Enhancing the Effectiveness of Simple Multi-Agent Systems through Stigmergic Coordination

Sorin C. NEGULESCU*, Boldur E. BĂRBAT**

"Lucian Blaga" University of Sibiu, Dept. of Computer Science and Automatic Control, 4 Emil Cioran Street 2400 Sibiu, ROMANIA;

**Phone: +40-745-209952; Fax: +40-269-212716; E-mail: sorin_negulescu@yahoo.com*
***Phone: +40-269-212523; Fax: +40-269-212716; E-mail: bbarbat@rdslink.ro*

The challenge is to accept a solution that you do not understand

CHRISTOPHER G. LANGTON

***Keywords*:** Multi-Agent Systems (MAS), Stigmergic Coordination (SC), Synergy, Optimisation, Travelling Salesperson Problem (TSP), Ant Systems (AS).

Abstract

After detailing the terminology employed, the paper explains the rationale of choosing stigmergic coordination to enhance the problem-solving power of simple MAS, emphasizing the engineering advantages (very simple entities, unaware of each other). To be affordable on usual configurations, a reduced number of agents and a problem-class of manageable complexity are needed. To be comparable against related work, an effectiveness measure is defined and two reference usual algorithms are chosen. All tests apply to the *Travelling Salesperson Problem* (TSP) solved with variants of the Elitist Ant System (*EAS*). The paper follows two paths: searching for local enhancements (based on the biological model) and creating a problem-solving method for TSP (searching for inter-paradigmatic synergy). The first path proved useful (from about 45 algorithm instances tested out, 22 gave relevant results): effectiveness is visible (albeit not much) enhanced by fine-tuning the *EAS*. The second (more promising) path is based on adding symbolic processing factors (adapting the environment and instituting a limited central coordination). Since now the added factors are controlled externally, quantitative evaluations are missing but improvements are apparent. The paper concludes that stigmergic coordination improves effectiveness on affordable configurations with simple MAS, classical problems and usual benchmarks.

1. Introduction

The concept of "stigmergy" is used in its initial meaning proposed by Grasse [9], to characterise the type of interaction taking place in biological insect societies. (In observing ant colonies, Grasse identified a coordination mechanism, based on the creation and placement of a dissipative field of smelling substances – the ant pheromones – in the environment; such "stigmas" alter the environment for other ants and influence their behaviour.) However, especially in the syntagma "stigmergic coordination" (SC) related to Multi-Agent Systems (MAS), it "describes a form of asynchronous interaction and information exchange between agents mediated by an "active" environment" [20], or "the production of certain behaviour in agents as a consequence of the effects produced in the local environment by previous behaviour" [19] (cited in [21]). In this context: "the agents are simple, reactive, and unaware of other agents or of the emerging complex activities of the agent society; the environment is an important mechanism to guide activities of these agents and to accumulate information about ongoing activities of the whole agent society" [20]. Therefore, ants – natural or artificial, alike – are too trivial to be genuine agents since they lack two fundamental agentsy features: *social ability* (being unable to interact with their peers and with humans) and *pro-activeness* (being unable to manifest teleological behaviour by taking initiative – since specialized knowledge is the source of goals and initiatives). Thus, they behave rather as robots than as agents. Yet, the *system* they belong to is not a "multi-*robot*" system but a "multi-*agent*" one; the marvel is due to the synergistic effect of their interaction: beyond the individuals (ant-like entities), the team (society) comes out [3] [11]. However, that is the pre-terminological meaning of "synergy", due to Aristotle: the whole is stronger that the sum of its parts. (Hence, at least in this paper, the term "synergy" is preferred to the more usual one: "emergent synthesis".) Thus, SC is not only a research field per se, but also a test bed for exploiting inter-paradigmatic synergy.

