

Analysis of Swarm Behaviors Based on an Inversion of the Fluctuation Theorem

Heiko Hamann*, Thomas Schmickl, and Karl Crailsheim

Artificial Life Lab
of the Department of Zoology,
Universitätsplatz 2,
Karl-Franzens University Graz,
8010 Graz, Austria

Phone: +43 316 380 3982

{heiko.hamann, thomas.schmickl, karl.crailsheim}@uni-graz.at

March 8, 2013

Abstract

A grand challenge in the field of artificial life is to find a general theory of emergent self-organizing systems. In swarm systems most of the observed complexity is based on motion of simple entities. Similarly statistical mechanics focuses on collective properties induced by motion of many interacting particles. In this paper we apply methods from statistical mechanics to swarm systems. We try to explain the emergent behavior of a simulated swarm by applying methods based on the fluctuation theorem. Empirical results indicate that swarms are able to produce negative entropy within an isolated sub-system due to ‘frozen accidents’. Individuals of a swarm are able to locally detect fluctuations of the global entropy measure and store them, if they are negative entropy productions. By accumulating these stored fluctuations over time the swarm as a whole is producing negative entropy and the system ends up in an ordered state. We claim that this indicates the existence of an inverted fluctuation theorem for emergent self-organizing dissipative systems. This approach bears the potential of general applicability.

Keywords: fluctuation theorem, swarm behavior, frozen accidents, fluctuations

*corresponding author

1 Introduction

One characteristic of living organisms is their metabolism. Living beings require energy in order to maintain their internal order. This is determined by the second law of thermodynamics that describes the ubiquitous decay of all things and does not allow the increase of order without the cost of dissipation. In the context of self-organizing systems one might cite Parunak and Brueckner [16]: “Emergent self-organization in multi-agent systems appears to contradict the second law of thermodynamics.” This is of course not the case, as discussed by Parunak and Brueckner [16], one has to distinguish between two kinds of sub-systems: one that hosts the self-organizing swarm and one in which disorder is increased. Hence, a swarm can be thought of as a heat pump that decreases entropy (entropy is defined in Sec. 2 and entropy measures are investigated in Sec. 5.1) in one basin in favor of increased entropy in another basin. However, the question of how the swarm manages to do that still persists. Whether thermodynamic properties are relevant and helpful in understanding such systems is currently discussed [17, 10].

The complexity of the animate world is explained by natural selection in combination with random events (natural evolution). It is one thing to select the adapted organism but the mutation, that results in an improved adaptivity, has to occur first. Concerning the genetic code Crick [3] phrased the term ‘Frozen Accident Theory’. While Crick was introducing this concept with focus on genetics, Gell-Mann [6] applied it to everything:

[...] the effective complexity [of the universe] receives only a small contribution from the fundamental laws. The rest comes from the numerous regularities resulting from ‘frozen accidents’.

We want to define a concept of frozen accidents within emergent self-organizing multi-agent systems [28] that explains how they can work as heat pumps in the sense as described above.

While a heat pump has to work against the second law (e.g., diffusion of heat) by expending energy, limited violations of the second law without the expenditure of energy [4] are also possible as, for example, indicated by Maxwell [14]:

The truth of the second law is ... a statistical, not a mathematical, ... for it depends on the fact that the bodies we deal with consist of millions of molecules.

Violations of the second law are possible for small systems and short time scales, that is, at atomic and micron scales over short times (up to two seconds), and were shown experimentally [26]. We claim that the reduction of entropy by emergent self-organizing systems could be explained by the ‘summation’ of such violations to the second law. As already noted, the second law is only statistical and, hence, allows spontaneous decreases of entropy in isolated systems with nonzero probability.

The possibility of a time-limited decrease in entropy exists because a system at a temperature above absolute zero according to statistical mechanics always shows thermal fluctuations, that are random deviations of a system from its equilibrium. Say x is a thermodynamic variable (i.e., it describes a state of a thermodynamic system at a given time), then the probability distribution $f(x)$ of this variable for a system at maximum entropy (i.e., at equilibrium state) turns out to be approximately Gaussian with mean $\mu = 0$:

$$f(x) = \frac{1}{\sqrt{2\pi\langle x^2 \rangle}} \exp\left(-\frac{x^2}{2\langle x^2 \rangle}\right), \quad (1)$$

for the variance defined by the mean square fluctuation $\sigma^2 = \langle x^2 \rangle$, which is an average over many ensembles (i.e., average over many realizations of the system). Hence, the probability of observing negative ($\int_{-\infty}^0 f(x)dx$) or positive fluctuations ($\int_0^{+\infty} f(x)dx$) is equal at equilibrium.

The fluctuation theorem [5, 4] quantifies the probability of violations to the second law. For short intervals it can be said that nature was running in reverse because the sign of entropy production is used to define the direction of the arrow of time. Even in reference to living systems this might be said. For example, small ‘machines’ within a cell (e.g., mitochondria) are likely to run in reverse from time to time. A transfer of this concept to the macro-world is typically denied categorically. In a review of Wang et al. [26], Gerstner [7] wrote: “For larger systems over normal periods of time, however, the second law of thermodynamics is absolutely rock solid.”

Generally the fluctuation theorem is said to be applicable only to the micro-world, where Brownian motion can be observed. Truly, this is a well chosen hypothesis. However, what if we allow dissipation of energy in the first place, separate the system in two sub-systems of the self-organizing part and a heat bath, and then observe only the behavior in the self-organizing half of the system? That way one could argue that we simulate the




Figure 1: Schematic of a system divided into a heat bath with increasing entropy and a self-organizing, dissipative sub-system with decreasing entropy.

micro-world by a macro-system at the cost of lost heat. This concept (see Fig. 1) is for example taken into account by Smith [24] when stating

$$dQ = -k_B T dS \equiv k_B T d\mathcal{I}, \quad (2)$$

for an increment of heat dQ rejected by the system to a thermal bath at temperature T , Boltzmann constant k_B , reduction in entropy of the (sub-)system's internal state $-dS$, and the increase in information $d\mathcal{I}$ (note that Smith [24] defines information as “the reduction in some measure of entropy”). Note that the mere property of being dissipative is not sufficient to explain a self-organizing system. In addition to squandering energy the system has to generate orderly structures. Dissipation is only a necessary condition for negative entropy production while additional sufficient conditions exist. In case of Rayleigh-Bénard convection [2], for example, initially fluctuating flows [29] occur that are enhanced and trigger the formation of Bénard cells in spontaneous symmetry breaking, cf. also [15, 8]. We want to point out the self-amplification of fluctuations as such a sufficient condition here.

