Explaining Emergent Behavior in a Swarm System Based on an Inversion of the Fluctuation Theorem

Heiko Hamann[1], Thomas Schmickl[1], Karl Crailsheim[1]

[1] Artificial Life Laboratory of the Department of Zoology, Karl-Franzens University Graz, Universitätsplatz 2, A-8010 Graz, Austria
heiko.hamann@uni-graz.at

Abstract

A grand challenge in the field of artificial life is to find a general theory of emergent self-organizing systems. In this paper we try to explain the emergent behavior of a simulated swarm by applying methods based on the fluctuation theorem. Empirical results indicate that the swarm is able to produce negative entropy within an isolated sub-system due to 'frozen accidents'. Individuals of the 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.

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 (2001): "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 (2001), 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[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 (Polani, 2008; Hamann et al., 2011a).

[1] In principle, we refer here to Gibbs entropy $S = -k_B \sum_i p_i \ln p_i$, for Boltzmann constant k_B and the sum over all microstates with probabilities p_i which applies especially to classical, finite systems far away from equilibrium. However, an intuitive understanding of entropy suffices in the following.

The emergence of life 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 (1968) phrased the term 'Frozen Accident Theory'. While Crick was introducing this concept with focus on genetics, Gell-Mann (1995) 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'.

The intuition of this theory is relatively clear in the context of the slow evolution of our universe. However, we want to define a concept of frozen accidents within emergent self-organizing multi-agent systems (De Wolf and Holvoet, 2005) 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 (Evans et al., 1993) are also possible as, for example, indicated by Maxwell (1878):

> 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 (Wang et al., 2002). 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. The second law is only statistical and, hence, allows spontaneous decreases of entropy in isolated systems with nonzero probability.

The possibility of temporal entropy decreases 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 (at equilibrium state) turns out to be Gaussian with mean $\mu = 0$:

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

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 (Evans and Searles, 2002; Evans et al., 1993) quantifies the probability of violations to the second law. For short intervals it can be said that nature was running in reverse. Even concerning living systems this might be true. 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. (2002), Gerstner (2002) 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 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 (2008) when stating

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

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 (2008) 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 (Bodenschatz et al., 2000), for example, initially fluctuating flows (Wu et al., 1995) occur that are enhanced and trigger the formation of Bénard cells in spontaneous symmetry breaking, cf. also (Nicolis and Prigogine, 1977; Haken, 1977). We want to point out the self-amplification of fluctuations as such a sufficient condition here.

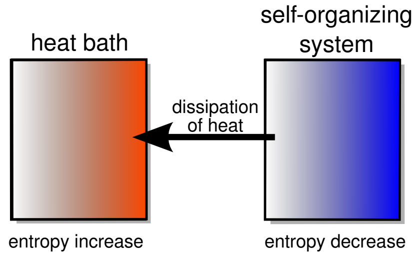

Figure 1: Schematic of a system divided into a heat bath with increasing entropy and a self-organizing, dissipative sub-system with decreasing entropy.

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 (Hamann et al., 2011a).

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 cognition 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 investigated scenario and the fluctuation theorem. 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.

BEECLUST algorithm

The BEECLUST algorithm can be considered a model algorithm for swarms. It is based on observations of young honeybees (Szopek et al., 2008), was analyzed in many models (Hereford, 2011; Schmickl and Hamann, 2011; Schmickl et al., 2009; Hamann et al., 2011b, 2010), and even implemented in a swarm of robots (Schmickl et al., 2008).

This algorithm allows a swarm to aggregate at a maximum of a gradient field although individual agents do not perform a greedy gradient ascent. Hence, it might be justified to