This paper – based on a diploma-work for graduating in Computer Science [10], having the co-author as adviser – aims at illustrating how SC can better problem-solving capacities of simple MAS, affordable on usual configurations. Its rest is organised as follows: Section 2 presents related work. Section 3 shows the rationale and the approach. On this ground, the next two sections describe search branches: Section 4 deals with local, minor, but clear-cut and measurable enhancements, while Section 5 examines the less quantifiable synergistic effect of some new factors - stemming from the symbolic paradigm – adapted to the biological model. Section 6 focuses on implementation aspects and comments upon the results. Conclusions and intentions close the paper.

2. Related work

Most basic work regarding SC, as well as its main application areas (e.g., adaptive routing in communication networks, combinatorial optimisation problems, flexible manufacturing systems) was referred to recently in [12] [20] [21]. Robots based on the model of desert ants[1] are discussed in [15] and [22]. Thus, to impair redundancy, here is presented only recent work directly related to the topics of this paper.

In all early *Ant System* (*AS*) algorithms, ants construct (candidate) solutions based on two main components: pheromone trails and problem-dependent heuristic information. These algorithms have suffered frequent modifications in order to improve their efficiency. Thus, the *AS* [6] developed into the *Elitist Ant System* (*EAS*) [4], because each ant that finds a better solution has the chance to deposit more pheromone. Other systems that emerged from the AS are the *Ant-Q* [8] where the deposited pheromone amount is directly proportional to the quality of the found solution, the *Max-Min Ant System* (*MMAS*) [18]. Dynamic TSP problems are approached with *AS* in [7] while the quadratic assignment problem is addressed in [17].

3. Rationale and approach

The first reason is obvious: any paradigm still in its syncretic stage is promising for both research and applications. Beyond this, there are other reasons – somehow interrelated but distinct: a) For ants, it works remarkably [9] [19]. b) Almost any biologically inspired model has proved to be useful to applied artificial intelligence (AI) [1]. c) SC follows the trend of other sub-symbolic paradigms (two well-known instances: artificial neural networks and evolutionary algorithms) [12]. d) Moreover, it can be regarded as very close to the physical-grounding (ethological) paradigm [13] [14] (although the agents are extremely simple). e) It offers a good test-bed for a (rather heterodox) idea: the strength of synergy seems to be proportional not only to the scale of parallelism itself (number of entities involved) but also to the extent of sub-symbolic depiction [3].

From an "Engineering in Intelligent Systems" point of view, there are other significant advantages of SC [12] [20]: a) *Global* information is made available *locally*. b) The positive feedback (due to pheromonic trails) allows the emergence of global *order* without global *coordination*. c) No direct agent-to-agent communication is needed, creating a threefold benefit in: *simplicity* (no languages, messages, awareness of partner agents, etc.), *robustness* (agents are not coupled, computation is off-loaded, and the negative feedback provides "forgetting" the fruitless paths), and *protection* (without explicitly conveyed information, confidentiality is preserved: a paramount asset for military applications).

Considering the objective and the rationale, the approach is based on the following *premises* and *criteria*:

a) To be affordable on usual configurations, the "artificial ant colony" has to be restricted to a reduced number of agents[2].
b) To be workable, as well as to allow assessment, the undertaking has to avoid starting from scratch (a "tabula rasa" stance impairs any genuine evaluation).
c) Corollary of a) and b): the problem-class chosen has to be of manageable complexity (to prevent failure under combinatorial explosion); moreover, it should be a familiar "workhorse".
d) Practicality entails that the trail-building behaviour of ants and their random movement shall be taken only as initial *model* not as inexorable *dogma*. Explicitly, agents, artificial pheromones, and their discrete environment must not necessarily simulate an ant society; in contrast, they shall be a compliant problem-solving tool.
e) Corollary of d): after "squeezing" the standard algorithms by fine-tuning their parameters (in Section 4), the symbolic paradigm is brought in to amend the pattern (in Section 5).

Thus, all tests apply to the (easy to understand but difficult to solve) *Travelling Salesperson Problem* (TSP): [4] [6] [7] [8] [18].

---

[1] *Cataglyphis fortis*. In contrast to most other ant species, they do not use pheromones to mark their path; instead, they navigate by path integration and by visual landmarks.

