



Contents lists available at ScienceDirect

Theoretical Computer Science

journal homepage: www.elsevier.com/locate/tcs

Swarm intelligence theory: A snapshot of the state of the art

Nature offers us many interesting and surprising examples in which the behaviour of a group of organisms seems to have some fundamentally distinct characteristics, not shared by the individuals in that group. Different species of birds flock together, and species of fish form schools, in groups that vary in size from a handful to many millions. Meanwhile it is well-known that most species of ants, bees and termites form swarms that perform many functions collectively, including hunting and gathering food, and building complex structures. In different scenarios, these groups may be called herds, flocks, schools, and so forth, but the convenient term that stands for all such cases is 'swarm'. The concept of 'swarm intelligence' captures the interest of many groups of (indeed, swarms of) academics and scientists, encapsulating the idea that the behaviour of a swarm often exhibits useful, functional and intelligent behaviours which seem well beyond the ability, as far as we know, of any of the individuals that together constitute the swarm.

Swarm intelligence therefore concerns systems in which a group of similar 'agents', each of which is relatively simple in its behavioural repertoire, is somehow coordinated in a way that leads to useful ('emergent') behaviour of the swarm itself. Certain structural aspects of swarms are commonly assumed when swarm intelligence is discussed: apart from the aforementioned 'simplicity' of the constituent agents, we also expect that a swarm has no central controller or 'master' agent that conducts the activities of others. Each agent is independent, but interacts with its fellow swarm-members (and other aspects of its environment) in simple ways. The fact that such a system can lead to interesting, useful and robust behaviour is in itself one of the appealing points that makes swarm intelligence an area of intense current study. This is partly because it suggests how we might build real-world systems of various types, that are more robust to damage and/or easier to construct or deploy than alternatives that rely on sophisticated central controllers. Meanwhile, of course, some swarm intelligence studies are undertaken with the goal of better understanding swarms in nature.

Some of the most useful outcomes from swarm intelligence research for computer science are a collection of novel optimization algorithms. Inspired in turn by the flocking behaviour of birds, and the pheromone-trail following behaviour of ants, *particle swarm optimization* (PSO) and *ant colony optimization* (ACO) have both found considerable success in addressing a wide range of optimization problems. The growth of interest in these algorithms, as well as other themes in swarm intelligence (such as collective robotics, foraging algorithms, swarm simulation, and more), presents a number of specific challenges for theoretical work.

We can conveniently class the theoretical challenges into two kinds. First, given specific new algorithm designs, we need to develop an understanding of their scalability, their convergence properties, the interactions between their parameters, and so forth – i.e., the typical range of issues that face us when a new algorithm is mooted, irrespective of its origins. Naturally, we may object that such work is really necessary or worthwhile for many algorithms, but in the case of algorithms that have proven excellent in practical and empirical work on real problems – which is certainly the case here – such work is clearly warranted, and has a lot to contribute when done well. The second kind of challenge concerns understanding the special features of swarm-intelligence-based methods in particular. Examples of questions in this area are: What is special about the nature-inspired interactions in particle swarm optimization that seems to accelerate and improve search on some problems? What particular role does (an abstraction of) pheromone laying/following in social insects play, in the context of an optimization algorithm that integrates this with heuristics and local search?

The present state of understanding, with regard to both kinds of questions, is still relatively immature. There have been many attempts already to undertake (for example) runtime analyses of simplified versions of ant colony optimization, for simple optimization landscapes. For ant colony optimization in particular, Dorigo and Blum [1] provides a comprehensive recent account of theoretical progress. However much remains to be done before we have built up a strong backbone of such theoretical results that underpin a real understanding of ant colony optimization's capabilities, and of how to set its parameters for particular cases. Similar things can be said of particle swarm optimization, and there is very little theoretical understanding so far of other algorithms in the swarm intelligence community. When it comes to understanding the specific benefits that arise from the novel features of swarm intelligence algorithms, again there has been very little progress so far overall. However there have been clear notable developments, such as Poli [3,4], which studies

the particular distribution of new sample points in the search space that emerges from the dynamics of particle swarm optimization.

As we will see, the papers in this special issue each add to this incomplete but growing body of knowledge, with contributions of both kinds – characterizations of the performance, convergence, and scalability of swarm intelligence algorithms, as well as focused insights into the nature and role of the special aspects of these algorithms that arise from the source natural inspiration. We now outline the contents of the issue, paper by paper.