In this paper, we report empirical evidence that the negative entropy production in emergent self-organizing systems is based initially on frozen accidents allowed by the original fluctuation theorem which, in turn, leads in the end to a global behavior that is described by an inversion of the fluctuation theorem in dissipative self-organizing systems. This concept might bear potential of embedding the concept of emergent behavior in multi-agent systems (swarms, self-propelled particles etc.) in a theoretical framework built on sound foundations of theories from physics. Hence, we propose an approach to understand emergent behavior through thermodynamics which follows up our earlier reported concept [10].

In addition, the relation to the fluctuation theorem might allow to define preconditions for effective self-organizing systems in the future. For example, one can define minimum requirements for the agents of the system concerning its cognitive abilities in order to be able to leverage fluctuations. The agent needs sensors that allow to estimate at least probabilistically whether the (local) entropy has just decreased. Furthermore, the system needs the ability to store such local fluctuations.

In the following we describe the fluctuation theorem and the investigated scenarios. We analyze the multi-agent system or swarm, discuss how the results could be viewed as obeying an inverted fluctuation theorem and conclude by giving a short summary and outlook.

2 Fluctuation Theorem and entropy measures

According to Evans and Searles [5] the group of fluctuation theorems “gives an analytical expression for the probability of observing Second Law violating dynamical fluctuations in thermostatted¹ dissipative non-equilibrium systems.” One of these theorems (steady state fluctuation theorems) applies to time-reversible, thermostatted, ergodic dynamical systems and defines the relation of fluctuations [5]

$$\lim_{t \rightarrow \infty} \frac{1}{t} \ln \frac{P[\bar{\Sigma}_t = A]}{P[\bar{\Sigma}_t = -A]} = A, \quad (3)$$

for the time averaged entropy production $\bar{\Sigma}_t = (1/t) \int_0^t \Sigma(s) ds$. The fluctuation theorem compares probabilities of observing a certain time averaged entropy production A to its negative value $-A$. The value $P(\bar{\Sigma}_t = A)$ describes the probability of finding the system initially in those states that subsequently generate bundles of

¹In a thermostatted system the temperature is kept constant, for example, by rescaling the particles' velocities. The system can be thought of as being in contact with a large heat reservoir in order to thermostat the system [1]. A thermostat method is applied in Sec. 4.1.

trajectory segments with the time averaged value A . The above theorem (eq. 3) predicts an exponential increase of the relation $P(\overline{\Sigma}_t = A)/P(\overline{\Sigma}_t = -A)$. Hence, with increasing time positive entropy producing trajectories become exponentially more likely than their negative entropy producing counterparts. As a consequence of the fluctuation theorem one obtains the Second Law Inequality

$$\langle \overline{\Sigma}_t \rangle \geq 0, \quad \forall t, \quad (4)$$

which states that the average over many ensembles, in which the time averaged entropy productions were measured, is positive. Hence, the fluctuation theorem is in accordance with the second law of thermodynamics. In the following we want to measure entropy productions and consequently entropy in swarm systems. It is important to note that there is not one generally applicable measure of entropy. Within statistical mechanics entropy is defined in general based on a set of microstates and their probabilities p_i

$$S = -k_B \sum_i p_i \ln p_i. \quad (5)$$

A microstate is said to be a detailed description of the system's current configuration. A macrostate is a much shorter description of such a configuration usually based on averages and, hence, summarizes a set of microstates. Macrostates are defined with an application in mind as stated by Jaynes [13]:

But it is also possible that two experimenters assign different entropies [...] to what is in fact the same microstate [...] without either being in error. That is, they are contemplating a different set of possible macroscopic observations on the same system, embedding that microstate in two different reference classes [...]

it is necessary to decide at the outset of a problem which macroscopic variables or degrees of freedom we shall measure and or control

In order to measure probabilities of microstates effectively one needs to define a countable set of microstates which is usually done by coarse graining (an example in the context of multi-agent systems is given by Parunak and Brueckner [16]). In the context of entropy production it is also important to determine whether negative or positive entropy production is observed. Please note that this 'sign of entropy production' might be opposite to the actual mathematical sign of a certain entropy measure. According to the statistical physics interpretation, this sign of entropy production is defined relative to the entropy of that macrostate which has the highest number of microstates. According to thermodynamics it is defined relative to the entropy of that state with the lowest amount of energy that can be used to do thermodynamic work. That macrostate is the equilibrium state, which would be achieved in the context of this paper by mere random behavior of the particles/agents, that is, a uniform distribution of agents. Thus, any process that leaves the equilibrium state and enters states of lower entropy shows negative entropy production.

3 Investigated scenarios

We investigate two scenarios. In the first one the swarm is controlled by the BEECLUST algorithm. The second scenario is a simple clustering process and is introduced to give evidence that our approach has potential for general applicability.

3.1 BEECLUST algorithm

The BEECLUST algorithm can be considered a model algorithm for swarms. It is based on observations of young honeybees [25], was analyzed in many models [12, 19, 20, 11, 9], and even implemented in a swarm of robots [21]. This algorithm allows a swarm to aggregate at a maximum of a scalar field although individual agents do not perform a greedy gradient ascent. In addition, a BEECLUST-controlled swarm is able to break symmetries [11] of equal maxima in the scalar field as also observed here (e.g., see Fig. 2). Hence, it might be justified to call this emergent behavior. Controlled by this algorithm three agents will stop (i.e., stopping threshold set to three; note that in previous works typically two agents were enough for stopping; the actual choice of this stopping threshold is, however, irrelevant in this paper) when they approach each other, measure the local value of the scalar field, and wait for some time proportional to this measurement. Clusters form and finally the swarm will be aggregated close to the global optimum of the scalar field (see the lower part of Fig. 2). See Fig. 3 for a definition of the BEECLUST algorithm.




Figure 2: Bottom: Typical state of a swarm controlled by BEECLUST; positions of stopped agents (circles) and moving agents (triangles) with trajectories of the last 20 time steps, contours show levels of the scalar field. Top: function used in eq. 8.

- 1.) Each agent moves straight until it perceives an obstacle O within sensor range.
- 2.) If O is a wall the agent turns away and continues with step 1.
- 3.) If O is another agent and there is a third agent as well, the agent measures the local scalar field value. The higher the scalar field value the longer the agent stays still. After this waiting period, the agent turns away from the other agent and continues with step 1.