[2] For a 1.5 GHz processor 200 ants need about 1 second to travel through 200 places.

4. Searching for Stigmergy: Local Enhancements

To increase effectiveness it has to be assessed. Therefore, three choices are needed:

a) *Performance Metrics*. Two dimensions are chosen: *optimality* (*O*, expressing the closeness to the optimal solution, measured by the ratio between the lengths of the shortest and the current path), and *simplicity* (*S*, the opposite of complexity, measured by the ratio between the speeds of the fastest and the current algorithm). To start with, the effectiveness measure, *E* is defined as:

$$E = w_1 * O + w_2 * S$$

where the w_i are (still) empirically chosen weighs.

b) *Reference Algorithms*. Because of the metrics, the two extreme algorithms must be conventional ones (i.e. not based on SC), used in operational research. Such algorithms could be the *Best-First Algorithm* (*BFA*) (starting from every town and choosing the nearest town to be the next to be visited) and an *Exhaustive Algorithm* (*XA*) that determines all solutions. The *BFA* is very fast but the solutions provided are quite far from the optimum, while the *XA* gives always the best solution, but with a very low speed.

c) *Stem Algorithm.* In the role of a stem cell, this algorithm will spawn all the variants taken into account. Hence, it has to be one already proven effective – at least according to the proposed metrics.

The variant chosen for testing is *EAS* that differs from the standard *AS* on four points: transition rule, pheromone-trail update rule, local updates of pheromone trails and the use of a candidate list, yielding an increase in exploration and a better performance on large problems. Having many parameters, this algorithm is very flexible and can be fine-tuned for specific map configurations. Some of the parameters are: number of ants used to solve the problem (*m*); the influence (α) of the pheromone intensity in the transition rule; the distance between two towns (*d*); the influence (β) of that distance in the transition rule; evaporation speed of the pheromone (ρ); pheromone intensity for elitist ants (*e*); the amount of pheromone that an ant deposits (*q*); the pheromone intensity between two towns (τ).

From those variables, the most sensitive proved to be: α, β, ρ and *e*. Usually a trade-off between α and β has to be found in order to achieve the best quality solution and a short response time. The importance of the evaporation speed of the pheromone (ρ) can be explained as follows: a high ρ value could trigger the need to re-explore the map, while a low ρ value could lead to the saturation of the paths, creating a general confusion in choosing the best way. Using elitist ants, the convergence is faster.

Noteworthy, even if this algorithm shifts away from the "pure" SC model, through such parameters as β, *e* and the vector with the towns that an ant has visited (this one imposed by TSP itself), it is still rather close to the biological model.

Applying the metrics is straightforward: after finding out for which parameters the stem algorithm is most sensitive to, these are modified "micro-continually" [2] and the path length and speed of each variant are introduced in the expression of *E*. The results are commented upon in Section 6 (to be compared to those obtained according to the approach taken in the next section).

5. Searching for Synergy: Outside the Paradigm

This section (first of all, its title) requires explanation: Why searching for synergy, when SC is intrinsically based on synergy? Why mixing further paradigms, when the algorithms applied before, already use artificial ants differing from the natural ones (e.g., elitist ants)?

Here the approach is very pragmatic, disregarding theoretical problems (even basic concepts as SC or synergy tend to be dealt with as labels). This involves a strategic shift: whereas Section 4 focused on using a (tailored) biological model for the problem at hand, Section 5 aims at creating a problem-solving method. Put bluntly: ants do not care about TSP while salespersons have no reason to smell pheromones. Thus, since the technological *constraints* (mainly, processing power and problem complexity) are unavoidable, the only way out is to exploit the technological *freedoms*. These are based on symbolic processing, are kept at minimum (to carry on the advantages of SC), and are grouped in three classes:

a) *Boosting the agents*. Artificial ants may be smarter than their natural counterparts, tending to become closer to agents (e.g., for other kinds of problems, they gain full autonomy being implemented as threads; however, they will not be genuine agents, since they still will not communicate directly.) Indeed, the natural SC is a process of self-organization that does not imply structural changes, so the future agent behaviour will not be different (ants do not change their transition rule – the criteria to choose the next town to be visited – from one place to another or from one tour to the next). However, it is obvious that if an entity changes over time its behaviour will improve. Thus, analysing different types of maps, it came out that choosing the closest town as the next one to be visited, will not always lead to a better solution – in some cases, this was even worse – whereas ants can choose sometimes less promising paths. Hence, the agents used to resolve the TSP are able to memorize the towns they have passed through, can be aware of the best tour found, and can make decisions accordingly.

