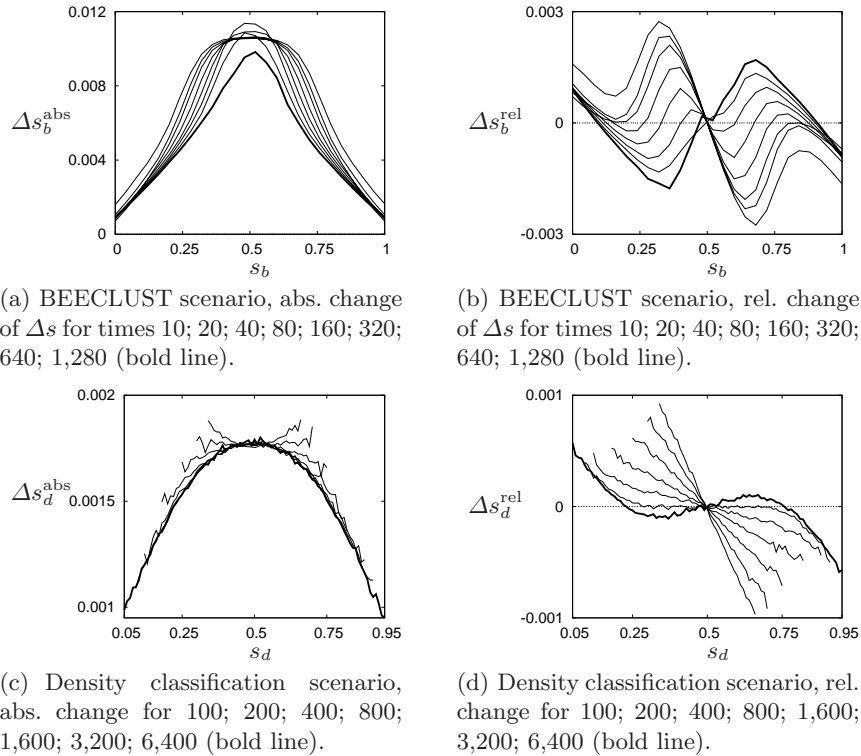


Fig. 1. Measured dynamics of the symmetry parameter; 2×10^5 samples per scenario.

the PDF for $s(t)$ obtained in different ways: once by simulating the full system and once by solving the Fokker–Planck equation for the 1-d system.

Fig. 2 shows the densities obtained via simulation and Fokker–Planck equation, respectively, together with a typical sample trajectory. These plots were obtained by solving Eq. 6 numerically for the initial value problem of an initial peak at $s = 0.5$ for $t = 0$ and with the corresponding full simulation. It shows that there is good qualitative correspondence between the two systems and thus that the simplified model with the chosen α, β on this level adequately describes the evolution of the symmetry parameter for these systems. We are thus justified in dispensing with the full model and conducting an analysis of symmetry breaking properties based only on the simplified model, which is amenable to a formal analysis.

5 Deriving Properties of the Collective Decision System

The simplified model can now be used to analyze the symmetry breaking properties of the system. To begin, consider the effectivity of symmetry breaking. It




Fig. 2. Density classification scenario; PDF measured in the simulation ($\hat{\rho}_s$, based on 2×10^5 samples) and obtained by solving Eq. 6 (using the converged measurements) and superimposed samples of trajectories of s from the simulation and by sampling Eq. 2.

can be quantified via the steady state of the Fokker–Planck equation. Assuming reflecting boundaries at $s = 0$ and $s = 1$, the steady state $\rho^*(s)$ is given by

$$\psi(s) = \exp\left(\int_0^s 2\frac{\alpha(y)}{\beta^2(y)} dy\right), \quad \rho^*(s) = C\frac{\psi(s)}{\beta(s)}, \quad (7)$$

with a normalization constant C [16, 4]. In Fig. 3(a), the results obtained from Eq. 7 are shown for both scenarios using the functions α and β as defined above in Eq. 2 and with the data shown in Fig. 1 (with noise). In addition, the positions of δ as defined by Eq. 1 are given and the associated areas are marked in gray. For both scenarios $\rho^*(s)$ is clearly bimodal and exhibits very low densities around $s \approx 0.5$, hence the symmetry breaking is effective.