Figure 3: The BEECLUST algorithm (stop threshold of 3 agents).

Table 1: Used parameter setting in this work.

arena dimensions	150×50 [length units] ²
proximity sensor range	3.5[length units]
max. waiting time	660[time units]
velocity	4[length units]/[time units]
number of agents	25

The collective aggregation close to the global optimum is achieved via a positive feedback process [11]. Clusters of 3 stopped agents will form by chance anywhere in the arena. The area covered by clusters grows with the number of contained agents and clusters covering a bigger area are more likely to be approached by chance by moving agents. Hence, bigger clusters will grow faster. The intensity of this positive feedback is inhomogeneous in the arena. Agents in clusters closer to the global optimum have longer waiting times. These clusters will exist longer than those that are farther away from the global optimum. Hence, the chance of growing is bigger for clusters closer to the global optimum. This process, typically, generates one big cluster close to the global optimum. The agents interact only locally and as noted above a BEECLUST-controlled swarm breaks symmetries. Hence, this behavior is different from other aggregation processes, for example, star formation which includes global interactions due to gravitation.

In the following experiments, the agents have initially random headings, are in the state ‘moving’, and are uniformly randomly distributed in the arena. The scalar field is bimodal with maxima of the same value and shape (see contours in Fig. 2). See Table 1 for the standard parameters used.

3.2 Clustering

The second scenario that we investigate is simple clustering of agents. It is similar to the BEECLUST-controlled swarm. However, the agents have constant waiting times independent of a scalar field, they move on a torus, and we introduce noisy sensors. The missing scalar field and missing wall effects remove the preferred regions of cluster formation from the system as seen in the BEECLUST scenario. The used parameters are the same as given above in Table 1. We implement noise in the sensors in a simple way by an agent-agent-recognition probability γ , that is, even if two agents are mutually within sensor range they might not perceive each other and might not stop. We set $\gamma = 0.75$, that is, an agent will perceive another agent in its sensor range in 75% of the time. The noise is uncorrelated in time. Noise was introduced because deterministic clustering seems to introduce artifacts into the entropy production histograms.

4 Analysis of scenarios

4.1 Analysis of BEECLUST

We consider a system of N agents that move in a two-dimensional box and scalar field. The particles are always found in one of two states: moving with constant velocity or stopped². We define the following equations of motion for each agent i

$$\dot{\mathbf{q}}_i = \mathbf{p}_i/m, \tag{6}$$

$$\dot{\mathbf{p}}_i = \mathbf{F}_i + \begin{cases} -\mathbf{p}_i, & \text{particle autonomously stops} \\ \mathbf{p}'_i, & \text{particle autonomously starts,} \\ 0, & \text{else} \end{cases} \tag{7}$$

where $\mathbf{q}_i = (x_i, y_i)^\top$ is the position of agent i , \mathbf{p}_i is the momentum, and \mathbf{p}'_i is the value of \mathbf{p}_i at the time the agent had stopped. Whenever the agent moves undisturbed by walls or by other agents we have $|\mathbf{F}_i| = 0$. If

²In any embodied implementation of this system (e.g., robots) the particles would be affected by friction and would consequently need to have a permanent acceleration compensating friction. This, in turn, means they would have an energy reservoir (cf. active particles [22]) and would permanently dissipate heat which would result in a situation as shown in Fig. 1. Energy costs have to be paid to allow self-organization and to comply with the second law of thermodynamics. Henceforth, we carry out the separation between these two sub-systems: the self-organizing sub-system containing the agents and the sub-system typified by the heat reservoir. Due to its energy dissipation the self-organizing sub-system does not have to obey the second law of thermodynamics.




Figure 4: Distribution of the entropy production for a swarm controlled by the BEECLUST algorithm, $t = 1500$, $\langle t\Omega_b \rangle \approx 15.77$, $T = 909.1$, number of samples $n \approx 5.0 \times 10^6$.

the agent bounces off the bounds or closely approaches another agent, we have a force $|\mathbf{F}_i| > 0$ that separates agents from each other or in case of approaching a wall it implements the regular behavior of a billiard ball (angle of incidence equals angle of reflection). This can be implemented, for example, via a WCA potential (Weeks-Chandler-Andersen potential, see [27]), which is a purely repulsive potential. As thermostat method we use velocity scaling, that is, based on the current \mathbf{p}_i we calculate a velocity scaling factor $\lambda \propto \sqrt{\beta}$ (cf. eq. 10). Consequently we can scale time $dt \propto 1/\sqrt{\beta}$ ($v = \frac{ds}{dt}$) which is then governed by the number of stopped agents. In particular, the special periods of time in which all agents are stopped are converted to time periods of no extend. Note that this is only our method of measuring the self-organizing system. It is not intrinsic to the system and the behavior of the agents is unconcerned by this method.

The system dynamics takes place in a high dimensional phase space $(\mathbf{q}_0, \mathbf{q}_1, \dots, \mathbf{q}_{N-1}, \mathbf{p}_0, \mathbf{p}_1, \dots, \mathbf{p}_{N-1}) \in \Gamma$. In the following we need to detect the essentials of this dynamics by a measure of entropy. First, we present an advantageous choice of an entropy measure and in Sec. 5.1 we give an alternative which is, however, less beneficial.

We ignore the momenta \mathbf{p} and also the y-positions because the main feature of the clusters is defined by the agents' x-positions (see Fig. 2). Ignoring the momenta does not hide entropy. Although we start with all nonzero momenta and during the experiments we have inhomogeneous momentum distributions but the experiments typically end with almost all agents stopped (i.e., a homogeneous momentum distribution). Similar to [5, Sec. 4.3] we measure the agent density modulation ρ_b of the BEECLUST scenario via

$$\rho_b(k, t) = \sum_{i=1}^N \sin(kx_i(t) + \frac{\pi}{2}), \quad (8)$$

where $x_i(t)$ is the x-position of agent i at time t , $k = 2\pi/L$, and $L = 150$ is the box length. The applied sine-function is shown in Fig. 2. Agents in the leftmost and rightmost quarters of the arena contribute positively, agents in the middle contribute negatively. In equilibrium, $x_i \in [0, L]$ is equally distributed averaged over many ensembles, yielding $\langle \rho_b \rangle = 0$. By applying the converse argument, averages of $\langle \rho_b \rangle \neq 0$ would correspond to unequal distributions of agents whereas negative and positive values indicate whether the main cluster is in the middle or at the ends.