b) *Adapting the environment*. Albeit having to return to their nest (the departure town), artificial ants are not obliged to mark two kinds of paths (towards food sources and towards the nest). Hence, the pheromones can convey other kind of information; likewise, the positive/negative feedback can be adjusted to the current state (e.g., when a path is clearly unpromising, evaporation can be instantaneous). Thus, artificial ants can mark their paths with two kinds of pheromone: their own individual pheromone and the colony's pheromone, trying to avoid the paths that have already been marked with their own pheromone at the previous tour(s). The result is a higher map exploration. Another way to achieve a better exploration is to "eat away" pheromone on the edges just crossed (down to a minimum value) like in the *Ant Colony System* [5].

c) *Instituting (limited) central coordination*. In this respect, the distance to the biological model can be large and will probably increase. Indeed, the MAS itself may have a vital role in guiding the search. In *EAS* the MAS reveals itself through the graph map (sometimes also including non-Euclidian elements for expressing the distance) and the rewarding for finding better paths.

To speed up the investigation, the current interface allows external control of the problem-solving process, using "user-driven heuristics" [3]. That means to give the user the power to guide dynamically the search process (i.e. to modify any parameter), after assessing the existing partial results. This method proved to be very effective despite an important drawback: because of the user interference, the speed of the variant cannot be measured, and as a result, its effectiveness is evaluated subjectively. Fortunately, this problem can be fixed in two steps: a) translating the user actions into an interface agent; b) converting this agent into a monitoring one (an adequate way to set up flexible central coordination).

6. Implementation and Results

Although for TSP the ants are too simple to require multithreading, to allow upgrade (i.e. more processing power for individual agents), the operating system is Windows 2000 (because its pre-emptive scheduling, kernel mode, and versatile application programming interface, allow affordable multithreading). The reason: since agents are intrinsically reactive and "run-time beings", they must be autonomous, lasting, and context-sensitive – i.e. interrupt-driven. All tests used the same map (benchmark *Eil45* [16] with 45 towns and an optimal path of 3,006.393 (Euclidian distance) units. Because of ease, the program was written at first in *C++Builder5* but, to gain speed (via dynamic tables), another version was written in *Delphi6*. From the about 45 algorithm instances tested out, 22 gave relevant results. The most important are presented below.

Varying the parameters (Section 4). Parameters α and β have a vital role in the ant decision-making. If either α or β equals zero, the ants will take into account only the distance between the towns ($\alpha = 0$) or the intensity of the pheromone trails ($\beta = 0$), leading thus to solutions far from optimum. Since these parameters are so important, they have to be carefully chosen. The optimal values found for the parameters are $\alpha = 1$, β between 5 and 10 covering a vast diversity of map configurations. The simulation results are grouped in Figure 1. In the same manner, several simulations have been carried out in order to determine the other parameters, the optimal values being: $e = 5$, $\rho = 5$, $q = 100$, and the number of ants (m) equal to the number of towns. After fine-tuning the algorithm, simulations were carried out and the values for speed (or simplicity, S) and optimality (O) determined using the *Eil45* benchmark (with an optimum tour of 3,006.393 units).

Since the speeds of different algorithms also depend on implementation, the S parameter must be determined using the algorithm's complexity formula. The complexities of the examined algorithms are *BFA* (n^2), *XA* ($n!$) and *EAS* ($m * n * i$, where m is the number of ants, n is the number of towns, and i is the number of iterations). Here $n = 45$, $m = 45$ and $i = 23$ (because the best solution was found after completing 23 iterations). Thus, $O = 3006.393 / 3006.393 = 1$ and $S = 45 * 45 / 45*45* 23 = 0.04347826$.

