Christian Blum · Daniel Merkle (Eds.)

Swarm Intelligence

Introduction and Applications

Springer

Christian Blum · Daniel Merkle (Eds.)

Swarm Intelligence

Introduction and Applications

🦅 Springer

Dr. Christian Blum
ALBCOM
Dept. Llenguatges i Sistemes
Informàtics
Universitat Politècnica de Catalunya,
Jordi Girona 1-3
Omega 112 Campus Nord
08034 Barcelona, Spain
cblum@lsi.upc.edu

Dr. Daniel Merkle
Dept. of Mathematics and Computer
Science
University of Southern Denmark
Campusvej 55
5230 Odense M, Denmark
daniel@imada.sdu.dk

Natural Computing Series
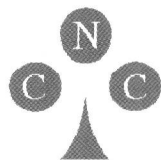
For María and Marc
(Christian Blum)

For Janine
(Daniel Merkle)

Preface

Swarm intelligence is a modern artificial intelligence discipline that is concerned with the design of multiagent systems with applications, e.g., in optimization and in robotics. The design paradigm for these systems is fundamentally different from more traditional approaches.

Instead of a sophisticated controller that governs the global behavior of the system, the swarm intelligence principle is based on many unsophisticated entities that cooperate in order to exhibit a desired behavior. Inspiration for the design of these systems is taken from the collective behavior of social insects such as ants, termites, bees, and wasps, as well as from the behavior of other animal societies such as flocks of birds or schools of fish. Colonies of social insects have mesmerized researchers for many years. However, the principles that govern their behavior remained unknown for a long time. Even though the single members of these societies are unsophisticated individuals, they are able to achieve complex tasks in cooperation. Coordinated behavior emerges from relatively simple actions or interactions between the individuals.

For example, ants, termites and wasps are able to build sophisticated nests in cooperation, without any of the individuals having a global master plan of how to proceed. Another example is the foraging behavior that ants or bees exhibit when searching for food. While ants employ an indirect communication strategy via chemical pheromone trails in order to find shortest paths between their nest and food sources, bee colonies are very efficient in exploiting the richest food sources based on scouts that communicate information about new food sources by means of a so-called waggle dance. For more examples and a more detailed description of the fascinating biological role models that inspired swarm intelligence applications see Chaps. 1 and 2 of this book.

Scientists have applied these principles to new approaches, for example, in optimization and the control of robots. Characterizing properties of the resulting systems include robustness and flexibility. The field of research that is concerned with collective behavior in self-organized and decentralized systems is now referred to as *swarm intelligence*. The term swarm intelligence was first used by Beni and colleagues in the context of cellular robotic sys-

tems where simple agents organize themselves through nearest neighbor inter-actions. Meanwhile, the term swarm intelligence is used for a much broader research field, as documented in the seminal book *Swarm Intelligence—From Natural to Artificial Systems* by Dorigo, Theraulaz, and Bonabeau, published by Oxford University Press. However, since the appearance of the above-mentioned book in 1999, the literature on swarm intelligence topics has grown significantly. This was the motivation for editing this book, whose intention is to provide an overview of swarm intelligence to novices of the field, and to provide researchers from the field with a collection of some of the most interesting recent developments. In order to achieve this goal we were able to convince some of the top researchers in their respective domains to write chapters on their work.

Introductory chapters in the first part of the book are on biological foundations of swarm intelligence, optimization, swarm robotics, and applications in new-generation telecommunication networks. Optimization and swarm robotics are nowadays two of the domains where swarm intelligence principles have been applied very successfully. A third and very popular application domain concerns routing and loadbalancing in telecommunication networks. The second part of the book contains chapters on more specific topics of swarm intelligence research such as the evolution of robot behavior, the use of particle swarms for dynamic optimization, organic computing, and the decentralized traffic flow in production networks.

Finally, we hope that the readers enjoy reading this book, and, most importantly, that they learn something new by seeing things from a new perspective.

