Energy, Entropy and Work in Computational Ecosystems: A Thermodynamic Account

Mariusz Jacyno and Seth Bullock

School of Electronics and Computer Science
University of Southampton
Southampton, UK
mj04r@ecs.soton.ac.uk

Abstract

Recently, computer scientists have begun to build computational ecosystems in which multiple autonomous agents interact locally to achieve globally efficient organised behaviour. Here we present a thermodynamic interpretation of these systems. We highlight the difference between the regular use of terms such as energy and work, and their use within a thermodynamic framework. We explore the way in which this perspective might influence the design and management of such systems.

Introduction

Modern IT systems are increasingly complex, in some cases resembling computational ecologies or ecosystems comprising myriads of interacting elements, each with their own thread of control and autonomy (Huberman and Hogg, 1993; Bullock and Cliff, 2004). As the scale, dynamism and interconnectedness of such systems increases, effective control via some central executive becomes nontrivial and eventually infeasible. Consequently, there is increasing interest in drawing inspiration from the homeostatic properties of some analogous large-scale, adaptive, decentralised natural systems, and extracting design principles for artificial computational ecologies (Parunak and Brueckner, 2004; Zambonelli and Parunak, 2004). Underpinning this paradigm are concepts and theories of self-organisation that derive from the study of physical systems in far-from-equilibrium conditions (Nicolis and Prigogine, 1977; Kay, 1984).

However, the man-made nature of computational ecosystems gives them a particular teleological status that differs somewhat from the physical (and biological) systems to which thermodynamic accounts of self-organisation are typically applied. Moreover, computational systems are instantiated as physical systems, giving rise to a problem of identifying the level of description at which to apply the self-organisation/thermodynamic interpretation. Partly as a consequence, a consistent thermodynamic account of their behaviour is not straightforward. As engineers aiming to build artificial systems exhibiting adaptive and organised behaviour, we must be careful in applying the concepts and tenets of self-organisation. There may be some value in the informal use of technical language (e.g., attractor, basin of attraction) to re-describe problems, etc., but there are risks of confusion when technical terms (that may also have lay meanings) come to be used as mere façons de parler.

The purpose of this paper is thus to analyse self-organisation in natural systems as governed by the physical laws of thermodynamics and, based on this, to clairfy and make explicit an analogous interpretation of the functioning of artificial self-organising computational ecosystems.

Thermodynamics in natural systems

One powerful strength of a thermodynamic account of self-organisation is its potential to apply across physical, chemical, biological, social, and socio-technological domains. However, it is most clearly and straightforwardly articulated in the absence of the beliefs, desires, and functions that are proper parts of the 'higher' systems. Here we first present the framework in the context of physical and then biological systems before demonstrating its application in the context of a particular class of socio-technological system.

Thermodynamics of self-organisation

Studies investigating the thermodynamics of self-organisation in far-from-equilibrium systems can be found in (Nicolis and Prigogine, 1977; Swenson, 1997; Kauffman, 2000). Irrespective of whether the investigated system is described in terms of 'dissipative structures' (Nicolis and Prigogine, 1977), autonomous agents (Kauffman, 2000) or an autocatakinetic system (Swenson, 1997), self-organisation is interpreted as a process of organised energy flow from which work can be extracted and employed by the system for its structure maintenance (Kay, 1984; Wicken, 1989; Swenson and Turvey, 1991). Central to understanding this process are the following concepts derived from thermodynamics: displacement from equilibrium, energy transfer, gradient dissipation, constraint formation and work.

Displacement from equilibrium

According to classical thermodynamics, the behaviour of physical systems can be explained as transformations of energy between the system and its surroundings. Hence, when both are allowed to interact, what is exchanged between them is energy (Kay, 1984). Energy, here, has a general meaning, defining the capacity of the system to perform work, and may be added to the system by increasing its temperature, pressure or a chemical potential.

Considering the energy of the system and its environment, we can measure the relative difference between both, often defined as a potential or gradient. If the gradient is equal to zero, meaning that both the system and its environment have the same energy (e.g., temperature, or pressure) we consider them to be at equilibrium. In this state, the system is indistinguishable from its environment and has no capacity to perform work. Any deviation from equilibrium implies that free energy is stored, and that there may be the potential to release this energy through useful work. The extent to which a system is displaced from equilibrium is reflected in the gradient (difference) between the state variables defining its energy state (e.g., temperature) and that of its environment.