For *XA*, $O = 1$ (obviously) and $S = 8.98052677\text{e-}54$ (very close to zero). As O values are equal in both algorithms (*XA* and *EAS*) by comparing the S values it is evident (and not surprising) that *EAS* is better.

Nonetheless, because *EAS* does not find always the best solution, it was compared also to *BFA*. For *BFA* $O = 0.84795737$ and $S = 1$. Choosing $w_1 = 0.9$ (for optimality) and $w_2 = 0.1$ (for speed), $E = 0.863161633$ ($0.9 * 0.84795737 + 0.1 * 1$), which compared to the E value of *EAS* (0.904347826) is smaller. The reason for choosing these weighs: optimality is much more important than speed (indeed, TSP is an optimisation problem).

Symbolic processing (Section 5). Because of the intrinsic heuristic nature of this search branch, it is much more difficult to do it systematically.

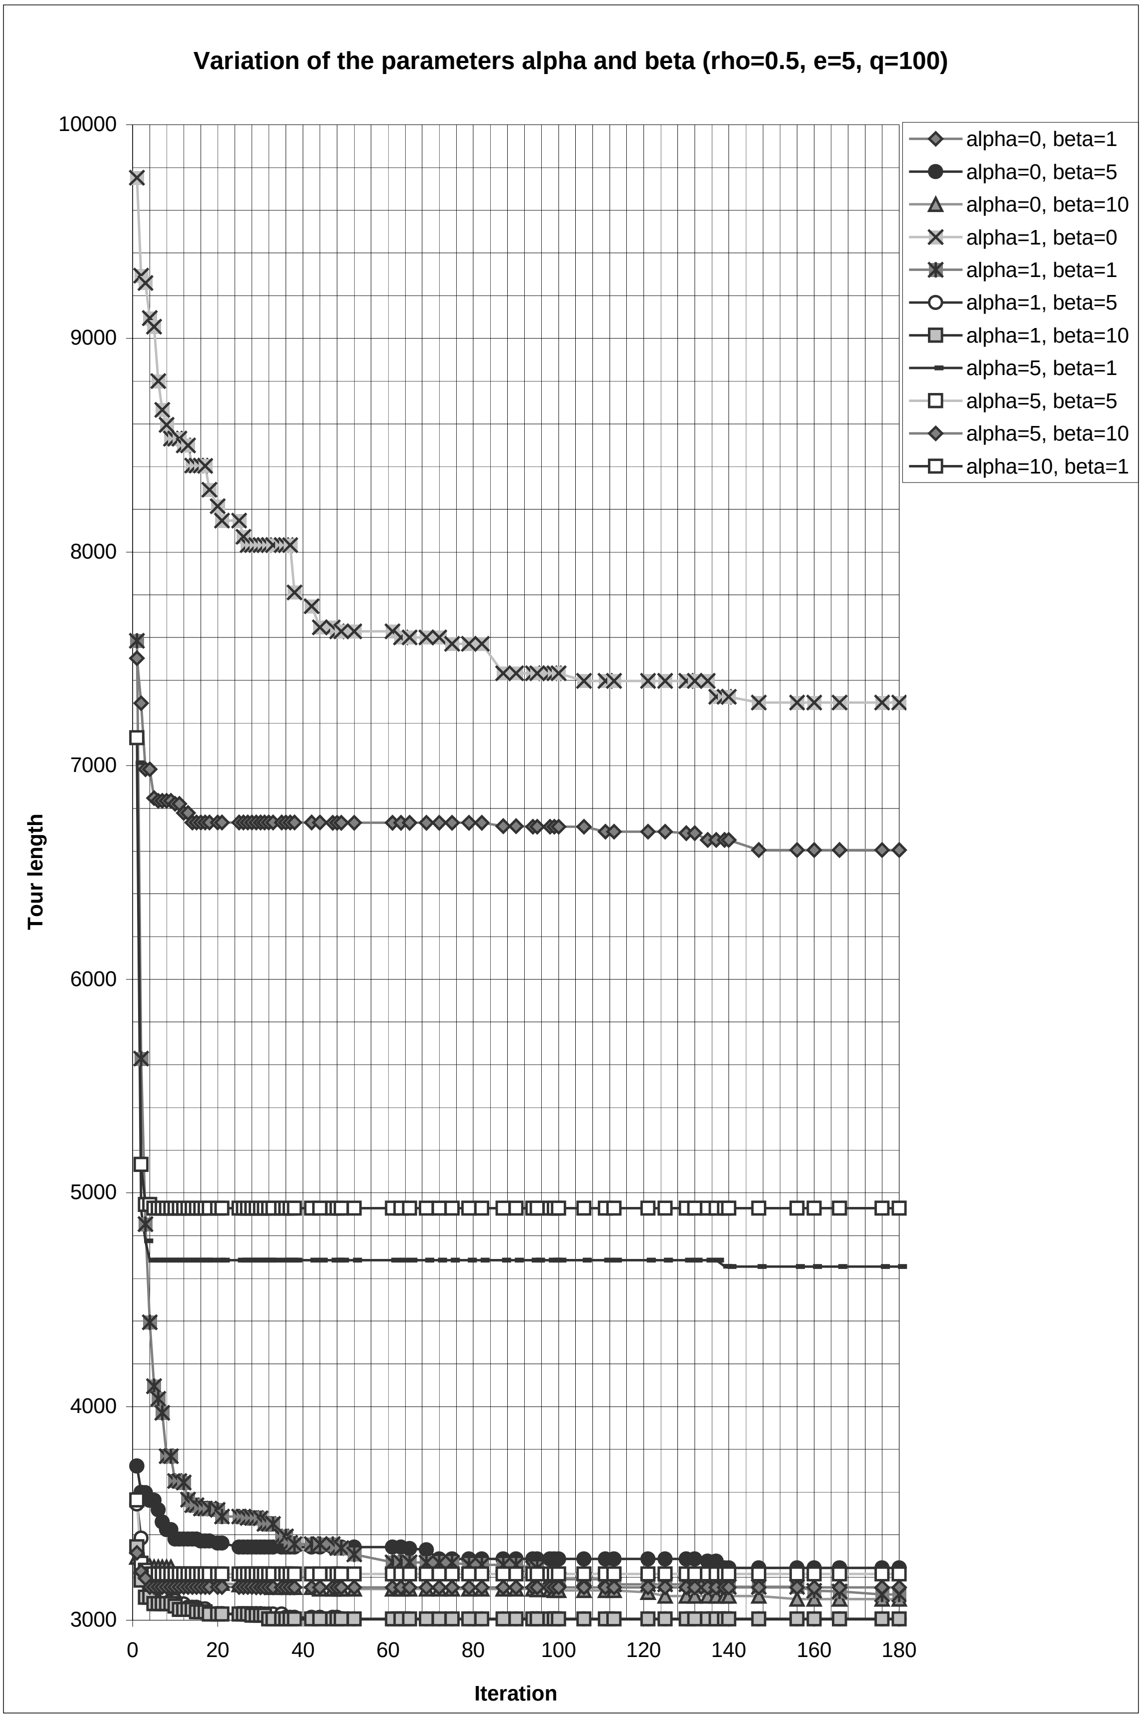

Figure 1: The results achieved by varying alpha and beta parameters

However, the best sequence seems to be "*bca*", i.e. to start with changing the environment (it is rather simple and the problem itself implies obvious differences compared to the ant environment). Boosting the ants is the last resort because: 1) At least in the simulation stage, each addition has to be multiplied by the number of ants before considering its worth. 2) It is sound to use the same ants for different problems. 3) Using complex agents would impair one of the key aims (achieving synergy using a *very* great number of *very* simple entities). Thus, avoiding boosting the agents, the results are:

