Homeotaxis: Coordination with Persistent Time-Loops

Mikhail Prokopenko, Astrid Zeman, and Rongxin Li

CSIRO Information and Communication Technology (ICT) Centre
Locked bag 17, North Ryde, NSW 1670, Australia
mikhail.prokopenko@csiro.au

Abstract. We present a novel approach to self-organisation of coordinated behaviour among multiple resource-sharing agents. We consider a hierarchical multi-agent system comprising multiple energy-dependent agents split into local neighbourhoods, each with a dedicated controller, and a centralised coordinator dealing only with the controllers. The coordinated behaviour is required in order to achieve a balance between the overall resource consumption by the multi-agent collective and the stress on the community. Minimising the resource consumption increases the stress, while reducing the stress may lead to unrestricted and highly unpredictable demand, harming the individual agents in the long-run. We identify underlying forces in the system's dynamics, suggest a number of quantitative measures used to contrast different strategies, and introduce a novel strategy based on persistent sensorimotor time-loops: *homeotaxis*. Homeotaxis subsumes the homeokinetic principle, extending it both in terms of scope (multi-agent self-organisation) and the state-space, and allows to select the best adaptive strategy for the considered system.

1 Introduction

In general, the ability to coordinate (e.g., synchronise) multiple individual actions within large multi-agent groups is an adaptive response observed in many biological systems. As noted by Trianni and Nolfi [17], "synchrony can increase the efficiency of a group by maximising the global outcome or by minimising the interference among individuals". An investigation of Baldassarre *et al.* [3] characterised coordinated motion in a swarm collective as a self-organised activity of the constituent independently-controlled modules, and measured the increasing organisation of the group on the basis of Boltzmann entropy. The emergent *common direction* of motion, with the chassis orientations of the robots spatially aligned, was related to high synchrony and coordination within the group.

While synchronisation has been extensively studied in a variety of applications, ranging from swarm robotics [3,17], to coordinating sensors in wireless networks [18], to models of fireflies flashing in unison in biology [11], the inverse problem of desynchronisation has received less attention, as noted by Patel *et al.* [13]. Patel *et al.* consider desynchronisation as the task of spreading a given set of identical oscillators throughout a time period, resulting in a round-robin schedule, and argue that this can be useful in several applications. For example, in wireless sensor networks sensor nodes can (i) desynchronise their sampling times to distribute the energy cost, while still providing efficient coverage, and (ii) desynchronise their transmission times to avoid collisions

M. Asada et al. (Eds.): SAB 2008, LNAI 5040, pp. 403–414, 2008.

and message loss. The study of Patel *et al.* is motivated by biology: "cells, acting as oscillators, control animal gaits and regulate heart valves through desynchronisation" [13], leading to biologically-inspired algorithms for achieving desynchronisation.

Thus, both synchronisation and desynchronisation may be desirable at different stages in a multi-agent system. We believe, synchronisation enables better exploitation, while desynchronisation allows the system to explore alternatives. In general, a system that is well-balanced in terms of synchronisation and desynchronisation (i.e., in terms of exploitation and exploration) achieves an adequate coordination of its components through space and time, and across multiple activities. In other words, synchronisation and desynchronisation are interleaved sub-tasks of a more generic task: multi-agent spatiotemporal coordination.

A coordinated system may involve a degree of centralisation. For example, a bio-inspired control architecture for artificial muscle materials is proposed by Odhner *et al.* [12]. The active material is broken up into many small cells, coordinated to produce a combined force or displacement. A single central controller uses only one input and one output: it measures only the summed displacement of all of the cells, producing a feedback signal that is broadcast to the cells. Each cell controls its displacement with a stochastic automaton, that is "a small local control automaton containing a pseudo-random number generator, so that it contracts and relaxes stochastically, with a probability distribution dictated by the input from the central feedback controller" [12]. This closed-loop system results in a smooth and predictable motion (e.g., tracking a desired position), and is scalable to many cells.

This paper solves the task of coordination in a model setting (distributed power load management system) using a novel domain-invariant and bio-inspired principle: *homeotaxis*, that incorporates homeokinesis and self-organisation of perception-action loops. The coordinated behaviour in our model system is required in order to balance the overall energy consumption by a multi-agent collective and the stress on the community. Minimising the energy consumption strains the system, while removing the stress typically leads to unrestricted and highly unpredictable demand, harming the individual agents in the long-run.