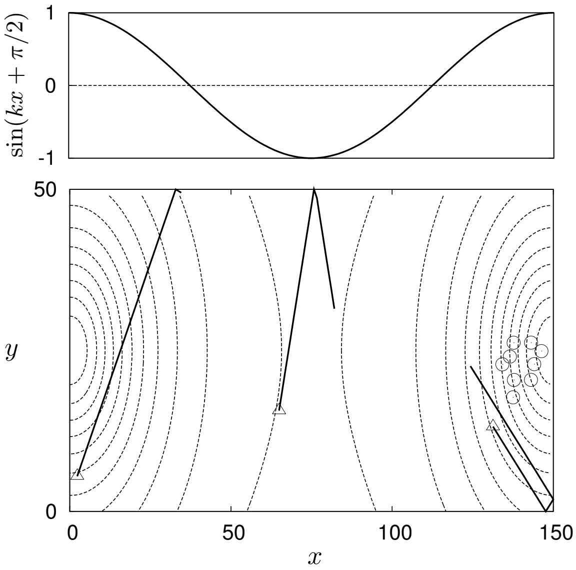

Figure 2: Bottom: Typical state of a swarm controlled by BEECLUST; positions of stopped agents (circles) and moving agents (triangles) with trajectories of the last 20 time steps, contours show levels of the gradient field. Top: function used in eq. 7.

```
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 gradient value.
    The higher the gradient 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).

call this emergent behavior. Controlled by this algorithm three agents will stop (note that in previous works typically a threshold of two was chosen, which is, however, irrelevant in this paper) when they approach each other, measure the local value of the gradient, 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 gradient field (see the lower part of Fig. 2). See Fig. 3 for a definition of the BEECLUST algorithm.

The collective aggregation close to the global optimum is achieved via a positive feedback process (Hamann et al., 2011b): Clusters of 3 stopped agents will form by chance anywhere 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 into a cluster of size 4 is bigger for clusters closer to the global optimum. 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. This process, typically, leads to just one big cluster close to the global optimum. The agents interact only locally and a BEECLUST-controlled swarm is able to break symmetries (Hamann et al., 2011b). Hence, this behavior is different from other aggregation processes, for example, star formation which includes global interactions due to gravitation.

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	3[length units]$/$[time units]
number of agents	25

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

Fluctuation Theorem

According to Evans and Searles (2002) the group of fluctuation theorems “gives an analytical expression for the probability of observing Second Law violating dynamical fluctuations in thermostatted dissipative non-equilibrium systems.” In 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. One of these theorems (steady state fluctuation theorems) applies to time-reversible, thermostatted, ergodic dynamical systems and defines the relation of fluctuations (Evans and Searles, 2002)

$$\lim_{t\to\infty} \frac{1}{t} \ln \frac{P[\overline{\Sigma}_t = A]}{P[\overline{\Sigma}_t = -A]} = A, \tag{3}$$

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

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

which states that the average over many ensembles, in which the time averaged entropy productions where measured, is positive. Hence, the fluctuation theorem is in accordance with the second law of thermodynamics.

Analysis of BEECLUST

We consider a system of N agents that move in a two-dimensional box and gradient field. We assume the particles to move frictionless which basically means they have a permanent acceleration compensating friction. This, in turn, means they have an energy reservoir (cf. active particles (Schweitzer, 2003)) and permanently dissipate heat which results in a situation as shown in Fig. 1. In addition, we allow infinite accelerations because the agents stop and start within one time step in our numerical simulation. Energy costs have to be paid to allow self-organization and to comply with the second law of thermodynamics. In the following 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. We define the following equations of motion for each agent i