- *Adapting the environment*. Some new rules imposed to *EAS* reads as follows: a) If the ants do not find a better solution after N iterations (the search stagnates), then the pheromone evaporation (ρ) will constantly be increased until a better solution will be found in a limited number of iterations (only then ρ will be reset to its initial value). b) If the ants do not find a better solution after N iterations the evaporation speed of the pheromone trails (ρ) will constantly be increased and the value of the parameter α will be constantly be decreased. When the ants will find a better solution, the parameters will be reset to their initial values. (Of course, other kind of rules could be followed, this proving to be a good research direction.)

- *Instituting (limited) central coordination*. Here two main directions can be followed: an auto-coordination from the MAS itself and user-driven heuristics. Since at this stage, because of user-driven heuristics (see Section 5), irrefutable results about auto-coordination are still missing, they will be presented in a future paper.

7. Conclusions and intentions

Since the paper treated the *potential of a paradigm* and one of its *problem-solving instances,* the review involves both A) *general* conclusions and B) *factual* ones.

A1. Despite being yet in a syncretic stage, stigmergic coordination shows an obvious potential as problem-solving tool, at least in the context of simple MAS.

A2. Relevant results can be achieved on affordable configurations, with simple MAS, classical problems and usual benchmarks.

A3. Synergistic effects can be reached deviating from the biological model by adding symbolic processing factors. It seems that the effectiveness of simple MAS is enhanced firstly *adapting the environment* and secondly *instituting some (limited) central coordination*. Expanding the agents is less promising.

B1. To assess very many (sometimes quite similar) algorithm variants, a *performance metrics* is vital (even if it is subjective and debatable as the one proposed here).

B2. In line with the applied metrics, the effectiveness of simple MAS in solving TSP, can be visibly (albeit not very much) enhanced by fine-tuning the *EAS*.

B3. Paradoxically, the most promising improvements seem to be reached where quantitative evaluations are yet missing, i.e. intruding into the way ants would act.

B4. This intrusion is carried out most efficiently (but in a yet not quantifiable way) through external control.

The *short-range intentions* are corollaries: a) improving the evaluation formula E; b) finishing the tests regarding the control of pheromones and search paths; c) speeding up the search by extending the "user-driven heuristics" d) developing the interface agent; e) converting the external control into a monitoring agent.

The *middle-range target* is to try stigmergic coordination in real-world problems, starting with the manufacturing control in a small consumer goods factory.

Acknowledgment

The authors thank C.B. Zamfirescu from the Catholic University Leuven for his valuable help – ranging from bibliography to ideas.

References

[1] Aleksander, I. Emergence, Evolution and Depiction: The Next 50 years of Artificial Intelligence. Keynote Speech, *International NAISO Congress on Information Science Innovations ISI'2001*, Dubai, 2001.

[2] Bărbat, B.E. Agent-Oriented Captology for Anthropocentric Systems. *Large Scale Systems: Theory and Applications 2001* (F.G. Filip, I. Dumitrache, S.S. Iliescu, Eds.), Elsevier, IFAC Publications, 214-219, 2001.

[3] Bărbat, B.E. *Agent-Oriented Intelligent Systems*. Romanian Academy Publishing House, Bucharest, 2002. (In Romanian.)

[4] Bonabeau, E., Dorigo, M., Theraulaz, G. *Swarm Intelligence: From Natural to Artificial Systems*, Oxford University Press, New York, 1999.

[5] Brueckner, S. *Return from the ant – Synthetic Ecosystems for Manufacturing Control*. Dissertation at the Humboldt University, Berlin. Published by the Humboldt University's Project Digital Dissertation.

[6] Dorigo, M., Maniezzo, V., Colorni, A. The Ant System: Optimisation by a Colony of Cooperating Agents. *IEEE Transactions on Systems, Man, and Cybernetics,* Part B, **26** (1): 29-42, 1996.