Section 2 introduces our model problem: resource-sharing agents connected to a power load management system, and presents the motivation for the study. Domain-invariant approaches to multi-agent coordination are reviewed in Section 3. The following sections describe the proposed approach, and present simulation results, followed by conclusions.

2 A Distributed Power Load Management System

Electricity distribution is a complex system, consisting of loads (appliances), generators, transformer stations and distribution networks, and influenced by a market. During most times of the year, the market price is low (e.g., less than AUD$20 per megawatt hour), but it can also be very high (near or at the maximum price, e.g., AUD$10,000 per megawatt hour) during a peak demand, when energy supply is under extreme pressure.

Energy demand management is a new technology proposed to cope with the unpredictability of the energy market and provide a rapid response when supply is strained by

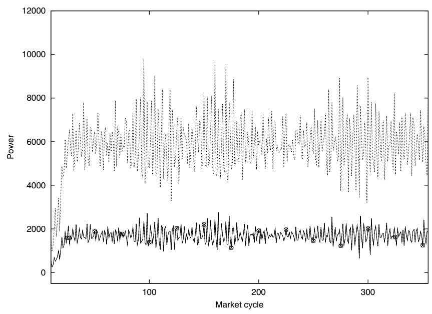

Fig. 1. Strategy $q_t = l_t$. Circles on solid line: x_t; crosses on dashed line: l_t; dotted line: u_t.

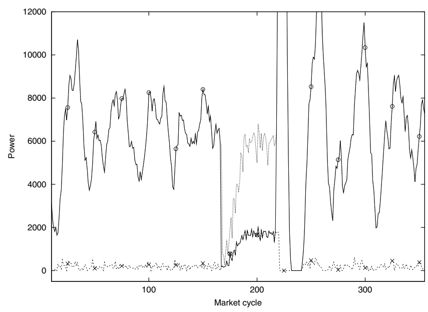

Fig. 2. Strategy $q_t = u_t$. Circles on solid line: x_t; crosses on dashed line: l_t; dotted line: u_t.

demand. Energy demand management essentially equips appliances with simple agents, and enables these agents to defer their electricity consumption when price suddenly skyrockets. We consider a hierarchical multi-agent system comprising multiple appliance-level agents split into local neighbourhoods, each with a dedicated neighbourhood's controller (middle-level), and a high-level centralised coordinator dealing only with the middle-level controllers. Appliances within a neighbourhood can be switched on or off by their middle-level controller, within the constraints defined by customer preferences. The neighbourhood's controller agent in turn receives an energy quota from the high-level coordinator — this energy quota is the limit on the total energy consumption for the neighbourhood.

In this paper, we focus on the behaviour of the high-level coordinator, i.e., on how to choose the optimal strategy setting the real-time quota to a middle-level controller. The high-level coordinator needs to incorporate information from the market (e.g., the current local price of energy p_t), as well as inputs from middle-level controller. The coordinator is required to balance the risk of exposure to volatile prices, and reduce strain in the load network.

The local energy demand limits for a neighbourhood are obtained by each middle-level controller as two values per interval: 1) the minimal consumption l_t of the local neighbourhood, and 2) the unconstrained (default) consumption u_t of the neighbourhood. The first value is defined as the total minimal consumption over an entire future market cycle (e.g., five minute interval). The minimal load requirements of appliances depend on the hardware and operating limits of individual appliances. The unconstrained consumption of the neighbourhood refers to its total consumption when the individual agents operate under normal conditions without any external influence. The limits l_t and u_t are computed by each middle-level controller in the beginning of a market cycle, using an optimisation procedure described elsewhere [9].

The coordinator sets the quota q_t for the next market cycle, as a value inclusively between the values of the minimal and unconstrained demand: $l_t \leq q_t \leq u_t$. In response to this quota the middle-level controller selects the best control plan for each load in its neighbourhood [9], resulting in actual total consumption x_t in the neighbourhood. These interactions essentially create a control loop where multiple individual agents plan their consumption for the next interval t, producing the aggregate limits l_t and u_t, while the coordinator reacts to these inputs (as well as the dispatch price p_t) by setting the quota q_t. This results in the actual consumption x_t, affecting the agents' plans for the next interval $(t+1)$ and the future limits l_{t+1} and u_{t+1}, and so on. Thus, the limits depend on the quotas set in the past.