$$\dot{\mathbf{q}}_i = \mathbf{p}_i/m, \tag{5}$$

$$\dot{\mathbf{p}}_i = \mathbf{F}_i + \begin{cases} -\mathbf{p}_i, & \text{particle autonomously stops} \\ \mathbf{p}'_i, & \text{particle autonomously starts} \\ 0, & \text{else} \end{cases}, \tag{6}$$

where $\mathbf{q}_i = (x_i, y_i)^\intercal$ 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. We have $\mathbf{F}_i > 0$ in case the agent bounces off the bounds or closely approaches another agent. This can be implemented, for example, via a WCA potential (Weeks et al., 1971), which is a purely repulsive potential. As thermostat method we use velocity scaling which is 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 $(\mathbf{q}_0, \mathbf{q}_1, \ldots, \mathbf{q}_{N-1}, \mathbf{p}_0, \mathbf{p}_1, \ldots, \mathbf{p}_{N-1}) \in \Gamma$. In the following we need to detect the essentials of this dynamics by a measure of entropy. 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

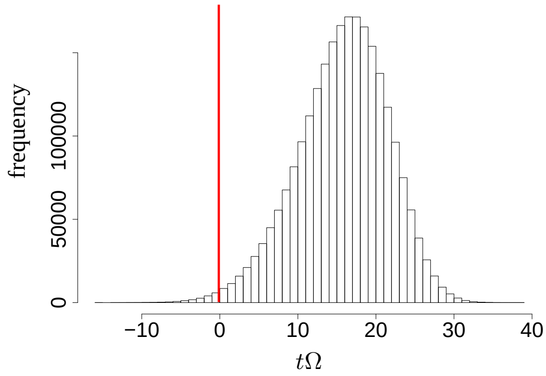

Figure 4: Distribution of the entropy production for a swarm controlled by the BEECLUST algorithm, $t = 1500$, $\langle t\Omega \rangle \approx 15.77$, $T = 909.1$, number of samples $n \approx 5.0 \times 10^6$.

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., again a homogeneous momentum distribution). Similar to (Evans and Searles, 2002, Sec. 4.3) we observe the agent density modulation via

$$\rho(k, t) = \sum_{i=1}^{N} \sin(k x_i(t) + \frac{\pi}{2}), \tag{7}$$

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, $x_i \in [0, L]$ would be equally distributed averaged over many ensembles, yielding $\langle \rho \rangle = 0$. By applying the converse argument, averages of $\langle \rho \rangle \neq 0$ 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 (2002) we define a 'dissipation function' $\Omega(\Gamma)$ that gives the entropy production for a given phase space trajectory. We integrate changes of ρ over a time interval $[0, t]$

$$t\Omega = \beta \int_0^t \dot{\rho}(k, s) ds = \beta(\rho(k, t) - \rho(k, 0)) \tag{8}$$

and

$$\beta = 1/T = \left(\frac{2}{k_B N_d} \left\langle \sum_{i \in [0, N-1]} \frac{\mathbf{p}_i^2}{2m} \right\rangle \right)^{-1} \tag{9}$$

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

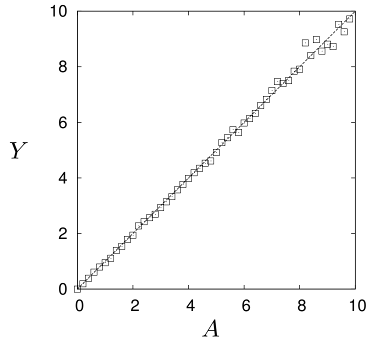

Figure 5: Test of the entropy production distribution of the BEECLUST-controlled swarm shown in Fig. 4 against the fluctuation theorem (eq. 10), $Y = \frac{1}{\beta t} \ln \frac{P[\rho(k,t)-\rho(k,0))=A]}{P[\rho(k,t)-\rho(k,0))=-A]}$, $t = 1500$, $T = 909.1$. Note that any $Y \neq 0$ corresponds to negative entropy production.

initially random uniformly distributed, is shown in Fig. 4 for $t = 1500$. The initial uniform distribution yields $\langle \rho(0) \rangle = 0$ which is the state of maximal entropy. Hence, any distribution of the entropy production with a mean of $\langle t\Omega \rangle \neq 0$ 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 $\langle t\Omega \rangle \approx 15.77$ which means that negative entropy is produced (initially at maximum entropy). Note that there is no direct influence by the gradient field to the entropy productions which are defined based on the agents' x-positions. Furthermore, the waiting times, that are determined by the gradient 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 (Hamann et al., 2011a). However, we argue that it is fair to assume approximate reversibility because the irreversibility vanishes, if the agents measure almost equal gradient values (typically the difference is only about $\pm 10\%$) determining almost equal waiting times and almost equal wake-ups. Applying the fluctuation theorem gives