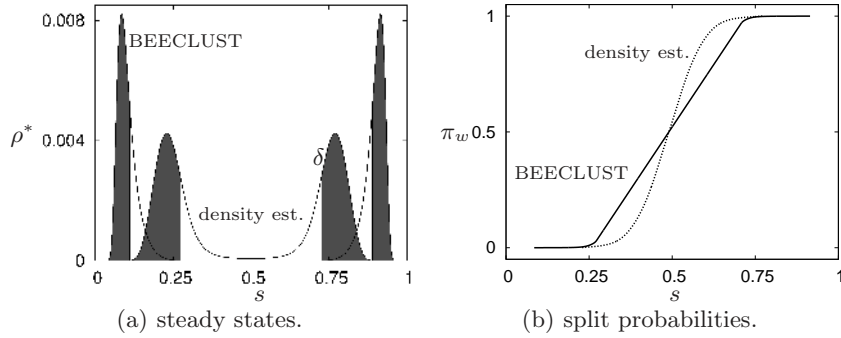


Fig. 3. Analytically and numerically obtained measures of symmetry breaking.

As a further example of a possible analysis we quantify the stability of decisions. An appropriate measure for this is the probability of revising a majority decision. This is similar to the concept of ‘splitting probabilities’ [16]. The split-

ting probability $\pi_w(x)$ gives the probability that the system initialized at $s = x$ will reach the state $s = w$ before the state $s = u$, i.e. it will make a majority decision for w . The corresponding probability $\pi_u(x)$ is defined symmetrically. The splitting probability can be calculated as

$$\pi_w(x) = \int_u^x \rho^*(s)^{-1} ds \left(\int_u^w \rho^*(s)^{-1} ds \right)^{-1}, \quad (8)$$

if we take u and w to correspond to the positions of the two peaks in the PDF (i.e. the decision states), we can immediately read off various important properties from a plot of $\pi_w(x)$ particularly how likely it is that the system reaches a given decision from a given start state and thus also how likely it is that a decision, once made, will be reversed giving a measure of the stability of the decision process.

The numerically obtained results for $\pi_w(x)$ are shown in Fig. 3(b). Based on this, the stability of the density classification scenario has to be classified as higher than that of the BEECLUST scenario, because its slope around the center point ($\pi = 0.5$) is steeper and it reaches saturation ($\pi = 0$ or $\pi = 1$) more quickly.

In the same way a host of other properties can be analyzed using the Fokker-Planck equation. In some cases, such as for the convergence time, a full analytic treatment may be difficult, particularly for time-variant functions α and β . However, a numeric estimate can still be obtained as shown in Fig. 2(b).

6 Conclusion

We have proposed a generalized approach to analyzing symmetry breaking in collective decision making. It is based on a simple stochastic process model that describes the development of a one-dimensional symmetry parameter. The model and its parameters are inferred by measurements from (real or virtual) experiments. The models obtained are simple enough to be amenable to powerful analytic techniques using Fokker-Planck equations. This allows us to calculate important properties, such as the effectiveness of the decision process, likelihood of decisions to be reached and to be revised, time to reach a consensus etc. without having to rely on further simulations. We suspect that our heuristics for inferring the model parameters (drift and diffusion coefficients) based on Δs^{rel} also applies to other self-organized collective decision making processes. This assumption receives some support from their successful use in the analysis of the two substantially different example scenarios. It is crucial where this heuristics cannot be used, the core of our approach remains applicable as it allows us to use a range of other methods to identify and verify candidates for α and β . We thus expect our approach to be in principle useful for a wide range of self-organized collective decision making scenarios.