Following Evans and Searles [5] we define a 'dissipation function' $\Omega_b(\Gamma)$ that gives the entropy production for a given phase space trajectory. We integrate changes of ρ_b over a time interval $[0, t]$

$$t\Omega_b = \beta \int_0^t \dot{\rho}_b(k, s) ds = \beta(\rho_b(k, t) - \rho_b(k, 0)) \quad (9)$$

and

$$\beta = 1/T = \left(\frac{2}{k_B N_d} \left\langle \sum_{i \in [0, N-1]} \frac{\mathbf{p}_i^2}{2m} \right\rangle \right)^{-1} \quad (10)$$

is the reciprocal temperature of the initial ensemble with Boltzmann constant k_B and degrees of freedom $N_d = 2N$. The distribution of the entropy production for $N = 25$ agents controlled by the BEECLUST algorithm, which were initially uniformly distributed, is shown in Fig. 4 for $t = 1500$. The initial random uniform distribution




Figure 5: Test of the entropy production distribution of the BEECLUST-controlled swarm shown in Fig. 4 against the fluctuation theorem (eq. 11), $Y = \frac{1}{\beta t} \ln \frac{P[\rho_b(k,t) - \rho_b(k,0) = A]}{P[\rho_b(k,t) - \rho_b(k,0) = -A]}$, $t = 1500$, $T = 909.1$. Note that any $Y \neq 0$ corresponds to negative entropy production.

yields $\langle \rho_b(0) \rangle = 0$ which is the state of maximal entropy. Hence, any distribution of the entropy production with a mean of $\langle t\Omega_b \rangle \neq 0$ indicates negative entropy production (i.e., averaged differences of the density modulation can have negative or positive signs but imply negative entropy production if they are nonzero). The ensemble average is $\langle t\Omega_b \rangle \approx 15.77$ which means that negative entropy is produced (initially at maximum entropy). Note that there is no direct influence by the scalar field to the entropy productions which are defined based on the agents' x-positions. Furthermore, the waiting times, that are determined by the scalar field, vary only by a factor of 5 between the minimum and the maximum.

Now we want to apply the fluctuation theorem (eq. 3) to this system. Especially we have to assume time-reversibility which is problematic because BEECLUST-controlled systems are in general not reversible [10]. However, we argue that it is fair to assume approximate reversibility because the irreversibility vanishes, if the agents aggregate, measure almost equal scalar field values (typically the difference is only about $\pm 10\%$) determining almost equal waiting times and almost equal wake-ups. Hence, in a time-reversed setting there would predominantly be valid cause-and-effect processes. It might appear counterintuitive to distinguish between $\rho_b(k,t) - \rho_b(k,0) = A$ and $\rho_b(k,t) - \rho_b(k,0) = -A$ because here both refer initially to negative entropy production. This indistinguishability, however, holds only for the initial phase. In later phases, $\rho_b(k,t) - \rho_b(k,0) > 0$ and $\rho_b(k,t) - \rho_b(k,0) < 0$ are to be distinguished because the system has in average a bias to non-zero values of $\rho_b(k,t)$, in particular positive values of $\rho_b(k,t)$ (i.e., the system operates most of the time in the positive half of Fig. 8). Consequently, a change of $\rho_b(k,t_1) - \rho_b(k,t_0) < 0$ takes the system back towards the maximal entropy state for $\rho_b(k,t_0) > 0$ corresponding to positive entropy productions (and for $\rho_b(k,t_1) - \rho_b(k,t_0) > 0$ vice versa) as discussed in Sec. 2. Applying the fluctuation theorem gives

$$\lim_{t \rightarrow \infty} \frac{1}{t} \ln \frac{P[\rho_b(k,t) - \rho_b(k,0) = A]}{P[\rho_b(k,t) - \rho_b(k,0) = -A]} = \beta A. \quad (11)$$

The data shown in Fig. 4 is tested whether it obeys eq. 11 in Fig. 5. The fluctuation theorem is satisfied for this system although the system is producing negative entropy and actually abandoning the equilibrium to which it was initialized. Hence, one could speak of an 'inverted fluctuation theorem' that is satisfied here.

In the following we want to investigate how it is possible for this self-organizing system to produce negative entropy. We hypothesize that the negative entropy production is based on fluctuations and the stopping behavior of the agents, hence, a process of frozen accidents.

We start our analysis with a measurement of the entropy production within a limited time interval [$t_0 = 15$, $t_1 = 20$] in the early transient. In addition, we classify for each measurement whether at least one agent changed its state from moving to stopped (starting agents do not occur that early in the simulation). The entropy production distributions for these two classes are shown in Fig. 6. For the measurements without a stopping agent the averaged change in the density modulation is about zero ($\langle (t_1 - t_0)\Omega_b \rangle \approx 0.06$). In contrast, for those measurements with stopping agents the averaged change of density modulation is negative ($\langle (t_1 - t_0)\Omega_b \rangle \approx -3.09$) indicating frozen accidents. For much later time intervals no difference between measurements with stopped and without a stopping agents are found. The negative value of $\langle (t_1 - t_0)\Omega_b \rangle$ demands for clarification because in the limit $t \rightarrow \infty$ the average density modulation is positive.

The explanation is based on a special feature of the BEECLUST-controlled swarm in this scenario which consists of three phases (see Fig. 7). In the short period before the first cluster forms, the average entropy production is $\langle \Omega_b = 0 \rangle$ indicating that the original fluctuation theorem holds for this phase. The first cluster usually does

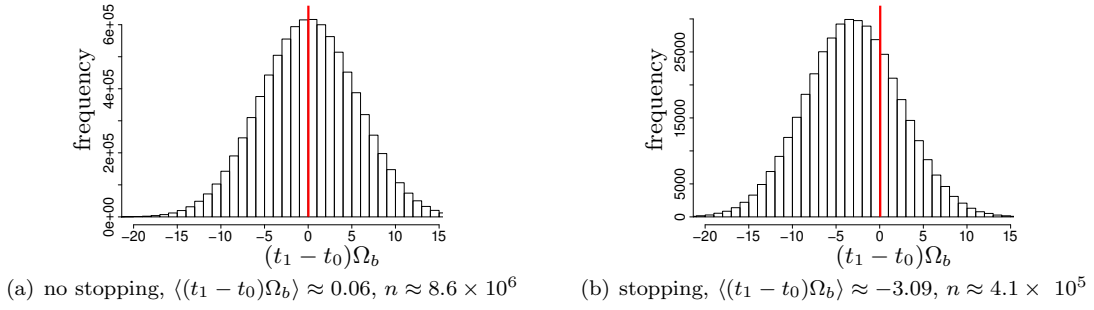


Figure 6: Distributions of the entropy production for the BEECLUST-controlled swarm and an early time interval during the transient ($t_0 = 15$, $t_1 = 20$, $T = 909.1$) classified according to whether a stopping agent was observed during the measurement.