$$\lim_{t\to\infty} \frac{1}{t} \ln \frac{P[\rho(k,t)-\rho(k,0)) = A]}{P[\rho(k,t)-\rho(k,0)) = -A]} = \beta A. \qquad (10)$$

The data shown in Fig. 4 is tested whether it obeys eq. 10 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. Note that such a mechanism is similar to the famous thought experiment 'Maxwell's Demon' (Maxwell, 1871). Furthermore, an implementation of Maxwell's Demon was reported (Bannerman et al., 2009) that is used as a cooling technique (cf. our metaphor of a heat pump in the introduction). Here we have rather a 'distributed demon' embodied by many autonomous agents that control themselves (Adami (1998) applies a similar argument to evolution). BEECLUST does not sort particles or agents as Maxwell's Demon but aggregate them (i.e., we generate uneven density distributions).

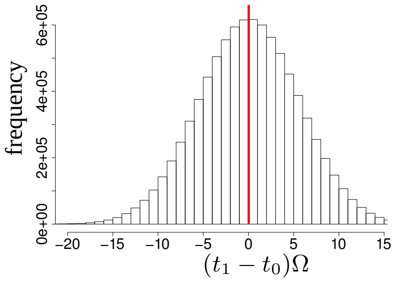

(a) no stopping, $\langle (t_1 - t_0)\Omega \rangle \approx 0.06$, $n \approx 8.6 \times 10^6$

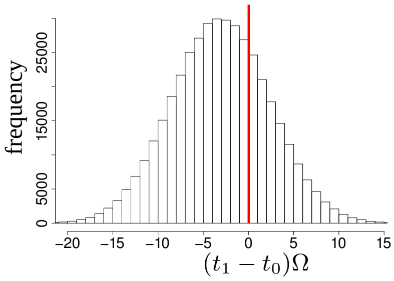

(b) stopping, $\langle (t_1 - t_0)\Omega \rangle \approx -3.09$, $n \approx 4.1 \times 10^5$

Figure 6: Distributions of the entropy production for 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.

We measure 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 distribution 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 0 ($\langle (t_1 - t_0)\Omega \rangle \approx 0.06$). In contrast, for those measurements with stopping agents the averaged change of density modulation is negative ($\langle (t_1 - t_0)\Omega \rangle \approx -3.09$) 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 $\langle (t_1 - t_0)\Omega \rangle$ demands for clarification because in the limit $t \to \infty$ the average density modulation is positive.

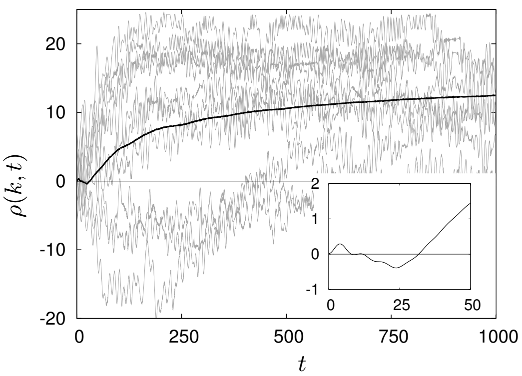

Figure 8: Evolution of the agent density modulation over time, black line shows ensemble average, gray lines show samples, insert shows details of the ensemble average within the first 250 time steps.

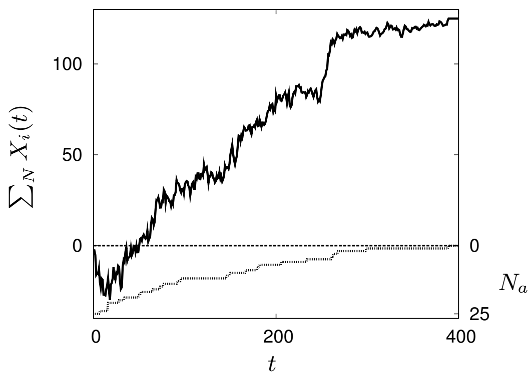

Figure 9: Sample run of a simple model based on summations of $N = 25$ random processes initialized to $X_i(0) = 0$ and based on normally distributed random variables ($\mu = 0$, $\sigma^2 = 1$).

The explanation is 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 $\Omega = 0$ indicating that the original fluctuation theorem holds for this phase. The first cluster usually does 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. 7) contributes negatively. In a second phase the average density modulation is negative ($\Omega < 0$) 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 250 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 due to wall effects, see Fig. 7(c) and consequently the average density modulation is positive ($\Omega > 0$).

Discussion

Note again that $\rho(k, t) = 0$ corresponds to maximum entropy. Therefore, any $\rho(k, t) \neq 0$ 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 these 'negative entropy production'-events are locally observable by the agents themselves because there are three agent-to-agent encounters with mutual perception. This local perception of the global measure of entropy is leveraged by stopping all three agents and stores 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 still produces positive entropy (e.g., due to accelerations of the agents) while the agent-position-based entropy is only reduced in the self-organizing sub-system.