[7] Eyckelhof, C.J., Marko, S. Ant Systems for a Dynamic TSP: Ants Caught in a Traffic Jam. *ACO 2002 Conference***.**

[8] Gambardella, L.M., Dorigo, M. Ant-Q: A Reinforcement Learning Approach to the Travelling Salesman Problem. *Proceedings of the Eleventh International Conference on Machine Learning*, Morgan Kaufmann, 1995.

[9] Grasse, P.P. La reconstruction du nid et les coordinations interindividuelles chez bellicositermes natalensis et cubitermes sp. La theorie de la stigmergie: essai d'interpretation du comportement des termites constructeurs. *Insectes Sociaux* **6**, 41-81, 1959.

[10]Negulescu, S.C. *Applying Stigmergic Coordination in Optimisation*. Diploma work. "Lucian Blaga" University, Sibiu, 2003. (In Romanian.)

[11]NERO – Division of Neuroinformatics, Department of Computer Science, University of Bonn – *http://habakuk.nero.uni-bonn.de/*

[12]Parunak, H.V.D., Brueckner, S. Ant-like missionaries and cannibals: synthetic pheromones for distributed motion control. *Proceedings of the 4th International Conference on Autonomous agents*, 467-474, Barcelona, ACM Press, 2000.

[13]Pentland, A. Perceptual intelligence. *Comm. ACM*, **43** (3), 35-44, 2000.

[14]Pfeifer, R. Embodied Artificial Intelligence: Dynamics, morphology and materials in the emergence of cognition. (Invited paper.) *Proceedings of the International NAISO Congress on Information Science Innovations ISI'2001* (M.F. Sebaaly, Ed.), 38-39, ICSC Academic Press Canada/The Netherlands, 2001.

[15]Roumeliotis, S.I., Pirjanian, P., Matarić M.J. Ant-inspired navigation in unknown environments. *Proceedings of the 4th International Conference on Autonomous agents*, 25-26, Barcelona, ACM Press, 2000.

[16]Ruprecht-Karls-Universität Heidelberg – *http://www.iwr.uni-heidelberg.de/groups/comopt/software/tsplib95/*

[17]Stützle, T., Dorigo, M. ACO Algorithms for the Quadratic Assignment Problem. *New Ideas in Optimisation* (D. Corne, M. Dorigo, F. Glover, Eds.) McGraw-Hill, London, UK, 1999.

[18]Stützle, T., Hoos, H.H. MAX-MIN Ant System. *Future Generation Computer Systems*, 2000.

[19]Theralauz, G., Bonabeau, E., Deneubourg, J.L. The origin of nest complexity in social insects. *Complexity,* **3** (6), 1998.

[20]Valckenaers, P., Kollingbaum, M., Van Brussel, H., Bochmann, O., Zamfirescu, C. The design of multi-agent coordination and control systems using stigmergy. *Proceedings of the 3rd International Workshop on Emergent Synthesis*, (P. Butala, K. Ueda, Eds.), Bled, 2001.

[21]Vaughan, R.T., Støy, K., Sukhatme, G.S., Matarić, M.J. Whistling in the Dark: Cooperative Trail Following in Uncertain Localization. *Proc. 4th Int. Conf. on Autonomous agents*, 187-194, Barcelona, ACM Press, 2000.

[22]Wehner, R. Ant navigation: mini brains – mega tasks – smart solutions. (Invited paper.) *Proceedings of the International NAISO Congress on Information Science Innovations ISI'2001*, Dubai, 2001.

Sorin C. NEGULESCU. Graduated in Computer Science in 2003 ("Lucian Blaga" University, Sibiu). Master student. He is now modelling (for stigmergic coordination) the manufacturing process of a shoe factory.

Boldur E. BĂRBAT. M.Sc. in Electronic Engineering, postgraduate specialising in Programming, Ph.D. in Digital Computers ("Politehnica" University Bucharest), full professor ("Lucian Blaga" University, Sibiu). Author/co-author of nine books, 11 textbooks, over 60 papers/articles (IPCs: eight), 30 RD technical reports. Main interests: agent-orientation and real-time systems.