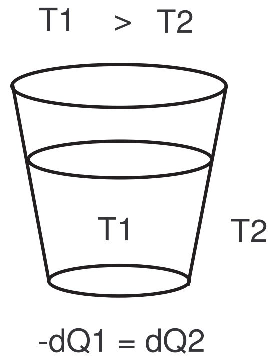

Figure 1: A glass of liquid at temperature $T1$ is placed in a room at temperature $T2$, where $T1 > T2$. The disequilibrium produces a field potential that spontaneously drives a flow of energy in the form of heat, $-dQ1$, from the glass to the room so as to drain the potential until it is minimized (the entropy is maximized). At this point thermodynamic equilibrium is reached and all flows stop. The expression $-dQ1 = dQ2$ refers to conservation of energy in that the flow of heat from the glass equals the flow of heat into the room.

Energy transfer

To displace a system from equilibrium requires that it be supplied with energy (be it thermal, mechanical or chemical), distinguishing it from its surroundings. According to the first law of thermodynamics, energy transfer can proceed in two different ways: through heat (Q) and work (W). This is captured in the formula summarising the first law:

$$dU = dQ + dW,$$

where dU is the infinitesimal increase in internal energy of the system, dQ is the infinitesimal amount of heat added to the system and dW is the infinitesimal amount of work done on the system. Although heating up a system and performing work on it will each increase its energy, each differs in the manner in which energy is being distributed in the system and thus whether the system moves away from equilibrium.

This difference is reflected through entropy (S) which can be interpreted as a measure of the uncertainty about how energy is distributed in the system (Jaynes, 1965, 1979). Adding heat (Q) to the system increases our overall uncertainty about the energy content of the system and causes proportional increase in entropy. This is manifested through the following relation:

$$dS = dQ/T,$$

where S is the entropy, dQ is the infinitesimal amount of heat added to the system and T is the absolute temperature of the system. For this reason, it represents the amount of energy that we lose information about when it is transferred and that we are thus unable to extract. When, on the other hand, work is done on the system (W) our knowledge about the energy content of the system increases, thus we are better able to distinguish between the system and its environment. In this case, work done on the system does not affect internal system entropy and thus represents the only way to move a system further from equilibrium (Kay, 1984).

Gradient dissipation

The second law of thermodynamics states that if two systems are allowed to interact and exchange energy, that is if the constraints imposed between them are removed, then the systems will evolve to equilibrium, a new state in which we cannot differentiate between the systems. A statistical consequence of this physical law is that entropy will increase.

The active nature of the second law is intuitively easy to grasp and empirically easy to demonstrate. Figure 1 shows a glass of hot liquid placed in a room at a cooler temperature. The difference in temperatures in the glass-room system constitutes a potential and induces a flow of energy in the form of heat. This ‘drain’ on the potential flows from the glass (source) to the room (sink) until the potential is minimized (the entropy is maximized) and the liquid and the room are at the same temperature. At this point, all flows and thus all entropy production stops and the system is at thermodynamic equilibrium. The same principle applies to any system where any form of energy is out of equilibrium with its surroundings (e.g., whether mechanical, chemical, electrical or energy in the form of heat).

The second law alone does not tell which of the available energy transfer paths the system will select in order to move back to equilibrium. The explanation to this can be provided in a classic experiment on self-organisation first devised by Henri Bénard in 1900 (Swenson and Turvey, 1991). A viscous fluid is held between a uniform heat source below and the cooler temperature of the air above. That is, there is a potential difference between fluid and air with a field force of a magnitude, F, determined by the difference between the two temperatures. When F is below a critical threshold heat flows from the source (fluid) to the sink (air) in the form of disordered collisions between the constituent molecules, and entropy is produced. If F exceeds the critical threshold Bénard 'cells' emerge spontaneously, each cell consisting of hundreds of millions of molecules moving collectively together in the form of rotating vertical convection columns. In this organised mode, the transfer of energy through the system and its dissipation to its surroundings is much more efficient than through unorganised collisions (Schneider and Kay, 1995). Such behaviour does not violate the second law. As long as a self-organising system produces entropy (minimises potentials) at a rate that is sufficient to compensate for its own ordering (persistence away from equilibrium) then the balance demanded by the equation of the second law is not violated (Kay, 1984; Swenson and Turvey, 1991).