Barcelona, Odense   *Christian Blum*
April 2008   *Daniel Merkle*

Contents

Part I

Introduction

Biological Foundations of Swarm Intelligence

Madeleine Beekman[1], Gregory A. Sword[2], and Stephen J. Simpson[2]

[1] Behaviour and Genetics of Social Insects Lab, School of Biological Sciences,
University of Sydney, Sydney, Australia
`mbeekman@bio.usyd.edu.au`
[2] Behaviour and Physiology Research Group, School of Biological Sciences,
University of Sydney, Sydney, Australia
`{greg.sword,stephen.simpson}@bio.usyd.edu.au`

Summary. Why should a book on swarm intelligence start with a chapter on biology? Because swarm intelligence is biology. For millions of years many biological systems have solved complex problems by sharing information with group members. By carefully studying the underlying individual behaviours and combining behavioral observations with mathematical or simulation modeling we are now able to understand the underlying mechanisms of collective behavior in biological systems. We use examples from the insect world to illustrate how patterns are formed, how collective decisions are made and how groups comprised of large numbers of insects are able to move as one. We hope that this first chapter will encourage and inspire computer scientists to look more closely at biological systems.

1 Introduction

"He must be a dull man who can examine the exquisite structure of a comb so beautifully adapted to its end, without enthusiastic admiration."

Charles Darwin (1872)

When the Egyptians first started to keep honeybees 5,000 years ago, they surely must have marveled on the beauty of the bees' comb. Not only is the honeycomb beautiful to look at, but how did the bees decide to build hexagonal cells and not cells of another form? Initially it was suggested that hexagonal cells hold the most honey, but the French physicist R.A.F. de Réaumur realized that it was not the content of the cells that counts, but the amount of material, wax, that is needed to divide a given area into equal cells. Obviously at that time it was assumed that the bees were "blindly using the highest mathematics by divine guidance and command" (Ball 1999). It was not until Darwin that the need for divine guidance was removed and the hexagonal cells were thought to be the result of natural selection. In this view the bees' ancestors 'experimented' with different shaped cells, but the bees that by chance

'decided' to build hexagonal cells did better and, as a result, the building of hexagonal cells spread. In Darwin's words, *"Thus, as I believe, the most wonderful of all known instincts, that of the hive-bee, can be explained by natural selection having taken advantage of numerous, successive, slight modifications of simpler instincts; natural selection having by slow degrees, more and more perfectly, led the bees to sweep equal spheres at a given distance from each other in a double layer, and to build up and excavate the wax along the planes of intersection."* (Chapter 7, Darwin 1872).

It was exactly such 'Darwinian fables' that inspired the biologist and mathematician D'Arcy Wentworth Thompson to write his book *On Growth and Form* (Thompson 1917). The central thesis of this book is that biologists overemphasize the role of evolution and that many phenomena can be more parsimoniously explained by applying simple physical or mathematical rules. Thompson argued that the bees' hexagonal cells are a clear example of a pattern formed by physical forces that apply to all layers of bubbles that are pressed into a two-dimensional space. Bees' wax is not different, the soft wax forms bubbles that are simply pulled into a perfect hexagonal array by physical forces. Hence, the pattern forms spontaneously and no natural selection or divine interference needs to be invoked (Ball 1999).

In fact, many instances of spontaneous pattern formation can be explained by physical forces, and given the almost endless array of patterns and shapes found around us, it is perhaps not surprising that such patterns are an inspiration for many people, scientists and non-scientists alike. Upon closer examination, amazing similarities reveal themselves among patterns and shapes of very different objects, biological as well as innate objects. As we already alluded above, the characteristic hexagonal pattern found on honeycombs are not unique; the same pattern can be obtained by heating a liquid uniformly from below. Autocatalytic reaction-diffusion systems will lead to Turing patterns (think stripes on tigers) in both chemical and biological mediums (Kondo and Asai 1995; Ball 1999), and minerals form patterns that have even been mistaken for extra-terrestrial fossils (McKay et al. 1996).