Acknowledgments

This work is supported by: EU-IST-FET project ‘SYMBRION’, no. 216342; EU-ICT project ‘REPLICATOR’, no. 216240; FWF research grant ‘Temperature-induced aggregation of young honeybees’, no. P19478-B16 and ARC DP0879239.

References

1. Camazine, S., Deneubourg, J.L., Franks, N.R., Sneyd, J., Theraulaz, G., Bonabeau, E.: Self-Organization in Biological Systems. Princeton Univ. P. (2001)
2. Franks, N.R., Mallon, E.B., Bray, H.E., Hamilton, M.J., Mischler, T.C.: Strategies for choosing between alternatives with different attributes: exemplified by house-hunting ants. *Animal Behavior* **65** (2003) 215–223
3. Dussutour, A., Beekman, M., Nicolis, S.C., Meyer, B.: Noise improves collective decision-making by ants in dynamic environments. *Proceedings of the Royal Society London B* **276** (December 2009) 4353–4361
4. Meyer, B., Beekman, M., Dussutour, A.: Noise-induced adaptive decision-making in ant-foraging. In: SAB 2008. Number 5040 in LNCS, Springer (2008) 415–425
5. Kernbach, S., Thenius, R., Kornienko, O., Schmickl, T.: Re-embodiment of honey-bee aggregation behavior in an artificial micro-robotic swarm. *Adaptive Behaviour* **17** (2009) 237–259
6. Reading, N.C., Sperandio, V.: Quorum sensing: the many languages of bacteria. *FEMS Microbiol Lett.* **254**(1) (January 2006) 1–11
7. Weisbuch, G., Stauffer, D.: Hits and flops dynamics. *Physica A*, 287:563576, 2000.
8. Vicsek, T.: A question of scale. *Nature* 411:421, 2001.
9. Jeanson, R., Deneubourg, J.L., Grimal, A., Theraulaz, G.: Modulation of individual behavior and collective decision-making during aggregation site selection by the ant *messor barbarus*. *Behav. Ecol. Sociobiol.* **55** (2004) 388–394
10. Garnier, S., Gautrais, J., Theraulaz, G.: The biological principles of swarm intelligence. *Swarm Intelligence* 1:3–31, 2007.
11. Nicolis, S.C., Deneubourg, J.L.: Emerging patterns and food recruitment in ants: an analytical study. *Journal of Theoretical Biology*, 198:575–592, 1999.
12. Hamann, H., Schmickl, T., Wörn, H., Crailsheim, K.: Analysis of emergent symmetry breaking in collective decision making. *Neural Comp. & Appl.* (2010) in press
13. Schmickl, T., Thenius, R., Möslinger, C., Radspieler, G., Kernbach, S., Crailsheim, K.: Get in touch: Cooperative decision making based on robot-to-robot collisions. *Autonomous Agents and Multi-Agent Systems* **18**(1) (February 2008) 133–155
14. Schmickl, T., Hamann, H.: BEECLUST: A swarm algorithm derived from honeybees. In Xiao, Y., Hu, F., eds.: *Bio-inspired Computing and Communication Networks*. Routledge (August 2010)
15. Hamann, H., Wörn, H.: Embodied computation. *Parallel Processing Letters* **17**(3) (September 2007) 287–298
16. Gardiner, C.W.: *Handbook of Stochastic Methods for Physics, Chemistry and the Natural Sciences*. Springer-Verlag (1985)
17. Hamann, H., Wörn, H., Crailsheim, K., Schmickl, T.: Spatial macroscopic models of a bio-inspired robotic swarm algorithm. In: *IEEE/RSJ 2008 Int. Conf. on Intel. Robots and Syst. (IROS’08)*, Los Alamitos, CA, IEEE Press (2008) 1415–1420
18. Packard, N.H.: Adaptation toward the edge of chaos. In: *Dynamic Patterns in Complex Systems*. World Scientific (1988) 293–301