Work

So far we have discussed displacement from equilibrium, constraint on energy transfer and gradient dissipation as distinct concepts describing the active nature of physical laws. But how can they be employed to control energy movement within systems, such that useful work could be extracted from their functioning (Jaynes, 1988)? Consider a system consisting of two connected tanks of equal volume but with different numbers of gas molecules. This difference defines a gradient between both tanks. As soon as a conduit between them is opened, gas whooshes through it, equalising the number of molecules in the tanks and erasing the gradient between them. Gas can rush through even if it has to turn a turbine along the way, thereby doing mechanical work. The energy to do that work came from the thermal energy of the environment, but the conversion from thermal to mechanical energy was paid for by the increase of disorder as the system equilibrated. Now, if we repeat the first process again by first closing the conduit and transferring energy from one tank to the other, we can repeat the same process of work extraction and gradient dissipation. Although simplified, this principle of work extraction constitutes a thermodynamic work cycle, which underpins the supply of most of the world's electric power and almost all motor vehicles.

Information

Within statistical mechanics, the entropy of a system at equilibrium can be recast in terms of the variety of microscopic states available to the system:

$$S \equiv k \ln \Omega,$$

where Ω is the number of states in which the system can be found when at equilibrium, and k is the Boltzmann constant, $1.38x10^{-16} J/K$. Consequently, entropy has been interpreted as a measure of macro-level disorder, formalised as Shannon entropy (Shannon, 1948) defined as:

$$S = -\sum p_i \log p_i,$$

where i ranges over the possible states of the system and p_i is the probability of finding the system in state i.

As such, it is possible to reinterpret the thermodynamic work cycle in information theoretic terms (Jaynes, 1988; Nelson, 2004). We have seen that the difference between doing work on a system and merely heating it up is the difference between how informed we are about the organisation of the system's energy. The potential gradient that must be established within a system before useful work can be extracted from it is thus also an informational property. Given that we are interested in computational systems that consume electricity and also process information, there is scope for the equivalences between information, energy and entropy to be useful, but also confusing.

Thermodynamics beyond physics

The application of thermodynamics is not limited only to physical systems (Jaynes, 1988). Ever since Alfred Lotka (1922) began writing about energy flows as the basis for natural selection, there has been a thermodynamic paradigm in evolutionary theory. Lotka observed that selection will favor those organisms that, in pulling resources into their own service, also increase the energy throughputs of their ecosystems (Wicken, 1989). What all organisms have in common is that they operate and evolve at some remove from thermodynamic equilibrium. By doing so they maintain the integrity of their organisational structures by irreversibly degrading free energy through informed kinetic pathways acquired through evolution. From this perspective, succession can be considered as the process by which an ecosystem moves away from thermodynamic equilibrium with its environment (Kay, 1984). By developing this account, the principles of variation and natural selection can be given a sound thermodynamic basis. The principle of variation derives from two sources: the entropic drive to generate configurational randomness and the quantum indeterminacy about where that randomness will occur. Natural selection follows from competition among alternative patterns of energy utilisation (Wicken, 1988).

One consequence of this perspective is an increasing appreciation that organisms can be viewed as more sophisticated 'engines' than the physical systems described so far

(Swenson, 1997). According to Kauffman (2000), for instance, life or its physical manifestation can be described in terms of an autonomous agent. This agent is a collectively autocatalytic system performing one or more thermodynamic work cycles that: (1) measures useful displacements from equilibrium from which work can be extracted; (2) discovers devices that couple to those energy sources such that work can be extracted; and (3) applies work to develop and maintain the constraints that enable the further extraction of work.

What are computational ecologies?

Figure 2 illustrates the architecture of a modern IT system. The infrastructure is an open system of interacting elements whose organisation is free to change and grow organically through the removal and introduction of components. Depending on the level of interpretation, these elements may be thought of either as physical servers or the software components hosted by them. However, these two levels of description encourage different thermodynamic accounts, and care must be taken in translating between them. While we may ultimately be interested in the physical time and energy required in order to achieve computational tasks, it has often been convenient to recast this problem in terms of efficient information processing with no explicit mention of energy, heat, etc.[1] However, while the 'motive force' driving the physical system stems from physical energy, the equivalent potential or gradient at the software level must be understood in terms of informational differences. Understanding how computation in such systems can be managed through self-organising mechanisms requires us to disentangle the physical system and software system levels.