The similarity of patterns found across a huge range of systems suggests that there are underlying principles that are shared by both biological and innate objects. Such similarities have been nicely illustrated by work on pattern formation in bacterial colonies. When one manipulates the amount of food available to bacteria and the viscosity of their medium, patterns emerge that are remarkably similar to those found in, for example, snowflakes (Ben-Jacob et al. 2000). In fact, the growth of bacterial colonies has proven to be an important playground for testing ideas on non-living branching systems (Ball 1999; Ben-Jacob and Levine 2001; Levine and Ben-Jacob 2004). As it turns out, many branching patterns found across nature can be explained by the same process, known as diffusion-limited aggregation, resulting from the interactions of the particles, be they molecules or individual bacteria (Ball 1999).

All patterns described above have been explained by approaching the systems from the bottom up: how do the particles interact with each other and with their immediate environment? One may not really be surprised by the fact that the same approach helps one to understand bacteria as well as molecules. After all, bacteria aren't really that different from molecules, are they? In the following we will illustrate how such a bottom-up approach can explain another remarkable feature of honeybees: the typical pattern of honey, pollen and brood found on combs.

The honeybee's comb is not only a marvel because of its almost perfect hexagonal cells, the bees also seem to fill the cells with brood (eggs that develop into larvae and then pupae and finally emerge as young workers or males), pollen (to feed the brood) and nectar (which will be converted into honey) in a characteristic pattern. This pattern consists of three distinct concentric regions: a central brood area, a surrounding rim of pollen, and a large peripheral region of honey (Fig. 1). If we envision the honeybee colony as a three-dimensional structure, this pattern is most pronounced in the central combs which intersect a large portion of the almost spherical volume of brood. How does this pattern come about? The storage of pollen close to the brood certainly makes sense as it reduces the time needed to get the pollen to the brood. But how do the bees know this? Do they use a blueprint (or template) to produce this characteristic pattern, implying that there are particular locations specified for the deposition of pollen, nectar and brood? Or is the pattern self-organized and emerges spontaneously from the dynamic interactions between the honeybee queen, her workers and the brood? Scott Camazine set out to determine which of these two hypotheses is the most parsimonious (Camazine 1991).

The beauty of working on macroscopic entities such as insects is that you can individually mark them. Honeybees are particularly suitable because we can then house them in what we call an observation hive, a glass-walled home for the bees. This means that we can study the interactions of the individually marked bees without taking them out of their natural environment (see Fig. 1).

Camazine did just that. He monitored the egg-laying behavior of the queen, of foragers that returned with pollen or nectar, and of nurse workers, those that feed the brood. The first thing that he observed was that the queen is rather sloppy in her egg-laying behavior, moving about in a zig-zag-like manner, often missing empty cells and retracing her own steps. Camazine further noticed that she has a clear preference to lay a certain distance from the periphery of the comb and never more than a few cell lengths of the nearest brood-containing cell. Interestingly, even though the queen somewhat has a preference for at least the middle of the comb and the vicinity of brood, bees returning with pollen or nectar did not seem to have a preference for specific cells at all. When an empty comb was left in the colony and the deposition of nectar and pollen observed, both could be found in any cell. Even though such absence of a preference clearly refutes the blueprint hypothesis, it does not explain how the characteristic pattern ultimately arises.

Fig. 1. Because of their relatively large size, we can easily mark individual bees in a colony. In this particular colony we marked 5,000 bees by combining numbered plates and different paint colors. This allowed us to study their behavior at an individual level. Photograph taken by M. Beekman.

As it turns out, bees do have a clear preference when they *remove* pollen or honey from cells. Both honey and pollen are preferentially removed from cells closest to the brood. By following the pattern of cell emptying during a period in which foraging activity was low (overnight or during rain), Camazine observed that all the cells that were emptied of their pollen or nectar were located within two cells or less from a cell containing brood. No cells were emptied that were further from brood cells. It is easy to see why the bees would have a preference for the removal (through use) of pollen that is found closest to the brood, as it is the brood that consumes the pollen. In addition, nurse bees are the younger bees which restrict most of their activity to the brood area (Seeley 1982).