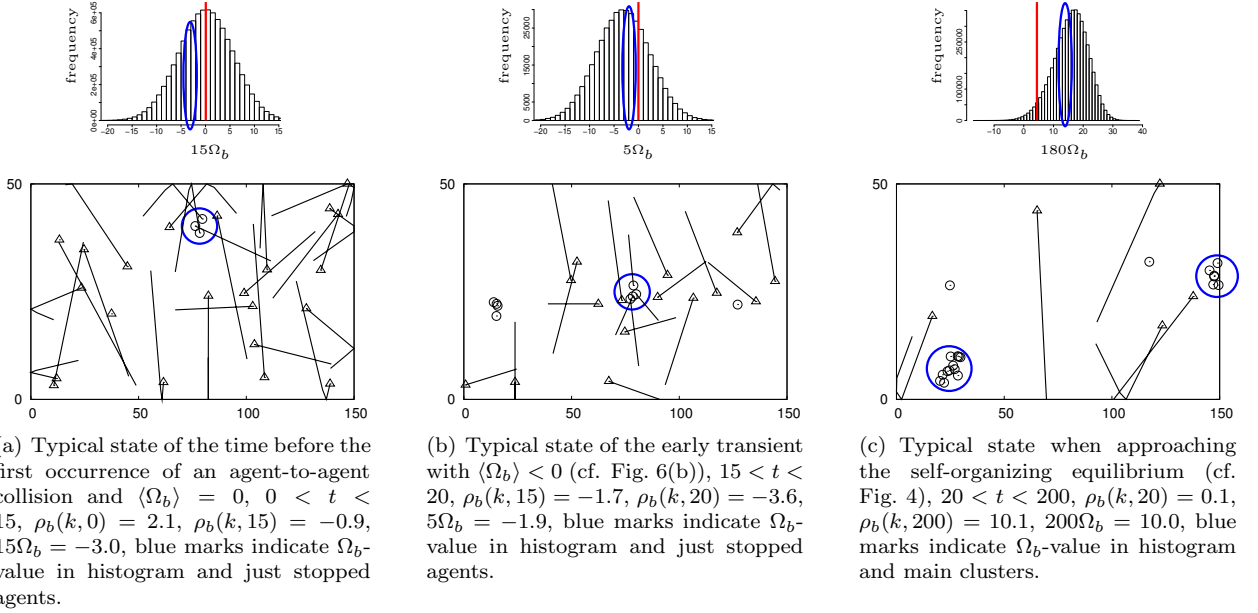


Figure 7: The three phases observed in the BEECLUST scenario each with a representative entropy production histogram and a plot of the arena showing moving (triangles) and stopped agents (circles) with a line indicating their most recent trajectory (histograms are meant to be qualitative).

not form close to the global optima but relatively close to the middle of the arena, see Fig. 7(a). In this area the agent density modulation (eq. 8) contributes negatively. In a second phase the average density modulation is negative ($\langle\Omega_b < 0\rangle$) because the density close to the middle of the arena increases further, see Fig. 7(b). This is also indicated by the evolution of the agent density modulation over time as shown in Fig. 8. Initially it stays close to 0 and only later it clearly takes a positive sign. The insert shows details of the first 50 time steps and indicates negative slope for the time interval [15, 20] (i.e., second phase) of Fig. 6. Only later the clusters ‘move’ towards the ends of the arena probably aided by wall effects (agents are more likely to approach the cluster again after having left it on the side of the cluster that is closer to one of the ends of the arena), see Fig. 7(c) and consequently the average density modulation is positive ($\Omega_b > 0$).

4.2 Analysis of clustering

The clustering scenario is simple and it serves in this paper as a second example to show the generality of our approach. Hence we reduce the analysis to a comparison of the entropy production and the test of the entropy production distribution to those obtained in the BEECLUST scenario.




Figure 8: Evolution of the agent density modulation over time in the BEECLUST scenario, black line shows ensemble average, gray lines show samples, insert shows details of the ensemble average within the first 50 time steps.

In analogy to the above density modulation we use a function

$$\rho_c(d_{\max}, t) = \frac{1}{N^2 d_{\max}} \sum_{i=1}^N \sum_{j=1}^N \text{dist}(\mathbf{q}_i(t), \mathbf{q}_j(t)) \quad (12)$$

defined by a normalized sum of all entries of the distance matrix to construct the dissipation function. The function $\text{dist}(\cdot, \cdot)$ gives the distance between two points on a torus and d_{\max} is the maximal distance between two points on the torus. Subsequently the dissipation function is

$$t\Omega_c = \int_0^t \dot{\rho}_c(k, s) ds = \rho_c(d_{\max}, t) - \rho_c(d_{\max}, 0). \quad (13)$$

We ignore any additional constants for this scenario. The agents are initialized to uniformly randomly distributed positions on the torus. In average this uniform distribution corresponds to high entropy which is reflected by high values of ρ_c . In average these high values of ρ_c cannot be increased by the defined swarm behavior. While forming clusters, distances between several agents decrease, and consequently ρ_c will decrease. Hence, we expect negative values of $t\Omega_c$ for the majority of runs. After $t = 80$ time steps the entropy production $t\Omega_c$ is measured. We choose this relatively early state because the system converges fast and consequently only few positive entropy productions are observed later which complicates the statistical analysis. The obtained distribution is shown in Fig. 9(a). In contrast to Fig. 4 the distribution is clearly asymmetrical, the average $\langle t\Omega_c \rangle \approx -0.015$ is closer to zero (median: -0.0091), and the average is negative indicating negative entropy production as expected.

The test of the entropy production distribution according to the fluctuation theorem and in analogy to eq. 11 is shown in Fig. 9(b). The data is close to linearity but shows systematic deviations (too small for $A < 0.01$ and too big for $0.15 < A < 0.03$). Given the fast convergence of the system to a majority of trajectories with negative entropy productions this is still a satisfactory result.

5 Generality

To support our claim that this approach bears the potential of general applicability, we report the influence of an alternative entropy measure and investigate the influence of temperatures in the following.