One may assume that an optimal quota-setting strategy would be the minimal achievable consumption for a neighbourhood: $q_t = l_t$, that is, restricting the agents' energy consumption levels to the lowest possible value that they can cope with. This does reduce the total cost $S(T) = \sum_t^T p_t x_t$. However, when the minimal quota is imposed, some of the energy which is not consumed in one period will need to be consumed at a later time. The longer a minimal quota is imposed, the more loads are pushed to a stressful limit. The results of such undesirable synchronisation are shown in Fig. 1: the minimal quota increases the minimal achievable consumption l_t after some initial period during which the agents can still defer their demand. The agents become eventually synchronised when none of them can defer the demand any longer, stressing the neighbourhood. Such stress would result in a peak demand if a quota is released, and can cause negative effects such as shortage of energy or a peak price, threatening system stability. In short, when $q_t = l_t$, the minimal consumption l_t itself grows over time, and while the cost S is low, the stress (the total uncompromisable "incompressible" demand), defined as $R(T) = \sum_t^T l_t$, is very high.

Similarly, setting the quota to another possible extreme $q_t = u_t$ (except during the period around cycle 200 when an extremely high price p_t forces the coordinator to use the default $q_t = l_t$) results in an unbalanced outcome (Fig. 2). In this case, the stress $R(T)$ is minimal, but the cost $S(T)$ is high. The stability of the actual consumption x_t, measured as the variance $\sigma_x^2(T)$, is also quite significant.

Hence the coordinator needs to define an optimal quota which not only reduces energy costs, but also avoids instability and stress within the system. More precisely, optimality of the quota over a time period T depends on (i) the total cost $S(T)$; (ii) stability $\sigma_x^2(T)$ of the actual consumption x_t; and (iii) the stress $R(T)$. In order to derive an optimal quota-setting strategy we turn to domain-invariant principles of multi-agent coordination.

3 Domain-Invariant Principles of Adaptive Behaviour

A few approaches were recently proposed in order to characterise and achieve spatiotemporal coordination in a general way. For example, a modular robotic system modelling a multi-segment snake-like organism, with actuators ("muscles") attached to individual segments ("vertebrae") was evolved according to generic information-theoretic measure (excess entropy or predictive information, defined in Shannon sense) [16]. In general, one may argue that *information-driven self-organisation* is one of the main evolutionary forces that may be used in both design (e.g., information-driven evolutionary design [15]) and biological evolution [14]. An example of such information-driven dynamics is the acquisition of information from the environment: there is evidence that pushing the information flow to the information-theoretic limit (i.e., maximisation of information transfer through the system's perception-action loop) can give rise to intricate behaviour, induce a necessary structure in the system, and ultimately adaptively reshape the system [8].

Information-driven self-organisation that relies on dynamics of predictive information [15,10] is related to the search of domain-invariant principles carried out by Der *et al.* [4,7,6,5]. While traditionally an objective function, measuring the distance between the current and a desired behaviour, is provided explicitly by the designers, Der *et al.* consider self-referential adaptive systems: the systems for which the objective function is derived from the dynamics of the system itself, i.e. "adaptive, embodied systems where the objective of adaptation is a function of the robots sensor values alone" [5]. In doing so they follow the principle of homeokinesis — the dynamical analogue of homeostasis [1]. The principle of homeokinesis was developed by Der *et al.* [4,7] as a general domain-invariant principle for self-organisation in robot behaviour. The principle is based on the assumptions that the robot is able to (i) learn an internal representation (self-model) of its current behaviour, and (ii) adapt its behaviour by minimising the difference between the self-model and true behaviour in the real world. Importantly, the principle does not lead to stabilisation of stationary states (that would result in a "do-nothing" behaviour), but rather suggests a smooth and predictable kinetic regime.

In developing the principle of homeokinesis, Der and Liebscher [7] have introduced a reversal of time in the modelling process, capturing a time-loop error. The time-loop error is formalised in Section 4 — at this stage we note, following [7], that it is driven by two opposite forces. On the one hand, the time-loop error is small if the current behaviour is well represented by the internal model, producing behaviours that correspond to smooth and predictable sensor values. The second tendency is that the time reversal in the modelling process inverts a stable behaviour into unstable, and vice versa. Specifically, if the behaviour is stable in the forward time direction, it is unstable if the time is reversed, and therefore, the time-loop error is minimised if the kinetic behaviour of the robot is unstable in the forward time direction. This feature eliminates trivial behaviours (e.g., a "do-nothing" behaviour). Nevertheless, the instability cannot continue unrestricted, as the first tendency (the faithful self-model) demands smooth trajectories in the sensor space. This closed-loop interplay between a smooth and predictable sensor space exploration and an unstable kinetic behaviour balances exploration and exploitation aspects of the behaviour.