The preferential removal of pollen and nectar from cells closest to cells containing brood and the queen's preference for laying eggs in cells close to brood made Camazine realism that this might explain the honeybee's characteristic comb pattern. But how to prove this? This is where the physicist's approach comes in. By constructing a simulation model based on his behavioral observations, Camazine was able to closely follow the emergence of the pattern. Initially, both pollen and nectar were deposited randomly throughout the frame with the queen wandering over the comb from her initial starting point. Despite the random storage of pollen and nectar, the queen's tendency to lay eggs in the vicinity of cells that already contain brood rapidly results in an area in which mostly brood is found. This is enhanced by the bees' pref-
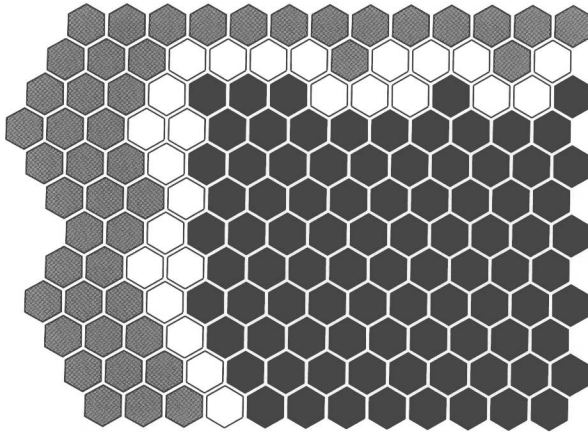
Fig. 2. The typical pattern of honey (grey cells), pollen (white cells), and brood (black cells) as seen on a honeybee's comb. Shown is the top-left corner of the comb

erence to remove honey and pollen from cells close to brood, which increases the availability to the queen of cells to lay eggs in. This further reduces the number of cells available for storage of honey or pollen. Thus, the brood area is continually freed of honey and pollen and filled with eggs resulting in a compact brood structure. But how do the pollen and nectar get separated (Fig. 2)?

Because initially both are deposited randomly, both pollen and nectar will be present in the periphery of the comb. However, most pollen that gets collected on a daily basis is consumed that same day. This means that given the normal fluctuations in pollen availability, there is often a net loss of pollen, with pollen present in the periphery being consumed at nearly the same rate as pollen being stored elsewhere. At the same time, these empty cells are most likely to be filled with nectar, as the nectar intake is much higher, and soon there is no longer space to store pollen. Where is pollen stored then?

Eventually the only place left for pollen to be stored is the band of cells adjacent to the brood. The developmental time from egg to adult is 21 days, meaning that for three weeks a brood cell cannot be used for anything else. But in the interface zone between the brood and the stores of honey at the periphery, the preferential removal of honey and pollen continuously provides a region in which cells are being emptied at a relatively high rate. And it is these cells that are available for pollen. Other cells that become available because bees emerge from them are found in the middle of the brood nest, but these will then be preferentially emptied and again filled with eggs.

Without his computer simulation Camazine would not have been able to fully understand how the behavior of the individual bees resulted in the organized pattern of brood, honey and pollen on the comb of the bees. And this

is a general principle for understanding collective animal behavior: without tools such as simulations or mathematics, it is impossible to translate individual behavior into collective behavior. And it is exactly with those tools that originally came from disciplines outside of biology, and with the view that interactions among individuals yield insights into the behavior of the collective, that we biologists have learned from physics. In fact, we began this chapter by illustrating that even biological phenomena can often more parsimoniously be understood using physical explanations, and that many systems, both innate and living, share the same physical principles. And it has exactly been these similarities and the wide applicability of the mathematical rules that govern diverse behaviors that have led to the field of Swarm Intelligence (e.g. Dorigo et al. 1996; Dorigo and Di Caro 1999).