In the first paper, by Sudholt and Witt, we find a state of the art theoretical analysis of a specific algorithm in the particle swarm optimization family: the binary particle swarm optimizer, which operates in a binary space (i.e., it optimizes over real-valued functions of bitstrings). This is the first attempt to develop a theoretical understanding of binary PSO, and Sudholt and Witt make substantial progress. They prove that if practitioners keep the v_{\max} parameter fixed as problem size grows, then results on larger problems can be greatly undermined; following on from this, they define a new formulation of binary PSO in which that parameter is adjusted according to the dimensions of the problem, and prove that this leads to efficient optimization. They move on to present lower bounds on runtime, and then, after showing how fitness-level arguments (proven effective in the analysis of certain evolutionary algorithms) can be applied to this case, they use such arguments to analyse the performance of binary PSO on unimodal functions. In particular, they prove that a simplified variant, 1-PSO, is competitive with more sophisticated evolutionary algorithms, and verify this finding empirically.

Given the prominence of PSO in current research and practice in the area of swarm intelligence, it is no surprise at all that this is also the topic of the next paper in this special issue, by Chen and Jiang. With a focus on how the particles in a PSO algorithm interact with each other, Chen and Jiang bring to bear the notion of using a statistical mechanics interpretation of the particle swarm. In this interpretation, the current state of the swarm can be considered as a ‘macrostate’, which in turn is a statistical abstraction of the detailed descriptions of each particle. The PSO algorithm then assumes a trajectory through the space of macrostates, which each characterize the positions and velocities of individual particles in terms of a distribution. After describing a way to model PSO in this way, Chen and Jiang move on to show how progress and convergence results can be sought in this formulation, and, in particular, they demonstrate conditions for convergence. Chen and Jiang’s approach is the first attempt to analyse PSO in this way, and considers a PSO algorithm which involves only ‘social interaction’ (i.e., particle parameters are updated only according to details of neighbouring particles, with no element involving the particle’s own history). In this way the statistical abstractions focus on modelling the interaction between particles. Chen and Jiang claim that the main advantage of their approach, in contrast to other approaches used so far in theoretical understanding of PSO, is that their framework does not assume fixed attractors in the search space; instead, considering the swarm itself as a position in a space of macrostates allows the developing influence of the objective function to be taken into account.

The next paper in the special issue takes us across to the other most prominent optimization algorithm in the swarm intelligence family: ant colony optimization (ACO). More generally, ACO and several other developments in swarm intelligence are characterized by an inspiration from *stigmergy* – the various mechanisms of indirect communication via which (mainly) social insects collaborate to perform complex tasks. In the paper by Vrancx, Verbeeck and Nowé, the authors consider the task of determining the global system behaviour that arises from local stigmergic interactions. The novel idea in this paper is to adopt game theory as the tool for analysing stigmergic interactions. They are able to show that a system in which agents are coordinated by stigmergic interactions can be approximated by what they call a limiting ‘pheromone game’. By casting a stigmergy-based system of agents in such a way, the authors can then use established techniques from game theory to describe the system’s properties. They demonstrate that, within this framework, the long term behaviour of a simple pheromone system can be determined, and they also show more sophisticated applications in which multiple colonies of agents can be characterized in terms of a ‘limiting colony game’. Although the work in this paper offers a novel framework for analysing systems in which agents are coordinated via stigmergic interactions, the authors point out that it is not straightforward to apply this framework directly to ant colony optimization algorithms, which usually incorporate sophistications such as local search and heuristic bias that confound the analysis. Nevertheless, further extension of this framework to incorporate these other aspects of ant colony optimization is an interesting and open challenge, while it remains the case that this new analysis framework could be useful in the design of such algorithms, since it provides a fresh perspective on the dynamics of their key stigmergy-based components.

The next paper in this special issue provides some theoretical progress in the understanding of the (relatively) ‘up and coming’ technique in swarm intelligence, concerning methods inspired by bacterial foraging. Passino’s seminal paper [2], described an algorithm for optimization inspired directly by the way in which colonies of bacteria of certain types seek and exploit food sources. In this process, known as ‘chemotaxis’, individual bacteria move under the influence of a chemical gradient indicating the presence of nutrient (or away from a toxin), but maintain a degree of noisy/random movement, especially when the concentration of nutrient is very low. Meanwhile, bacteria also seem to attract their sisters, and in particular a local concentration of bacteria will attract others. As well as echoing these elements, the classic bacterial foraging optimization algorithm also incorporates a mechanism inspired by bacterial reproduction, in which the better-fed bacteria (i.e., better positioned in terms of nutrient availability) will reproduce, hence adding greater resource in nutrient-rich areas; similarly, poorly fed individuals will be eliminated. In the paper by Biswas, Das, Abraham and Dasgupta, the authors focus on the main steps of the bacterial foraging optimization algorithm, with a particular focus on the dynamics induced by cycles of reproduction and chemotaxis. By setting up a simple two-bacterium system on a one-dimensional landscape, the authors identify conditions under which the main steps of the bacterial foraging algorithm lead to stability (settling towards the optimum position) or instability (diverging from the optimum position) when both bacteria are close to the optimum.