For our purposes, the physical level of description can be stated rather straightforwardly. Autonomic computing systems are made up of a large-scale network of interconnected clusters of machines, each offering computational or storage resources. These functions are dependent on the constant supply of energy that is being fed to these machines in order to maintain their on-line functioning. The outcome of this energy consumption in the physical world is heat, generated in proportion to the intensity of computation. In an efficient system, this will in turn be related to system throughput, defined in terms of the number of computational tasks achieved per unit time.

By contrast, the software level of description, which will be our primary concern in the paper, is a little more complicated. We assume that the computational power offered by a system's physical servers constitutes a limited capacity raw resource, and that the efficient distribution of this resource to meet the needs of a set of users makes it natural to decentralise the control over their management to software elements (Kephart and Chess, 2003). This may be realised by applying a decentralised multi-agent architecture (Sycara, 1998) comprising a population of agents possessing their own thread of control and autonomy but perhaps lacking access to some central repository of system information. Figure 3 depicts the roles of software agents in managing system resources for such a system. The process is initiated by allocation requests from infrastructure users, U, illustrated by dashed arrows crossing the system boundary. Requests arrive in parallel and are intercepted by "user" or "consumer" software agents responsible for resource allocation (depicted as open circles). This incoming information 'agitates' the allocation process in resource consumers, inducing them to discover and select amongst available resource providers (represented as solid circles) until one agrees to execute the requested job (solid connecting line). In addition to executing the tasks for which they are currently configured, resource providers may also adapt to locally perceived demand by reconfiguring to offer the most demanded kinds of service.

[1]However, heat management in both large-scale and micro-scale systems is a growing concern (Skadron et al., 2003; Sharma et al., 2005).

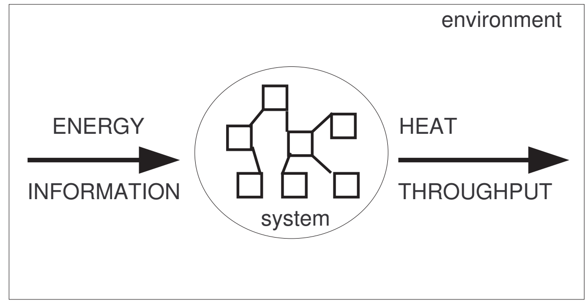

Figure 2: General architecture of the modern IT system.

It is important to note that the physics of the system imposes constraints on the software level. Allocation and reconfiguration decisions are only necessary as a consequence of the assumption that each resource provider is physically constrained such that it may only serve a limited number of consumers at the same time, and, furthermore, may only offer a limited set of services at any one time. Allocation and reconfiguration decisions must be efficient only as a consequence of the assumption that each interaction between agents, and each reconfiguration event incur associated physical costs in terms of time or power consumption.

The co-adaptation of resource providers and resource consumers takes place under conditions in which the demand for particular resource types may vary unpredictably. This requires providers to reconfigure their provision and consumers to track these reconfigurations. Consequently, the stability of the whole 'ecosystem' is dependent on the establishment of the information flow pathways that enable localised system elements to efficiently adapt and adjust their behaviours to the current system state. As the processing and

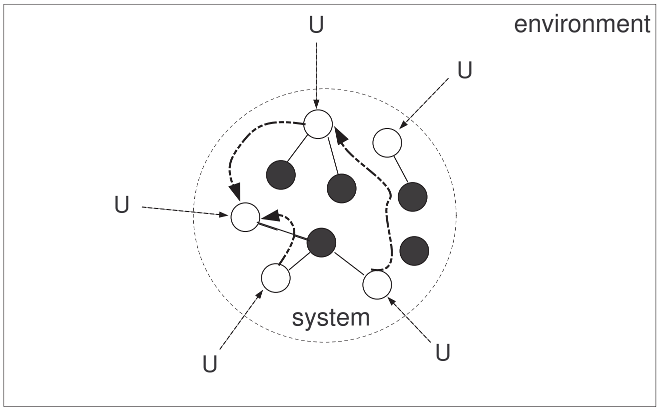

Figure 3: Resource allocation process conducted on the IT system software level.

propagation of information (represented by dotted arrows) is fully decentralised across the population of autonomic elements, understanding how this process may self-organise resource management is non-trivial. Nevertheless, there exists a range of studies focusing on information flows in different decentralised architectures demonstrating that effective decentralised control can be achieved if localised elements organise their information exchange (Packard, 1988; Guerin and Kunkle, 2004; Brueckner and Parunak, 2003). In the next section we will present a thermodynamical interpretation of this kind of self-organisation.