5.1 Independence from entropy measure

The entropy measure that might be the most obvious to apply in the context of this paper is probably the information theoretic entropy of Shannon [23]. Shannon entropy, however, is discrete and our system has many continuous features besides the discrete concept of agent numbers. The continuous extension of Shannon entropy (differential entropy) is also not directly applicable because we would have to sample a probability distribution from only N samples represented by agent positions. To get from continuous state space to discrete microstates we have to do some form of coarse graining. Unfortunately there is an ambiguity because there are many

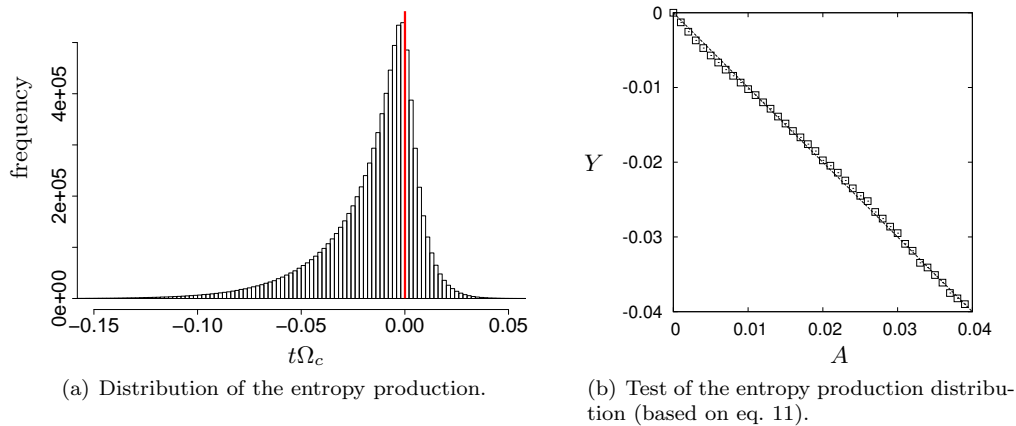


Figure 9: Distribution and test of the entropy production for the clustering swarm on a torus (described in Sec. 3.2), $t = 80$, $\langle t\Omega_c \rangle \approx -0.015$, $T = 3.1$, number of samples $n \approx 8.5 \times 10^6$.

possibilities to implement coarse graining. Still, at least two conditions can be fixed. First, the dynamics of the obtained entropy measure should parallel those of the system and, second, it should catch its main features.

We implement the coarse graining by a standard method. We put a grid over the arena with tile side length 10 generating 75 tiles (cf. Table 1). Hence, we do not ignore the y-positions as above but still we ignore the momenta. We count two micro-states: unoccupied and occupied tiles (ignoring the actual number of agents in one tile). This way we sample two probabilities $p_0(t)$ and $p_1(t) = 1 - p_0(t)$. The Shannon entropy is given by $S(t) = -\sum_i p_i(t) \log_2 p_i(t)$ and the entropy production during a time interval $[0, t]$ is given by

$$t\Omega_c = \beta \int_0^t \dot{S}(u) du = \beta(S(t) - S(0)). \quad (14)$$

Initially the agents are uniformly randomly distributed which causes about as many tiles being occupied as there are agents (for $N \ll 75$) and we get $p_1(0) \approx N/75$. After some time the agents cluster and the number of occupied tiles typically decreases (a typical value is $p_1(t) = 5/75$ for $N = 25$). Hence, the average entropy production will be negative ($\langle t\Omega_c \rangle < 0$).

The above experiments can be repeated with this entropy measure and no qualitative change is detectable. However, there are some technical problems. Due to the discrete measure, abnormalities in the histograms occur and the number of trajectories showing positive entropy production decreases fast with increasing time intervals which complicates the statistical analysis for which we need well filled histogram bins on both the positive and the negative side. Hence, the obtained data will not fully converge to a line according to eq. 11 and we cannot sample within long time intervals. The results are shown in Fig. 10. The shapes of the curves are systematic due to coarse graining. Still, the overall trend to linearity is seen. We conclude that the above entropy measure based on a density modulation was a better choice.

5.2 Influence by temperature

Both investigated scenarios are certainly more complex than a mere physical multi-particle system because they incorporate also autonomous stops of agents, sensor ranges, and waiting times in addition to particle velocities. Still, we can catch most of them with the concept of temperature as defined in eq. 10. When we reduce the agents' velocities we alter the temperature. However, we have to extend the concept of temperature to allow the comparison of systems at different agent velocities. To scale the system behaviors at different agent velocities completely, it is clear that we also have to scale waiting times and sensor ranges which are not considered in eq. 10. Hence, we have to apply an extended concept of temperature which includes these velocity-dependent properties of agents as well. However, the definition and analysis of a general temperature concept for multi-agent systems is beyond this paper. Here, we just scale speeds (halved), waiting times (doubled), and sensor ranges (halved) and obtain two systems with temperatures $T_1 = 2T_2$. We compare the entropy production distributions of these two systems in Fig. 11 for a late state at time $t = 6000$. The system at half temperature (T_2) shows systematic deviations because it converges slower. Still, the deviations are small enough to conclude that at least for some scenarios an extended concept of temperature can be found that allows to scale these systems and compare them across different temperatures.




Figure 10: Test of the entropy production distribution for Shannon entropy according to eqs. 3 and 14 for the BEECLUST scenario in analogy to Fig. 5 for different times $t \in \{60, 100, 1000, 2000\}$; for $t \geq 1000$ only values for $A \leq 0.07$ are given because for $A > 0.07$ fast convergence does not allow a statistically sound analysis, number of samples $> 1.5 \times 10^6$ each.

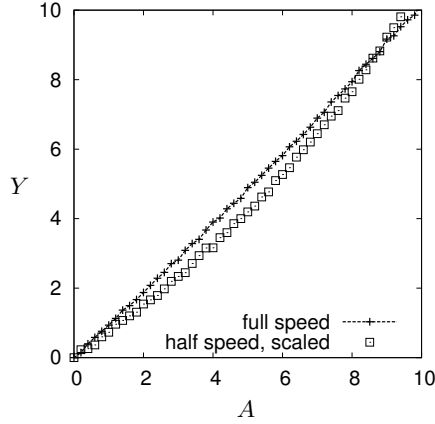


Figure 11: Comparison of systems at two different temperatures $T_1 = 2T_2$ for the BEECLUST scenario; The first system is parameterized as above (Table 1). The second system has halved speeds, doubled waiting times, and halved sensor ranges. The plot shows $Y = \frac{T_1}{t} \ln \frac{P[\rho_b^{T_1}(k,t) - \rho_b^{T_1}(k,0) = A]}{P[\rho_b^{T_1}(k,t) - \rho_b^{T_1}(k,0) = -A]} \approx \frac{2T_2}{t} \ln \frac{P[\rho_b^{T_2}(k,t) - \rho_b^{T_2}(k,0) = A]}{P[\rho_b^{T_2}(k,t) - \rho_b^{T_2}(k,0) = -A]}$; number of samples $n \approx 2.5 \times 10^6$ each, $t = 6000$.