This is a new style of analysis in bacterial foraging optimization, which to date has received little attention from theorists, and the authors needed several simplifications in order to make progress. However the outcomes provide some potential directions for variants on the bacterial foraging algorithm that introduce control actions for achieving desired behaviour in certain situations; this paper also lays a foundation for extending the analysis towards higher dimensional landscapes and multi-bacteria swarms.

The final paper in this special issue presents a fresh perspective in swarm intelligence research; it is not in the vein of a traditional theory paper, and nor is it an empirical account of swarm intelligence optimization algorithms. Instead, Gökçe and Şahin investigate flocking behaviour, one of the key biological inspirations of swarm intelligence algorithms. Although flocking behaviour is one of the inspirations behind several optimization algorithms that have emerged in swarm intelligence, including particle swarm optimization, its roles in nature are (or so we believe) quite different; flocking is most usually associated with migration, and in that context many have theorized that flocks can achieve greater accuracy in their migration activities than can individuals. Gökçe and Şahin provide a systematic study of this hypothesis via both computer simulations and swarm robotics experiments. Specifically they investigate the performance of flocking as a function of four factors, which include the tendency for individuals to try to align with their neighbours, sources of noise, and the diversity of individuals in the flock. A broad summary of their findings is that they confirm the well-known ‘many wrongs’ principle, which indicates that a flock’s migration accuracy benefits from the averaging out of error over several individuals in a flock. However, they are also able to shed light on observations in nature that seem to confound this principle when small flocks are involved. By analysis of how the four factors studied affect the overall flock performance, this paper has made a step towards a better understanding of elements of flocking in nature, as well as the exploitation and control of flocking in swarm robotics (also called collective robotics) applications. This also presents an opportunity for new theoretical accounts of flocking behaviour, where models are influenced by what Gökçe and Şahin have found out about the influence of the four factors studied.

Thereby, this special issue presents a snapshot of ongoing theory-centred and ‘foundations’-centred research in swarm intelligence. Naturally the papers herein are a small sample. Without going into the intricacies of Student’s T distribution, it is sensible to note that it would be unwise to regard these as a highly representative sample of the real distribution of current research in swarm intelligence theory; however there are certain elements that are reliably reflective. In particular, the strong focus on swarm intelligence optimization algorithms in practical research is reflected by the relatively high number of theoretical attempts to understand the most prominent algorithms in this field, namely particle swarm optimization and ant colony optimization. Hence it is fitting that two of the five papers in this special issue concern particle swarm optimization, and another concerns (via analysing its signature component) ant colony optimization. The other two papers represent areas which are continual satellite themes in both swarm intelligence theory and practice. First, there are several other algorithms in the swarm intelligence arena that are less studied than the main two mentioned, yet promising and intriguing from the viewpoints of their design and the empirical experience so far; bacterial foraging optimization, the topic of the fourth paper, is an excellent example in this theme. Finally, the paper that investigates robotic and simulated flocking behaviour is representative of a continuing thread of research that analyses matters close to the sources of natural inspiration that underpin swarm intelligence.

In concluding, we wish to state that we are most grateful indeed to the several academics who supported us by providing reviews of papers in this special issue.

References

- [1] M. Dorigo, C. Blum, Ant colony optimization theory: a survey, *Theoretical Computer Science* 344 (2–3) (2005) 243–278.
- [2] M. Passino, Biomimicry of bacterial foraging for distributed optimization and control, *IEEE Control Systems Magazine* (2002) 52–67.
- [3] R. Poli, Dynamics and stability of the sampling distribution of particle swarm optimisers via moment analysis, *Journal of Artificial Evolution and Applications* (January) (2008).
- [4] R. Poli, Mean and variance of the sampling distribution of particle swarm optimizers during stagnation, *IEEE Transactions on Evolutionary Computation* 13 (4) (2009) 712–721.

Eric Bonabeau
Icosystem, Cambridge, MA, USA

David Corne*
Heriot-Watt University, Edinburgh, UK
E-mail address: dwcorne@macs.hw.ac.uk.

Joshua Knowles
University of Manchester, Manchester, UK

Riccardo Poli
University of Essex, Colchester, UK

* Corresponding editor.