Energy, entropy and work in computational ecosystems

From the considerations outlined above, we might expect to find that the continued efficiency of self-organising computational ecosystems depend on them tending to establish and maintain information flows that bring about informational gradients that constrain agent behavioural choice. Under these constraints, agent behaviour results in efficient resource consumption, doing work for the user, but also re-establishes the constraints that enable the further extraction of work maintaining the system far from equilibrium. In what follows we will elaborate on this picture, focusing on the role of local information exchange.

Equilibrium

Each agent within a computational ecosystem may be characterised by its behavioural *repertoire*, the set of actions that are currently available to it. During each decision-cycle, an agent is required to select one action from the set of available ones and, by executing it, act upon its environment. The behaviour of an agent will exhibit the highest Shannon entropy when selection of any action is equally probable during each decision cycle, and the agent behaves randomly. Since the entropy of the whole population can be measured as the average over individual agent entropies, a multi-agent system can be said to be at equilibrium when all agent decisions are made at random.

Work

Whereas the establishment of an energy gradient is a precursor for useful work in a physical system, here it is useful to consider an *information gradient*. Recall that it is not the mere injection of energy that allows a physical system to perform work, but the *organisation* of this energy, which displaces the system from equilibrium. The same can be said of the distribution of information within a computational system. When one agent is informed such that it can be distinguished from the rest of the system, there is the potential for it to act in a manner that is constrained by its information, perhaps performing useful work. However, in a self-organising system, an agent's actions are liable to propagate information to other agents ensuring that informational disparities tend to be extinguished as they are exploited. Both energy and information flows are the result of local interactions between system components (molecules in a physical system; agents in a computational system) and, under the right conditions, both are 'motive forces' for achieving spontaneous system organisation.

Notice that information flow within a system constrains agent behaviour only if it creates a gradient between that agent and its surroundings. For instance, incoming information must perform work on agents rather than merely raising system temperature. While a constant supply of organised information is required to drive a system far from equilibrium, notice also that if the subsequent propagation of information between agents within the system is also organised, then agents can do work on (organise, inform) one another as they perform useful work for us, extinguishing their own potential gradient in the process. Recall the Bénard cells described earlier. There, molecules of a fluid, across which an energetic gradient is imposed, spontaneously organise each other such that they convey heat more efficiently than would be achieved by a random organisation of molecules.

Entropy

We have seen that information gradients may allow agents to make useful decisions, but that, for a system in flux, any collectively arising gradient informing agents about an available resource will eventually become 'dissipated'. That is, a flow of information about this resource will attract agents to consume it, extinguishing the original gradient, releasing constraints on agent behaviour and increasing system entropy. A computational ecosystem in which information propagates amongst agents is thus one in which there is a tendency for the system to equilibrate to an inefficient, essentially random state. However, it is precisely this tendency for information to propagate that can give rise to the possibility of efficient, persistent, self-organised behaviour.

Without information flows (of the right kind), agents cannot inform one another, organising or constraining each other's behaviour in a manner that is capable of achieving efficient work.

Case studies

Here, we provide three examples of decentralised system architectures (Parunak and Brueckner, 2001; Gambhir et al., 2004; Jacyno and Bullock, 2008), the functioning of which can be interpreted from the thermodynamical viewpoint outlined above. In each case, the local decision-making of individual system elements is achieved through the creation and destruction of gradients, work done on the system and by the system is manifested through the imposition and gradual release of behavioural constraints, and there is an important role for information flow.