Importantly, both the information-driven self-organisation and the principle of homeokinesis emphasise the role of behaviour's predictability as well as non-stationary sensorimotor dynamics in achieving the desired balance between exploitation and exploration. A possible unification is described by Ay *et al.* [2] and Der *et al.* [5].

4 Homeotaxis as Coordination with Persistent Time-Loops

In deriving the optimal quota-setting strategy we extend the homeokinetic principle to the task of multi-agent coordination, by considering the consumption x_t within a multi-agent system to be among sensory inputs of the high-level coordinator, and the coordinator's output q_t to be its actuation. This extends the principle to *homeokinetic coordination*. Homeotaxis is achieved by enhancing the time-loop error with persistence error, and is motivated by *taxis*: "an innate behavioural response by an organism to a directional stimulus" [19][1].

4.1 Example

Let us begin by exemplifying the original use of the time-loop error, following a simple homeokinesis model of Der and Liebscher [7]. Consider the linear system with dynamics

$$x_{t+1} = c_t x_t + \xi_t \tag{1}$$

where c_t is the controlling variable, and ξ_t is the part not handled by the model. The system can be modelled forward in time:

$$\hat{x}_{t+1} = c_t x_t \ , \tag{2}$$

as well as backward in time:

$$\breve{x}_t = x_{t+1}/c_t = (c_t x_t + \xi_t)/c_t = x_t + \xi_t/c_t \tag{3}$$

Here $\hat{x}_{t+1}$ is the predicted value given the current observation x_t, while $\breve{x}_t$ is the reconstructed value, given the latest observation x_{t+1}. Traditionally, having observed x_{t+1}, one uses $E = (x_{t+1} - \hat{x}_{t+1})^2$ as the error to minimise, i.e. as a feedback signal. The time-loop error is defined as

$$W = (\breve{x}_t - x_t)^2 \ , \tag{4}$$

that is, it is obtained by going forward in time from x_t to x_{t+1}, followed by a step (3) from x_{t+1} to the reconstructed state $\breve{x}_t$ backward in time. The full sequence $x_t \rightarrow x_{t+1} \rightarrow \breve{x}_t$ is called the time loop. Using gradient descent to minimise the time-loop error (4) $W = \xi_t^2/c_t^2$, with respect to the controlling variable c_t, yields the update rule (for a small number $\epsilon > 0$ determining the rate of descent)

$$c_{t+1} = c_t + \epsilon \, \xi_t^2/c_t^3 \tag{5}$$

As mentioned in the previous section, the time-loop error W is small if the current behaviour is well represented by the internal model (2), producing behaviours that obtain smooth and predictable ("exploitable") sensor values. On the other hand, the time-loop error is minimised if the kinetic behaviour is unstable in the forward time direction, eliminating trivial stationary behaviours, and encouraging exploration.

4.2 General Case

In general, one considers a system with an adaptive controller, defined by the controller's parameter vector $\mathbf{c}_t$, and output that depends on the sensor values $\mathbf{x}_t$ observed at time t. The adaptive model M aims to predict the true sensor values $\mathbf{x}_{t+1}$ at $(t+1)$:

$$\hat{\mathbf{x}}_{t+1} = \mathbf{x}_t + M(\mathbf{x}_t, y_t; m) \ , \tag{6}$$

[1] Taxis differs from *kinesis*, "a non-directional change in activity in response to a stimulus that results in the illusion of directed motion due to different rates of activity depending on stimulus intensity" [19].

where m is the predictor's parameter. Abbreviating $M(\mathbf{x}_t, y_t; m)$ as $M(\mathbf{x}_t)$, and minimising the prediction error

$$E = (\mathbf{x}_{t+1} - \hat{\mathbf{x}}_{t+1})^2 = (\mathbf{x}_{t+1} - \mathbf{x}_t - M(\mathbf{x}_t))^2 = (\triangle\mathbf{x}_t - M(\mathbf{x}_t))^2 \qquad (7)$$

by gradient descent yields an update rule for the controller and the predictor, respectively (the latter does not have to be updated in real-time), where $\epsilon > 0$, $\eta > 0$:

$$\mathbf{c}_{t+1} = \mathbf{c}_t - \epsilon \frac{\partial E}{\partial \mathbf{c}_t} \qquad m_{t+1} = m_t - \eta \frac{\partial E}{\partial m} \qquad (8)$$