Figure 12: Sample run of a simple model based on summations of $N = 25$ random processes initialized to $X_i(0) = 0$, based on normally distributed random variables ($\mu = 0$, $\sigma^2 = 1$); a random process i becomes inactive once $X_i(t) = 5$ is reached, N_a gives the number of active random processes.

6 Discussion

Note again that $\rho_b(k, t) = 0$ corresponds to maximum entropy. Therefore, any $\rho_b(k, t) \neq 0$ in Fig. 8 indicates negative entropy production. We conclude that the negative entropy production of this system is initiated by entropy fluctuations which are normally distributed and are negative/positive with about the same probability respectively according to the original fluctuation theorem and as seen in Fig. 6(a). Some of the ‘negative entropy production’-events are locally observable by the agents themselves because there are simultaneous agent-to-agent encounters of three agents with mutual perception. This local perception of the global measure of entropy is leveraged by stopping all three agents and consequently storing the local entropy fluctuation. Cascades of such stopping behaviors generate a positive feedback (self-amplification of fluctuations as in Rayleigh-Bénard convection). In the end a system dynamics is generated that can be described by an inverted fluctuation theorem, which dictates an exponentially increasing probability of low entropy states. Hence, this emergent self-organizing swarm does indeed rely on frozen accidents. Note that the overall system (i.e., including the heat bath) still produces positive entropy (e.g., due to accelerations of the agents) while the agent-position-based entropies are only reduced in the self-organizing sub-system.

The effectiveness of the frozen-accidents concept can easily be made clear by constructing a simple model. We represent the entropy contribution of each agent i by a random process $X_i(t)$. The total entropy is just the sum $\sum_{i=0}^N X_i(t)$ over all agents N . The restriction of all random processes to the interval $[-5, 5]$ is essential and we define $X_i(t) = 5$, $\forall t > t_0$ with t_0 is the first time agent i achieved $X_i(t_0) = 5$. That is, once a random process reaches $X_i(t_0) = 5$ (a local property) it stays there forever—a frozen accident. As a consequence the number of active random processes N_a will decrease monotonically. A sample run of this simple model for $N = 25$ based on Gaussian distributed X_i and initialization $X_i(0) = 0$ is shown in Fig. 12. The bias in the otherwise random trajectory is noticeable. Note that the summation of Gaussian distributed random variables $\sum_N X_i$ with each having a variance of σ_i^2 results in a random variable that is also Gaussian distributed with a variance of $\sigma^2 = \sum_N \sigma_i^2$. With decreasing number of active processes N_a more and more variances vanish ($\sigma_i^2 = 0$). Hence, also the variance of the sum will decrease which is the macroscopic effect of the frozen accidents and ensures that states of low entropy are much more likely to be maintained. The challenge in investigating self-organizing systems is, however, represented by the reverse engineering of such a model. One would start with an observed macroscopic effect as shown in Fig. 12 and then decompose it into a set of sub-systems which accidentally achieve low-entropy states by noise and then preserved them by ‘freezing’-events.

The reported results indicate that this emergent self-organized system obeys an inversion of the fluctuation theorem which could be stated as

$$\lim_{t \rightarrow \infty} \frac{1}{t} \ln \frac{P[\overline{\Sigma}_t = -A]}{P[\overline{\Sigma}_t = A]} = A, \quad (15)$$

following eq. 3. We get an immediate interpretation of this self-organizing system by inverting the interpretation of the fluctuation theorem. A self-organizing system that is started with high entropy will produce negative entropy with an exponentially increasing probability over time. As a consequence there is a ‘self-organization equilibrium’ of lower entropy to which the system will converge. As a second consequence the self-organizing entropy-reduction behavior is a transient phenomenon, cf. [18, p. 62].

7 Conclusion

In this paper, we have analyzed emergent self-organizing multi-agent (or swarm) systems with methods based on and suggested by the fluctuation theorem. The results provide empirical evidence for the existence of an inverted fluctuation theorem that applies for such dissipative self-organizing systems. A theory based on an inverted fluctuation theorem could prepare a wide basis for the analysis of self-organizing systems. We claim these methods have a potential for general applicability. This claim of generality is supported by the reported results with an alternative entropy measure and the results using a broader interpretation of the concept of temperature.

Specific exemplary benefits of such a theory could be the definition of preconditions for self-organization, for example, concerning the cognitive abilities of the agents. Statistical properties of fluctuations describe the time-scales on which negative entropy productions can be observed locally. The agents need to perceive local samples of this global property of negative entropy production and need to react within these time-scales. Hence, conditions for controller sampling rates could be derived. The agents need appropriate sensors that allow local measurements of entropy with an accuracy that is sufficiently higher than the rate at which events of negative entropy production occur. Thus, preconditions for successfully generating positive feedbacks could be derived.

Especially the origin of BEECLUST confirms the possibility of applying the proposed methods to natural systems such as clustering behaviors in young honey bees [25] or other social insects, as well as flocks, herds, and shoals. Hence, the same methods could be used for artificial and natural systems which could, in turn, enrich primarily biological studies.

This work proved again that thermodynamics and statistical physics offer many fully developed methods which can often be applied even unmodified to problems of emergent behavior (cf. Hamann et al. [10]). Pursuing this research track might be a promising way of achieving general insights to still rather fuzzy concepts such as emergence or self-organization.

Finally, it is clear that the reported approach is truly interdisciplinary in combining methods and problems from physics, biology, and computer science. It is obvious that, at least in the field of artificial life, any future research success has to be founded on a collection of several scientific fields. In our future work, we hope to continue this approach by generalizing the concept of an inverted fluctuation theorem for emergent self-organizing multi-agent systems.

Acknowledgments

The authors thank Payam Zahadat and the anonymous reviewers of the preceding paper at the ECAL'11 conference for helpful comments that improved this paper. This reported research is supported by: EU-ICT 'CoCoRo', no. 270382; EU-IST-FET project 'SYMBRION', no. 216342; EU-ICT project 'REPLICATOR', no. 216240; FWF (Austrian Science Fund) 'REBODIMENT', no. P23943-N13; Austrian Federal Ministry of Science and Research (BM.W.F).