Entropy in a two-agent system

A thermodynamic account of self-organisation within a multi-agent system is presented by Parunak and Brueckner (2001). The authors consider a simple coordination problem between two agents who desire to be together, one a mobile walker, the other in a fixed location. Both agents are embedded within a spatial environment with neither knowing the location of the other. The coordination problem for the walker is to locate the other agent and move towards it. An intelligent observer capable to seeing the state of both agents could send instructions to direct the movement of the walker. However, in this model Parunak and Brueckner investigate stigmergic coordination inspired by organisation in insect colonies. For this purpose, the stationary agent deposits pheromone molecules at its location. Initially, the walker is unable to sense any molecules and performs unguided movements. However, once pheromone molecules diffuse through the environment and are detected by the walker, it follows the gradient formed by them, thus reaching the target. We can understand how self-organised system behaviour emerges from the random processes of pheromone molecule diffusion on two levels: a macro-level at which coordinated behaviour of the walker agent arises; and a micro-level represented by a random motion of pheromone molecules that diffuse through the environment. An analysis of system organisation at both levels based on Shannon's entropy reveals that an increase in the micro-level entropy (as pheromone molecules diffuse to occupy an increasing number of locations) is accompanied with a decrease in entropy at the macro-level (as the movement of the walker is increasingly informed by the pheromone gradient).

This simple example illustrates not only how 'intelligent' behaviour emerges from a simple, entropy increasing processes, but also that the resulting self-organisation does not defy the second law of thermodynamics since the price paid for the entropy reduction at the macro system level is the increase in entropy generated by the random process that produces and maintains the gradient.

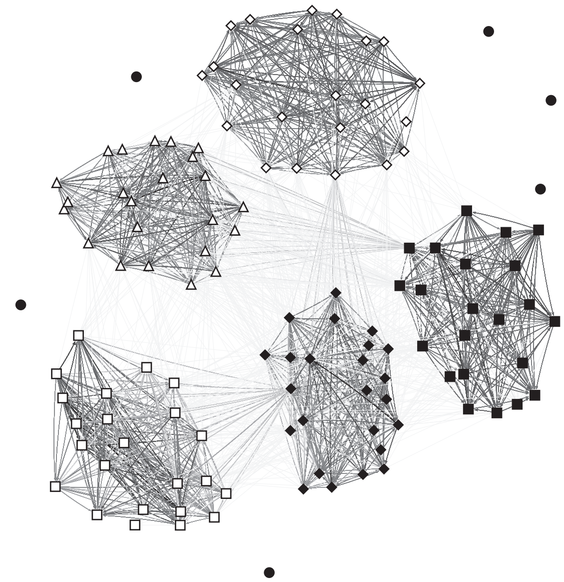

Figure 4: Communities of agents formed as a population of agents self-organise to reliably match consumption and provision of different resource types.

A full population model

A continuation of Parunak and Brueckner's work is presented by Gambhir et al. (2004). Here, the authors apply a computational model of an ant foraging system to demonstrate how complex organisation of interacting agents can be explained in terms of ideas from equilibrium and non-equilibrium thermodynamics. Their analysis of this classic example of self-organisation distinguishes three distinct modes of system behaviour: structure formation, structure maintenance and structure decay. During structure formation, some members of a population of agents diffusing over the environment discover a food source and establish a pheromone distribution instructing other agents to organise their activities into a foraging trail. By maintaining this structure, the population achieves reliable transport of food to the nest. Once the food source becomes depleted, the structure beings to decay and the agents return to their initial disorganised state.

To interpret how the system is displaced from equilibrium and how work is extracted from these conditions, the authors evoke ideas of unconstrained and constrained transfers of energy that are responsible for thermodynamical organisation and work extraction. Within a computational system, unconstrained flow of heat is considered as a diffusive, entropy producing process of agents performing random walks. By contrast, constrained transfer of energy, in the form of interactions with an organised pheromone distribution, is interpreted as work done on agents, constraining their behavioural degrees of freedom (i.e., agent movements are directed to climb the pheromone intensity gradient, as in the case of the walker agent discussed above). The insights drawn from this model are similar to those arrived at by Parunak and Brueckner. An initial increase in entropy, during which agents explore state space, enables the formation of organisation, imposing constraints on agent behaviour through interaction with the pheromone field. To measure construction and destruction of constraints in this self-organising system, Shannon's entropy is applied. The measure of *useful work* done by the system is represented by the number of pieces of food taken from the food-source to nest over a run.

So far we have considered a population of agents that move from thermodynamic equilibrium to a constrained state and then back to equilibrium. However, in order to characterise a computational ecosystem that organises itself such that it remains far from equilibrium, in a dynamic, poised state where constraints are continually formed and released in a reflexive, self-perpetuating manner, we need to go a step further.