Assuming the dynamics

$$\mathbf{x}_{t+1} = \mathbf{x}_t + M(\mathbf{x}_t) + \xi_t \ , \qquad (9)$$

where ξ_t is the vector of perturbations not covered by the model M, the model can be written backward in time [7] as

$$\check{\mathbf{x}}_t = \mathbf{x}_{t+1} + M^{(-)}(\mathbf{x}_{t+1}) \ , \qquad (10)$$

where we define the reverse model $M^{(-)}$ as follows:

$$\mathbf{x}_t = \hat{\mathbf{x}}_{t+1} + M^{(-)}(\hat{\mathbf{x}}_{t+1}) \qquad (11)$$

The definition (11) is symmetric to the expression (6). This definition corrects the one given in [7] that specifies $\mathbf{x}_t = M^{(-)}(\mathbf{x}_t + M(\mathbf{x}_t))$, which is equivalent to $\mathbf{x}_t = M^{(-)}(\hat{\mathbf{x}}_{t+1})$. The time-loop error then is given by

$$W = (\check{\mathbf{x}}_t - \mathbf{x}_t)^2 = (\xi_t + M^{(-)}(\mathbf{x}_{t+1}) - M^{(-)}(\hat{\mathbf{x}}_{t+1}))^2 \qquad (12)$$

yielding the update rules for the controller and the predictor, respectively:

$$\mathbf{c}_{t+1} = \mathbf{c}_t - \epsilon \frac{\partial W}{\partial \mathbf{c}_t} \qquad m_{t+1} = m_t - \eta \frac{\partial W}{\partial m} \qquad (13)$$

4.3 Extension

In developing the principle of homeotaxis and applying it to the distributed energy management, we consider a system with an adaptive controller and output

$$q_{t+1} = \begin{cases} l_{t+1} & \text{if } \ c_{t+1}x_t < l_{t+1} \\ u_{t+1} & \text{if } \ c_{t+1}x_t > u_{t+1} \\ c_{t+1}x_t & \text{otherwise} \end{cases} \qquad (14)$$

during a transition from t to $(t+1)$. Specifically, x_t is the last observed actual consumption (the sensor value) achieved in response to the quota q_t; the current limits l_{t+1} and u_{t+1} are known as well. The quota-controlling parameter c_{t+1} is to be updated before setting the quota (the output) q_{t+1}.

The adaptive model aims to predict the true sensor value

$$x_{t+1} = c_{t+1}x_t + \xi_t \qquad (15)$$

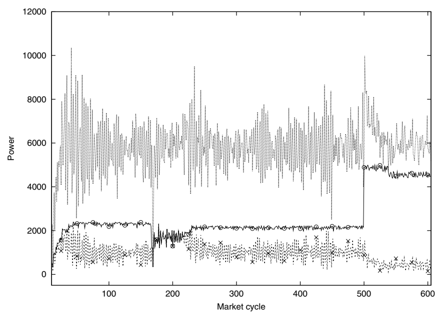

Fig. 3. Minimising error Q. Circles on solid line: x_t; crosses on dashed line: l_t; dotted line: u_t.

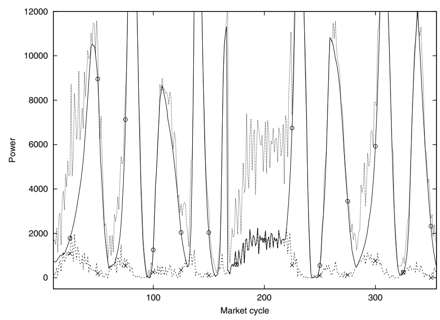

Fig. 4. Basic homeokinesis. Circles on solid line: x_t; crosses on dashed line: l_t; dotted line: u_t.

that will result when the quota is set:

$$\hat{x}_{t+1} = q_{t+1} \tag{16}$$

The quota dynamics may also be described recursively:

$$q_{t+1} = q_t + \zeta_t \tag{17}$$

for some perturbation ζ_t that needs to be minimised to maintain persistent quotas. The persistence error can be formulated, by applying relationships (17) and (14), as

$$Q = (q_t - q_{t+1})^2 \tag{18}$$

In our case, $Q = (q_t - c_{t+1}x_t)^2$ is minimised by

$$\hat{c}_{t+1} = q_t / x_t \tag{19}$$

The control strategy using this simplistic update rule balances cost, stability and stress reasonably well, but is not responsive to sudden changes in the underlying demand. Fig. 3 shows the dynamics for a scenario where there is a spike in demand at market cycle 500 — it is evident that the optimal balance is lost after the spike. However, it can be complemented by a signal produced by the time-loop error:

$$W_Q = (\check{x}_{t;Q} - x_t)^2 = (x_{t+1}/\hat{c}_{t+1} - x_t)^2 = ((\hat{c}_{t+1}x_t + \xi_t)/\hat{c}_{t+1} - x_t)^2 = \xi_t^2/\hat{c}_{t+1}^2 \tag{20}$$