The effectiveness of the frozen–accidents concept can easily be made clear by 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. 9. 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 results shown in Figs. 4, 5, and 6(b) indicate that this emergent self-organized system obeys an inversion of the fluctuation theorem which could be stated as

$$\lim_{t\to\infty} \frac{1}{t} \ln \frac{P[\overline{\Sigma}_t = -A]}{P[\overline{\Sigma}_t = A]} = A, \qquad (11)$$

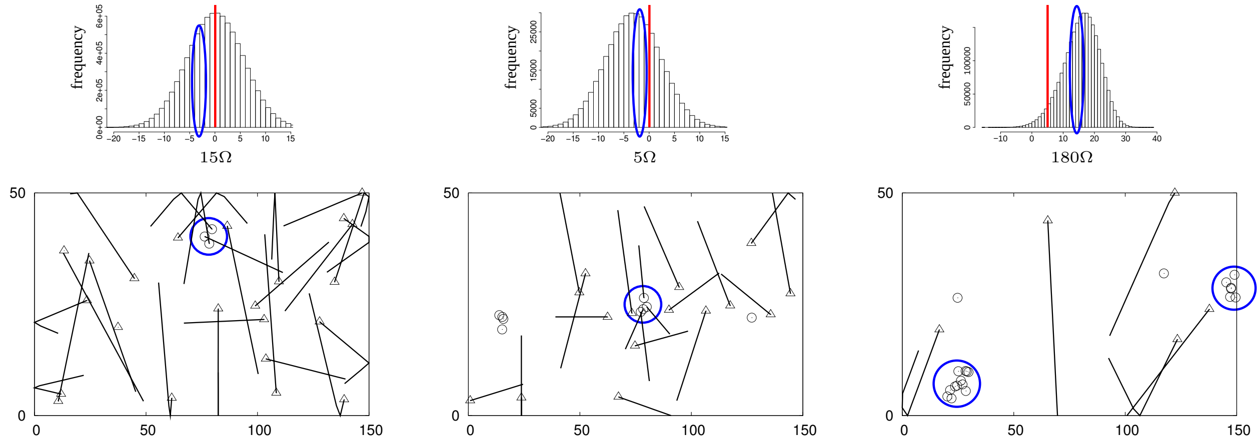

(a) Time until first occurrence of agent-to-agent collision and $\langle\Omega\rangle = 0$, $0 < t < 15$, $\rho(k, 0) = 2.1$, $\rho(k, 15) = -0.9$, $15\Omega = -3.0$, blue marks indicate Ω-value in histogram and just stopped agents.

(b) Early transient with $\langle\Omega\rangle < 0$ (cf. Fig. 6(b)), $15 < t < 20$, $\rho(k, 15) = -1.7$, $\rho(k, 20) = -3.6$, $5\Omega = -1.9$, blue marks indicate Ω-value in histogram and just stopped agents.

(c) Approach of self-organizing equilibrium (cf. Fig. 4), $20 < t < 200$, $\rho(k, 20) = 0.1$, $\rho(k, 200) = 10.1$, $200\Omega = 10.0$, blue marks indicate Ω-value in histogram and main clusters.

Figure 7: The three phases observed in the investigated 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).

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. (Prigogine, 1997, p. 62).

Conclusion

In this paper, we have analyzed an emergent self-organizing multi-agent (or swarm) system controlled by the BEECLUST algorithm 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. In addition, this work suggests the rich and thought-provoking metaphor of considering emergent swarm systems as implementations of a 'distributed Maxwell's demon' because random events are leveraged by autonomous decisions of embodied agents based on locally measured samples of a global entropy change. 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. For example, in flocking dissipation occurs due to rotational accelerations and averaging of directions (loss of information). Potential generalization is also indicated by preliminary results in other scenarios which will be reported in future work.

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 production can be observed. 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 then the rate at which events of negative entropy production occur. Thus, conditions 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 (Szopek et al., 2008) 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 offers many fully developed methods which can often be applied even unmodified to problems of emergent behavior (cf. Hamann et al. (2011a)). 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

Supported by: EU-IST-FET project 'SYMBRION', no. 216342; EU-ICT project 'REPLICATOR', no. 216240.

References

Adami, C. (1998). *Introduction to artificial life*. Springer-Verlag.

Bannerman, S. T., Price, G. N., Viering, K., and Raizen, M. G. (2009). Single-photon cooling at the limit of trap dynamics: Maxwell's demon near maximum efficiency. *New Journal of Physics*, 11(6):063044.