A self-organising computational ecosystem

Here, the system is an example of a computational ecosystem consisting of a population of consumer and provider agents responsible for reliable and efficient management of the resources offered by the system. Consumer agents belong to distinct groups, each characterised by the type of resource they are interested in allocating. Providers, on the other hand, are capable of offering any type of resource in general, but at any one time they are configured to offer only one type. As the interaction between agents and reconfiguration of offered resources has an associated cost (time), the efficiency of the system depends on discovering an organisation of agents that maximises the system's allocative throughput, at the same time avoiding unnecessary reconfiguration of providers or resource competition on the consumer side. This is achieved when consumers and providers self-organise into communities within which providers reliably offer a certain type of resource and consumers are biased towards selection of available providers (Jacyno and Bullock, 2008). Example communities are depicted in Figure 4.

Initially, the population of agents is uninformed and behaves randomly: consumers choose providers at random, and propagate information to one another at random. Organisation of the ecosystem into stable communities of agents is achieved through the formation and maintenance of information gradients between agents. These gradients are established through "gossiping", e.g., the local exchange of information about providers by individual consumers. As a consequence of sensing some gradient, an agent's initially unbiased selection of resources becomes constrained (work is done on the agent by the gradient). Agent behaviour is constrained in two ways. First, just as the agents in the previous case studies were able to exploit a pheromone gradient to discover food, here, consumer agents are constrained such that they tend to choose suitable providers. In doing so, they consume resource, and as a side-effect tend to dissipate the gradient that they were informed by. Second, the same gradient constrains agents such that they now tend to propagate information to a non-random sub-set of agents. By organising the information flows that propagate gossip such that agents form communities with shared interests, the system can maintain itself in a far from equilibrium organisation that allows useful work to be undertaken efficiently.

A complete analysis along these lines would clearly be more involved than in the previous examples, since here structure formation, maintenance and decay are ongoing processes that are capable of maintaining global system stability far from equilibrium. In particular, we have had to identify the manner in which the system organises the propagation of information in addition to merely establishing and releasing a constrain in order to achieve a piece of work. Here, we have attempted to lay some groundwork for further analysis of such systems by articulating the way in which thermodynamic ideas can offer a framework that focuses engineers on critical aspects of the system design.

Discussion and Conclusions

The aim of this paper is not to provide a ready-to-apply solution to the control of decentralised IT systems, but to point to and organise important work that has already been done in other research areas focusing on self-organisation and the homeostatic properties of natural systems. If we aim to engineer self-organising IT systems, we must understand the underlying thermodynamic principles of natural self-organisation, and, in particular, how to apply these principles in the context of open IT systems.

We have described how information disparity drives self-organisation in a population of software agents and that random behaviour is an integral part of the maintenance of information flows that allow such a population to organise effectively. This contrasts starkly with the (sometimes implicit) assumption present in the multi-agent system community that software agents share complete knowledge of the system, and make decisions as a result of joint deliberation, or at the behest of a central executive charged with deducing optimal behaviour. This approach is analogous to relying on a kind of maxwell demon to control a computational ecosystem. The demon knows the position and state of every element in the system and is able to impose/remove constraints that allow the system to do useful work. However, thermodynamic considerations imply that, even if such a demon could be implemented, it would be extremely costly.

The interpretation provided here should not be considered exclusive. While thermodynamics and self-organisation have been the object of extensive research, there are still open questions with respect to the application of these ideas to systems that are far from equilibrium but capable of maintaining steady state (Kay, 1984). In such cases, considerations of thermodynamical systems at, close to, or moving towards their equilibrium state are insufficient, making far-from-equilibrium thermodynamics an open and active area of study with direct implications for engineering open computational ecosystems.

References

Brueckner, S. and Parunak, H. V. D. (2003). Self-organizing MANET management. In Marzo, G. D., Karageorgos, A., Rana, O. F., and Zambonelli, F., editors, *Engineering Self-Organising Systems*, pages 1–16. Springer.

Bullock, S. and Cliff, D. (2004). Complexity and emergent behaviour in ICT systems. Technical Report HP-2004-187, Hewlett-Packard Labs.

Gambhir, M., Guerin, S., Kauffman, S., and Kunkle, D. (2004). Steps toward a possible theory of organization. In A. Minai, Y. B.-Y., editor, *Proceedings of International Conference on Complex Systems*, page 9. W. H. Freeman.

Guerin, S. and Kunkle, D. (2004). Emergence of constraint in self-organizing systems. *Journal of Nonlinear Dynamics, Psychology, and Life Sciences*, 8:2:131–146.