The error W_Q exemplifies *persistent time-loop error*:

$$W_Q = (\check{x}_{t;Q} - x_t)^2 = (\xi_t + M^{(-)}(x_{t+1};\hat{c}_{t+1}) - M^{(-)}(\hat{x}_{t+1}))^2 \tag{21}$$

where $\hat{c}$ is determined by minimisation of the persistence error Q (18). The controller update rule that minimises both the persistence error Q and the usual time-loop error is given by

$$c_{t+1} = \hat{c}_{t+1} - \epsilon \frac{\partial W_Q}{\partial \hat{c}_{t+1}} = \hat{c}_{t+1} + \epsilon' \frac{\xi_t^2}{\hat{c}_{t+1}^3} \tag{22}$$

In summary, the *persistent time-loop error* W_Q is obtained by (i) minimising the persistence error (18), and (ii) using the corrected controlling variable in the time-loop reconstruction. Essentially, it explicitly demands smoothness and predictability of trajectories in both the sensor-space and action-space. It is interesting to note, at this stage, that a basic homeokinetic strategy where update rule is given by (5) does not work in this case: as shown in Fig. 4, the actual consumption is too unstable, while being reasonably smooth and predictable on average.

The next step is to estimate the perturbation ξ_t. We used the following approximations:

$$\xi_t^{(1)} = |\check{x}_t - x_t| \qquad \xi_t^{(2)} = |x_{t+1} - x_t| \qquad \xi_t^{(3)} = |\check{x}_t - q_t| \tag{23}$$

Perfect control is indicated by $c = 1$. According to (20), $\xi_t^{(1)}$ yields $\hat{c} = 1$, while $\xi_t^{(2)}$ yields $c = 1$ according to (15), and $\xi_t^{(3)}$ yields $\hat{c} = 1$ according to (20) and (19) considered together.

5 Results

In this section we present experimental results of three different homeotaxic strategies that minimise persistent time-loop error with approximations $\xi_t^{(1)}$ (Fig. 5), $\xi_t^{(2)}$ (Fig. 6), and $\xi_t^{(3)}$ (Fig. 7). Each experiment involved 100 individual agents in the neighbourhood coordinated over a 5000 cycle run, including a short period around cycle 200 during which the coordinator used the default $q_t = l_t$, and an instantaneous spike in demand to $min(5000; u_t)$ at cycle 500. The performance was measured in terms of the cost S, stability σ_x^2, stress R over $T = 5000$ cycles, and their linear combination

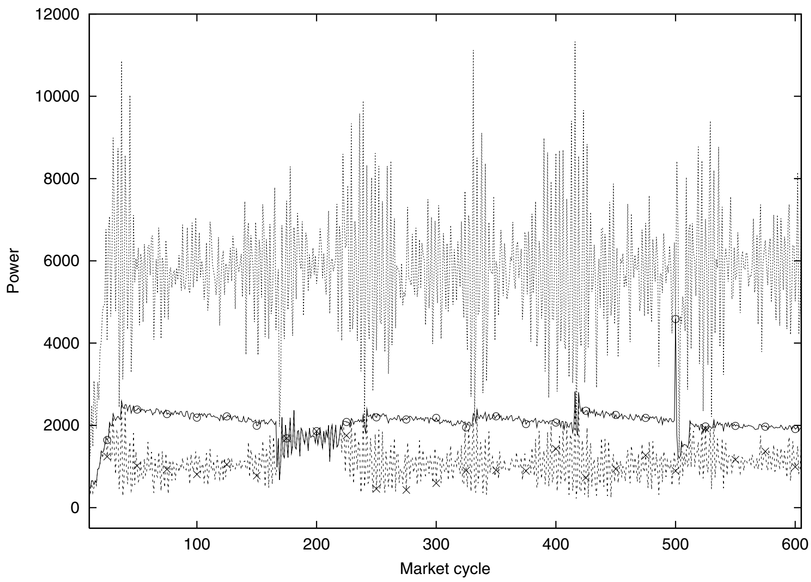

Fig. 5. Minimising $\xi_t^{(1)}$. Circles on solid line: x_t; crosses on dashed line: l_t; dotted line: u_t.