Bodenschatz, E., Pesch, W., and Ahlers, G. (2000). Recent developments in Rayleigh-Bénard convection. *Annual review of fluid mechanics*, 32(1):709–778.

Crick, F. H. (1968). The origin of the genetic code. *Journal of Molecular Biology*, 38(3):367–379.

De Wolf, T. 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 Engineerings Self Organising Applications*, volume 3464 of *Lecture Notes in Computer Science*, pages 1–15. Springer-Verlag.

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.

Evans, D. J. and Searles, D. J. (2002). The fluctuation theorem. *Advances in Physics*, 51(7):1529–1585.

Gell-Mann, M. (1995). Plectics. In Brockman, J., editor, *The Third Culture: Beyond the Scientific Revolution*, pages 316–332. Touchstone Press, New York.

Gerstner, E. (2002). Second law broken: Small-scale energy fluctuations could limit minaturization. Nature online news.

Haken, H. (1977). *Synergetics - an introduction*. Springer-Verlag, Berlin, Germany.

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., and Mouret, J., editors, *From Animals to Animats 11*, volume 6226 of *LNAI*, pages 639–648. Springer-Verlag.

Hamann, H., Schmickl, T., and Crailsheim, K. (2011a). Thermodynamics of emergence: Langton's ant meets Boltzmann. In *IEEE Symposium on Artificial Life (IEEE ALIFE 2011)*, pages 62–69. IEEE.

Hamann, H., Schmickl, T., Wörn, H., and Crailsheim, K. (2011b). Analysis of emergent symmetry breaking in collective decision making. *Neural Computing & Applications*. in press.

Hereford, J. M. (2011). Analysis of BEECLUST swarm algorithm. In *Proc. of the IEEE Symposium on Swarm Intelligence (SIS 2011)*, pages 192–198. IEEE.

Maxwell, J. C. (1871). *Theory of Heat*. Dover Publications.

Maxwell, J. C. (1878). Tait's 'Thermodynamics' (I). *Nature*, 17:257–259.

Nicolis, G. and Prigogine, I. (1977). *Self-organization in nonequilibrium systems*. Wiley New York.

Parunak, H. V. D. and Brueckner, S. (2001). Entropy and self-organization in multi-agent systems. In *AGENTS'01: Proceedings of the fifth international conference on Autonomous agents*, pages 124–130, New York, NY, USA. ACM Press.

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.

Prigogine, I. (1997). *The End of Certainty: Time, Chaos, and The New Laws of Nature*. Free Press.

Schmickl, T. and Hamann, H. (2011). BEECLUST: A swarm algorithm derived from honeybees. In Xiao, Y. and Hu, F., editors, *Bio-inspired Computing and Communication Networks*. Routledge.

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.

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.

Schweitzer, F. (2003). *Brownian Agents and Active Particles. On the Emergence of Complex Behavior in the Natural and Social Sciences*. Springer-Verlag, Berlin, Germany.

Smith, E. (2008). Thermodynamics of natural selection I: Energy flow and the limits on organization. *Journal of Theoretical Biology*, 252(2):185197.

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., Ballintijn, M., Bogers, N., Grieco, F., Loijens, L., Noldus, L., Smit, G., and Zimmerman, P., editors, *Proceedings of Measuring Behavior 2008, Maastricht, August 26-29*, volume 6, page 327.

Wang, G. M., Sevick, 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. *Phys. Rev. Lett.*, 89(5):050601.

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.

Wu, M., Ahlers, G., and Cannell, D. (1995). Thermally induced fluctuations below the onset of Rayleigh-Bénard convection. *Physical review letters*, 75(9):1743–1746.