Huberman, B. A. and Hogg, T. (1993). In Nadel, L. and Stein, D., editors, *Lectures in Complex Systems*, chapter The Emergence of Computational Ecologies, pages 185–205. Addison-Wesley.

Jacyno, M. and Bullock, S. (2008). Energy, entropy and work in computational ecosystems: A thermodynamic account. In *In Proceedings of the 11th Conference on Artificial Life (ALIFE 2008)*.

Jaynes, E. T. (1965). Gibbs vs Boltzmann entropies. *American Journal of Physics*, 33(5):620–630.

Jaynes, E. T. (1979). Where do we stand on maximum entropy? In Levine, R. D. and Tribus, M., editors, *The Maximum Entropy Formalism*, pages 15–127. MIT Press.

Jaynes, E. T. (1988). The evolution of Carnot's principle. In Erickson, G. J. and Smith, C. R., editors, *Maximum-Entropy and Bayesian Methods in Science and Engineering*, pages 267–284. Kluwer.

Kauffman, S. (2000). *Investigations*. Oxford University Press.

Kay, J. J. (1984). *Self-Organization In Living Systems*. PhD thesis, Deparment of Systems Design Engineering, University of Waterloo.

Kephart, J. O. and Chess, D. M. (2003). The vision of autonomic computing. *IEE Computer*, 36(1):41–50.

Lotka, A. J. (1922a). Contribution to the energetics of evolution. *PNAS USA*, 8:145–151.

Lotka, A. J. (1922b). Natural selection as a physical principle. *PNAS USA*, 8:151–154.

Nelson, P. (2004). *Biological Physics: Energy, Information, Life*. W. H. Freeman.

Nicolis, G. and Prigogine, I. (1977). *Self-organization in Non-equilibrium Systems: From Dissipative Structures to Order Through Fluctuations*. J. Wiley & Sons.

Packard, N. H. (1988). Adaptation toward the edge of chaos. In J.A. Kelso, A. M. and Shlesinger, M., editors, *Dynamic patterns in complex systems*, pages 293–301. World Scientific.

Parunak, H. V. D. and Brueckner, S. (2001). Entropy and self-organization in multi-agent systems. In J. Müller, E. Andre, S. S. C. F., editor, *Proceedings of the fifth international conference on Autonomous agents*, pages 124–130. ACM Press.

Parunak, H. V. D. and Brueckner, S. A. (2004). Engineering swarming systems. In Bergenti, F., Gleizes, M.-P., and Zambonelli, F., editors, *Methodologies and Software Engineering for Agent Systems*, pages 341–376. Kluwer.

Schneider, E. D. and Kay, J. J. (1995). Order from disorder: The thermodynamics of complexity in biology. In Murphy, M. P. and O'Neill, L., editors, *What Is Life: The Next Fifty Years. Reflections on the Future of Biology*, pages 161–172. Cambridge University Press.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423 and 623–656.

Sharma, R. K., Bash, C. E., Patel, C. D., Friedrich, R. J., and Chase, J. S. (2005). Balance of power: Dynamic thermal management of internet data centers. *IEEE Internet Computing*, 9(1):42–49.

Skadron, K., Stan, M. R., Huang, W., Velusamy, S., Sankaranarayanan, K., and Tarjan, D. (2003). Temperature-aware microarchitecture. In *Proceedings of the 30th International Symposium on Computer Architecture*, pages 2–13. IEEE Computer Society.

Swenson, R. (1997). Autocatakinetics, evolution, and the law of maximum entropy production: A principled foundation towards the study of human ecology. *Advances in Human Ecology*, 6:1–47.

Swenson, R. and Turvey, M. (1991). Thermodynamic reasons for perception-action cycles. *Ecological Psychology*, 3(4):317–348.

Sycara, K. (1998). Multi-agent systems. *AI Magazine*, 10(2):79–93.

Wicken, J. S. (1988). Evolution, thermodynamics, and information: Extending the darwinian program. *The Quarterly Review of Biology*, 63(1):84–85.

Wicken, J. S. (1989). Evolution and thermodynamics: The new paradigm. *Systems Research and Behavioral Science*, 6(3):181–186.

Zambonelli, F. and Parunak, H. V. D. (2004). Towards a paradigm change in computer science and software engineering: A synthesis. *The Knowledge Engineering Review*, 18(4):329–342.