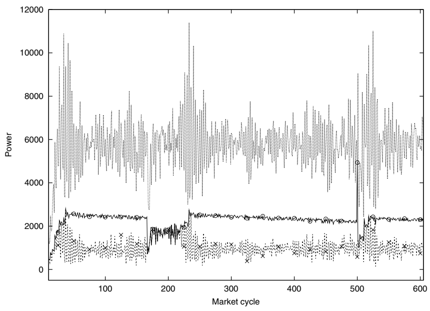

Fig. 6. Minimising $\xi_t^{(2)}$. Circles on solid line: x_t; crosses on dashed line: l_t; dotted line: u_t.

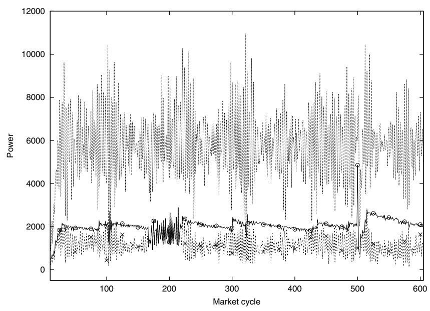

Fig. 7. Minimising $\xi_t^{(3)}$. Circles on solid line: x_t; crosses on dashed line: l_t; dotted line: u_t.

Table 1. Comparative results (bold: best in each category)

Strategy	$S(T)$	$\sigma_x^2(T)$	$R(T)$	Loss	RMSE$_c$
homeotaxis $\xi_t^{(1)}$	376,477	**38,806**	5,405,235	95,581	0.04371
homeotaxis $\xi_t^{(2)}$	385,528	41,967	5,101,962	**93,769**	0.04440
homeotaxis $\xi_t^{(3)}$	373,629	47,850	5,551,528	97,663	0.04388
persistent quota	547,045	394,167	2,941,011	123,531	**0.03922**
$q_t = l_t$	**334,945**	109,560	8,498,646	129,437	c_t not used
$q_t = u_t$	833,075	3,454,456	**973,996**	438,493	c_t not used
basic homeokinesis	723,124	20,900,555	1,751,512	2,179,883	0.52995

(the "loss"): $0.1(S + \sigma_x^2) + 0.01R$ (assuming, e.g., that cost and stability contribute ten times as much as stress). All homeotaxic strategies achieve good results (Table 1), easily outperforming the strategies considered above (minimisation of persistence error; $q_t = l_t$; $q_t = u_t$; and basic homeokinesis). We used the root mean square error (RMSE$_c$) in estimating how close the controller c_t is to the perfect case ($c_t = 1$), i.e. RMSE$_c = \sqrt{\sum_1^T (c_t - 1)^2/T}$. The success can be explained by the combination of homeokinetic exploration and exploitation of persistent actions. The optimal strategies based on $\xi_t^{(1)}$, $\xi_t^{(2)}$ and $\xi_t^{(3)}$ have also shown robustness to selfishness of individual agents (selfishness is defined here as a fraction of agents that always refuse to follow a control plan), and to additive noise in actual consumption, that can be due to a fraction of agents that intermittently decide against following a control plan (these experimental results are omitted due to a lack of space).

6 Conclusions

We argued that adaptive coordination of multiple resource-sharing agents requires both synchronisation and desynchronisation, balancing exploitation and exploration stages.

This balance can be derived (learned) from dynamics of the multi-agent system itself, if there is a feedback measuring the distance between current and a desired behaviour. A model hierarchical system comprised multiple energy-consuming agents split into local neighbourhoods, each with a dedicated controller, and a centralised coordinator dealing only with the controllers. The overall consumption within a neighbourhood was considered as a sensory input of the high-level coordinator, while the coordinator's output formed its actuation. This interpretation allowed us to develop homeotaxis: a generic domain-invariant approach, based on the homeokinetic principle, extended in terms of scope (multi-agent self-organisation) and the state-space (predictable perception and action spaces). A number of homeotaxic strategies, based on persistent sensorimotor time-loops, were introduced and experimentally verified, achieving a balance between resource consumption and stress within the multi-agent community.