References

- [1] Andersen, H. C. (1980). Molecular dynamics simulations at constant pressure and/or temperature. *Journal of Chemical Physics*, 72(4):2384–2393.
- [2] Bodenschatz, E., Pesch, W., and Ahlers, G. (2000). Recent developments in Rayleigh-Bénard convection. *Annual review of fluid mechanics*, 32(1):709–778.
- [3] Crick, F. H. (1968). The origin of the genetic code. *Journal of Molecular Biology*, 38(3):367–379.
- [4] Evans, D. J., Cohen, E. G. D., and Morriss, G. P. (1993). Probability of second law violations in shearing steady states. *Physical Review Letters*, 71:2401–2404.
- [5] Evans, D. J. and Searles, D. J. (2002). The fluctuation theorem. *Advances in Physics*, 51(7):1529–1585.
- [6] Gell-Mann, M. (1995). Plectics. In Brockman, J., editor, *The Third Culture: Beyond the Scientific Revolution*, pages 316–332. Touchstone Press, New York.
- [7] Gerstner, E. (2002). Second law broken: Small-scale energy fluctuations could limit minaturization. Nature online news.

- [8] Haken, H. (1977). *Synergetics - an introduction*. Springer-Verlag, Berlin, Germany.
- [9] Hamann, H., Meyer, B., Schmickl, T., and Crailsheim, K. (2010). A model of symmetry breaking in collective decision-making. In Doncieux, S., Girard, B., Guillot, A., Hallam, J., Meyer, J.-A., and Mouret, J.-B., editors, *From Animals to Animats 11*, volume 6226 of *Lecture Notes in Artificial Intelligence*, pages 639–648, Berlin, Germany. Springer-Verlag.
- [10] Hamann, H., Schmickl, T., and Crailsheim, K. (2011). Thermodynamics of emergence: Langton’s ant meets Boltzmann. In *IEEE Symposium on Artificial Life (IEEE ALIFE 2011)*, pages 62–69. IEEE.
- [11] Hamann, H., Schmickl, T., Wörn, H., and Crailsheim, K. (2012). Analysis of emergent symmetry breaking in collective decision making. *Neural Computing & Applications*, 21(2):207–218.
- [12] Hereford, J. M. (2011). Analysis of BEECLUST swarm algorithm. In *Proc. of the IEEE Symposium on Swarm Intelligence (SIS 2011)*, pages 192–198. IEEE.
- [13] Jaynes, E. T. (1992). The Gibbs paradox. In Smith, C. R., Erickson, G. J., and Neudorfer, P. O., editors, *Maximum Entropy and Bayesian Methods*, pages 1–22. Kluwer Academic, Dordrecht.
- [14] Maxwell, J. C. (1878). Tait’s ‘Thermodynamics’ (I). *Nature*, 17:257–259.
- [15] Nicolis, G. and Prigogine, I. (1977). *Self-organization in nonequilibrium systems*. Wiley, New York.
- [16] Parunak, H. V. D. and Brueckner, S. (2001). Entropy and self-organization in multi-agent systems. In André, E., Sen, S., Frasson, C., and Müller, J. P., editors, *Proceedings of the fifth international conference on Autonomous agents (AGENTS’01)*, pages 124–130, New York, NY, USA. ACM Press.
- [17] Polani, D. (2008). Foundations and formalizations of self-organization. In Prokopenko, M., editor, *Advances in Applied Self-organizing Systems*, Advanced Information and Knowledge Processing. Springer-Verlag.
- [18] Prigogine, I. (1997). *The End of Certainty: Time, Chaos, and The New Laws of Nature*. Free Press.
- [19] Schmickl, T. and Hamann, H. (2011). BEECLUST: A swarm algorithm derived from honeybees. In Xiao, Y., editor, *Bio-inspired Computing and Communication Networks*. CRC Press.
- [20] Schmickl, T., Hamann, H., Wörn, H., and Crailsheim, K. (2009). Two different approaches to a macroscopic model of a bio-inspired robotic swarm. *Robotics and Autonomous Systems*, 57(9):913–921.
- [21] Schmickl, T., Thenius, R., Möslinger, C., Radspieler, G., Kernbach, S., and Crailsheim, K. (2008). Get in touch: Cooperative decision making based on robot-to-robot collisions. *Autonomous Agents and Multi-Agent Systems*, 18(1):133–155.
- [22] Schweitzer, F. (2003). *Brownian Agents and Active Particles. On the Emergence of Complex Behavior in the Natural and Social Sciences*. Springer-Verlag, Berlin, Germany.
- [23] Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27:79–423.
- [24] Smith, E. (2008). Thermodynamics of natural selection I: Energy flow and the limits on organization. *Journal of Theoretical Biology*, 252(2):185–197.
- [25] Szopek, M., Radspieler, G., Schmickl, T., Thenius, R., and Crailsheim, K. (2008). Recording and tracking of locomotion and clustering behavior in young honeybees (*Apis mellifera*). In Spink, A. J., Ballintijn, M. R., Bogers, N. D., Grieco, F., Loijens, L. W. S., Noldus, L. P. J. J., Smit, G., and Zimmerman, P. H., editors, *Proceedings of Measuring Behavior 2008, Maastricht, August 26-29*, volume 6, page 327.
- [26] Wang, G. M., Sevcik, E. M., Mittag, E., Searles, D. J., and Evans, D. J. (2002). Experimental demonstration of violations of the second law of thermodynamics for small systems and short time scales. *Physical Review Letters*, 89(5):050601.
- [27] Weeks, J. D., Chandler, D., and Andersen, H. C. (1971). Role of repulsive forces in determining the equilibrium structure of simple liquids. *Journal of Chemical Physics*, 54(12):5237.
- [28] Wolf, T. D. and Holvoet, T. (2005). Emergence versus self-organisation: Different concepts but promising when combined. In Brueckner, S., Serugendo, G. D. M., Karageorgos, A., and Nagpal, R., editors, *Proceedings of the workshop on Engineering Self Organising Applications*, volume 3464 of *Lecture Notes in Computer Science*, pages 1–15. Springer-Verlag.
- [29] Wu, M., Ahlers, G., and Cannell, D. S. (1995). Thermally induced fluctuations below the onset of Rayleigh-Bénard convection. *Physical Review Letters*, 75(9):1743–1746.