However, it is important to realise that our biological 'particles' are more complex than molecules and atoms and that the 'simple rules of thumb' of self-organization (Nicolis and Prigogine 1977) have only limited explanatory power when it comes to biological systems (Seeley 2002). Bacterial colonies may grow in a similar pattern as minerals, Turing patterns may be found on fish, in shells and in chemical reactions, and we can understand the bees' hexagonal cells using physics, but when it comes to biological systems, an extra layer of complexity needs to be added. Besides the complexity of the individuals, we cannot ignore natural selection acting on, for example, the foraging efficiency of our ant colony, or the building behavior of our termites. If the underlying principles that govern the building behavior of termites results in colony-level behavior that is far from functional, this would be rapidly selected against. Moreover, it is of no use to assume that certain systems must behave similarly simply because they 'look' similar. It is true that if the same mathematical model or behavioral algorithm captures the behavior of different systems, then we can talk about similarities between systems that go beyond simple analogy (Sumpter 2005). However, as we will explain in the concluding section of this chapter, true biological inspiration needs to come not from the superficial similarities between systems, but from the intricate and often subtle differences between them. We shall illustrate this standpoint by drawing examples from our own study systems: decentralized decision making in social insects and the coordinated movement of animal groups.

2 Decentralized Decision Making

The evolution of sociality, the phenomenon where individuals live together within a nest such as is found in many bees and wasps, and all ants and termites, has created the need for information transfer among group members. No longer can each individual simply behave as if solitary, but actions by different group members need to be carefully tuned to achieve adaptive behavior at the level of the whole group. Insect colonies need to make many collective decisions, for example where to forage, which new nest to move to, when to

reproduce, and how to divide the necessary tasks among the available work-force. It is by now well known that such group-level decisions are the result of the individual insects acting mainly on local information obtained from interactions with their peers and their immediate environment (Bonabeau et al. 1997; Camazine et al. 2001). In other words, decision making in insect societies is decentralized. To illustrate how insect colonies achieve this, we will describe foraging and nest site selection in ants and honeybees.

2.1 Where to Forage?

In order to organize foraging, social insects need a form of recruitment. Recruitment is a collective term for any behavior that results in an increase in the number of individuals at a particular place (Deneubourg et al. 1986), and allows insect societies to forage efficiently in an environment in which food sources are patchily distributed or are too large to be exploited by single individuals (Beckers et al. 1990; Beekman and Ratnieks 2000; Detrain and Deneubourg 2002). In addition, social insects that transfer information about the location of profitable food sources can exploit an area much larger than those that lack such a sophisticated recruitment mechanism. Honeybees are a prime example. Their sophisticated dance language (von Frisch 1967) allows them to forage food sources as far as 10 km from the colony (Beekman and Ratnieks 2000).

Exact recruitment mechanisms vary greatly among the social insects but can be divided into two main classes: direct and indirect mechanisms. Mass recruitment via a chemical trail is a good example of indirect recruitment. The recruiter and recruited are not physically in contact with each other; communication is instead via modulation of the environment: the trail. The recruiter deposits a pheromone on the way back from a profitable food source and recruits simply follow that trail. In a way such a recruitment mechanism is comparable to broadcasting: simply spit out the information without controlling who receives it. The other extreme is transferring information, figuratively speaking, mouth to mouth: direct recruitment. The best-known example of such a recruitment mechanism is the honeybees' dance language. Successful foragers, the recruiters, perform a stylized 'dance' which encodes information about the direction and distance of the food source found and up to seven dance followers (Tautz and Rohrseitz 1998), potential recruits, are able to extract this information based upon which they will leave the colony and try to locate the advertised food source. Recruitment trails and the honeybee dance language can be seen as the two extremes of a whole range of different mechanisms used by social insects to convey information about profitable food sources.

Many computer scientists are familiar with the double bridge experiment as an example of the means by which foraging is organized in ant colonies. In this experiment a colony of trail-laying ants is offered two equal food sources located at the end of two paths of different lengths. After some time the