Acknowledgments. The authors are grateful to Oliver Obst for suggesting the term *homeotaxis*, and to Peter Corke for pointing out the work of Odhner *et al.* [12].

References

1. Ashby, W.R.: Design for a Brain. Chapman and Hill, London (1954)
2. Ay, N., Bertschinger, N., Der, R., Güttler, F., Olbrich, E.: Predictive information and explorative behavior of autonomous robots. European Journal of Physics: Complex Systems (submitted, 2007)
3. Baldassarre, G.: Self-organization as phase transition in decentralized groups of robots: A study based on boltzmann entropy. In: Prokopenko, M. (ed.) Advances in Applied Self-organizing Systems, pp. 127–146. Springer, Heidelberg (2007)
4. Der, R.: Self-organized robot behavior from the principle of homeokinesis. In: Proceedings of the SOAVE Workshop, Ilmenau (2000)
5. Der, R., Güttler, F., Ay, N.: Predictive information and emergent cooperativity in a chain of mobile robots. Santa Fe Institute Working Paper 08-02-006 (2008)
6. Der, R., Hesse, F., Martius, G.: Rocking stamper and jumping snake from a dynamical system approach to artificial life. J. Adaptive Behavior 14, 105–116 (2005)
7. Der, R., Liebscher, R.: True autonomy from self-organized adaptivity. In: Proceedings of The Workshop on Biologically Inspired Robotics, The Legacy of W. Grey Walter, August 14-16, 2002, HP Bristol Labs (2002)
8. Klyubin, A.S., Polani, D., Nehaniv, C.L.: All else being equal be empowered. In: Capcarrère, M.S., Freitas, A.A., Bentley, P.J., Johnson, C.G., Timmis, J. (eds.) ECAL 2005. LNCS (LNAI), vol. 3630, pp. 744–753. Springer, Heidelberg (2005)
9. Li, R., Poulton, G.: Optimal management of multiple resource consuming devices during periods of supply restriction. Technical report, CSIRO ICT Centre, Macquarie University (May 2007)
10. Lizier, J.T., Prokopenko, M., Zomaya, A.Y.: Local information transfer as a spatiotemporal filter for complex systems. In: Phys. Rev.E (accepted, 2008)
11. Mirollo, R., Strogatz, S.: Synchronization of pulse-coupled biological oscillators. SIAM Journal of Applied Math 50, 1645–1662 (1990)
12. Odhner, L., Ueda, J., Asada, H.: Broadcast feedback of stochastic cellular actuators. In: Proc. of The 10th International Symposium on Experimental Robotics (ISER), Rio de Janeiro (2006)

13. Patel, A., Degesys, J., Nagpal, R.: Desynchronization: The theory of self-organizing algorithms for round-robin scheduling. In: Proceedings of The First IEEE Conference on Self-Adaptive and Self-Organizing Systems (SASO), pp. 87–96 (2007)
14. Piraveenan, M., Polani, D., Prokopenko, M.: Emergence of genetic coding: an information-theoretic model. In: eCosta, F.A., Rocha, L.M., Costa, E., Harvey, I., Coutinho, A. (eds.) ECAL 2007. LNCS (LNAI), vol. 4648, pp. 42–52. Springer, Heidelberg (2007)
15. Prokopenko, M.: Design vs self-organization. In: Prokopenko, M. (ed.) Advances in Applied Self-organizing Systems, pp. 3–17. Springer, Heidelberg (2007)
16. Prokopenko, M., Gerasimov, V., Tanev, I.: Evolving spatiotemporal coordination in a modular robotic system. In: Nolfi, S., Baldassarre, G., Calabretta, R., Hallam, J.C.T., Marocco, D., Meyer, J.-A., Miglino, O., Parisi, D. (eds.) SAB 2006. LNCS (LNAI), vol. 4095, pp. 558–569. Springer, Heidelberg (2006)
17. Trianni, V., Nolfi, S.: Minimal communication strategies for self-organising synchronisation behaviours. In: Proceedings of The First IEEE Symposium on Artificial Life, pp. 199–206 (2007)
18. Werner-Allen, G., Tewari, G., Patel, A., Nagpal, R., Welsh, M.: Firefly-inspired sensor network synchronicity with realistic radio effects. In: Proceedings of The Third ACM Conference on Embedded Networked Sensor Systems (SenSys), San Diego, pp. 142–153 (2005)
19. Wikipedia,: Taxis — wikipedia, the free encyclopedia (2007) [Online; accessed January 10, 2008]