# User-Driven Heuristics for Nondeterministic Problems

**Sorin C. Negulescu*, Constantin B. Zamfirescu **, 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: +40-269-216062; Fax: +40-269-212716; E-mail: zbc@acm.org

*** Phone: +40-269-212523; Fax: +40-269-212716; E-mail: bbarbat@rdslink.ro

*Abstract: Multi-agent systems are winning more and more battles in the artificial intelligence field. A feasible and successful approach proved to relay on small and simple agents inspired from biological models. This approach emphasises the significance of the environment in the agent system, the attention that must be paid to the dynamics that emerges from interactions and the importance of growing rather then building agents. Finding out the form and parameters that influence the system behaviour is non-trivial for humans, even more difficult than knowledge engineering. The paper proposes and investigates the methods that are enabling user-driven solutions in dealing with this challenge. The methods are based on two essential functional requirements: a) at macro level, the need to monitor and represent in an intuitive way the system behaviour and b) at micro level, the need to control and track the system state space. The approach proved to be workable on usual configurations and effective in dealing with combinatorial explosion.*

*Keywords: Stigmergic Coordination, Multi-Agent Systems, Synergy, Symbolic Paradigm, Operations Research, Travelling Salesperson Problem.*

## 1 Introduction

Multi-agent systems (MAS) are winning more and more battles in the artificial intelligence (AI) field due to their ability in solving difficult problems. Some approaches are based on complex agents that are trying to emulate human intelligence and reason explicitly about their coordination. Other approaches draw on small and simple agents inspired from biological models where the far-most known multi-agent system seems to be the ant-system (AS). This approach emphasises the significance of the environment in the agent system, the attention that must be paid to the dynamics that emerges from interactions and the importance of growing rather then building agents. However, such an approach reintroduces many of the problems of complex system design. As long as there is no known algorithm to derive a specification of individual behaviour from overall system behaviour, a fundamental challenge in AS is to understand the relation between individual agent behaviours and the behaviour of the system as a whole. This is yet a challenge: as stated in [12], "much current work on constructing systems of this sort is more art than science" because "they often exhibit regularities (such as exponential convergence) that we do not understand, and we do not know how to improve their functioning in a disciplined manner". A most appropriate such manner is to study distributed decentralized systems (first of all, their convergence) via a simple statistical model of local behaviour [12].

On the other hand, Prigogine's idea that the most interesting scientific activities seem to occur at domain interfaces is a confirmed – and, because of financial reasons, often (mainly, now, in Romanian universities) the only affordable – path for applied research. Thus, combining the fields of Manufacturing Systems and Artificial Intelligence is already a customary undertaking; the outcome: Intelligent Manufacturing Systems. Still, on this path, rather incremental advancements are likely.

However, trying a "second degree" reading of Prigogine's idea, a more promising path is to prefer newer paradigms – for instance, stigmergic coordination (SC) – since they are closer to a syncretic stage, able to provide significant progress. Moreover, inter-paradigmatic synergy – e.g., by integrating stigmergic control with symbolic coordination – can be reached more easily.

Since local behaviour cannot be predicted analytically, it must be detected in an operational system. Thus, control elements must incorporate hooks for monitoring the system performance in real time. As a result, the paper proposes a new approach for a reduced class of problems that do not demand full autonomy and are able to accept user-driven solutions in real time. The aim is to faster AS deployment in real-life applications.

In Section 2, the paper introduces the rationale, related work and its history. Section 3 presents the SC model together with its strengths and weaknesses. Section 4 investigates and proposes two novel approaches that proved to be efficient in tuning the overall system behaviour. They are based on two essential functional requirements that we built in the system in order to assess their effectiveness. Firstly, there is the need to monitor and represent the system behaviour in an intuitive way for the designer. Secondly, there is the need to control and track the trajectory of the system by employing user-driven heuristics (UDH) and recording state-space data. The last section draws conclusions and intentions.

## 2 Rationale, Related Work, and History

Since the reasons for choosing SC have been detailed in [4], to impair redundancy, they are skipped over. Here the focus is on user-driven heuristics [3]:

Why humans? The term has to be read in three tones, epitomizing three rationales:

a) Humans – seen as the apex of symbolic reasoning – act as counterparts of sub-symbolic ants, in achieving inter-paradigmatic synergy (to take the best from two quite dissimilar worlds and to get the most from intersecting them).

b) Primarily, when complexity (of all kinds) is high, software effectiveness becomes uncertain and direct human intervention is rather welcomed (in this context, through UDH).

c) It is still helpful to remind that modern IT systems must be anthropocentric – a chief raison d'être of agents. Thus, the users (no matter whether system engineer, application developer, manager, and so on) should be allowed to monitor the system and to communicate with, using familiar semantics – i.e., symbol-based languages, humans are accustomed to.

From an engineering point of view, even the fundamental drawback of any sub-symbolic paradigm, i.e. the trouble to understand what is actually occurring in the lower system levels, is less disturbing than when using paradigms that are more familiar (as artificial neural networks or evolutionary algorithms) since stigmergic behaviour is easier to follow due to its simplicity – at least in the case of common ants. (Desert ants – cataglyphis fortis – show a much more complex behaviour: they do not use pheromones to mark their path; instead, they navigate by path integration and by visual landmarks [21] [17] [18].)

Considering the objective (adapting the Ant System for a class of problems that are able to accept user-driven solutions) and the rationale (the reasons above), the approach is based on the following premises and criteria [20] [4]:

a) To be reasonable on usual configurations, the “artificial ant colony” has to be restricted to a limited number of agents (for a 1.5 GHz processor, 200 ants needed about 1 second to travel through 200 places [19]).

b) To allow assessment (a genuine comparative evaluation), the undertaking has to avoid starting from scratch.

c) The problem-class chosen has to be of manageable complexity (to prevent combinatorial explosion) [15]; moreover, it should be a familiar “workhorse”.

d) Beyond the paper’s aim, practicality itself entails that the behaviour of ants shall be taken only as initial model: agents, pheromones, environment, behaviour, and so on, must not simulate an ant society; in contrast, they shall become a compliant problem-solving tool.

e) The symbolic paradigm has to be directly brought in several variants, to amend the pattern.

All tests apply to Ant System TSP [16] [14] [7] [13] [22] and have been carried out within the same environment as in [19] [20] to be able to draw conclusions.

**History**. The paper stems from two threads of history:

*Synergy*. As regards inter-paradigmatic synergy [1], suggests a triangular synergistic approach, considering the way multi-agent systems and holonic systems have to act in real-time (i.e., should be based on threads).

*Stigmergy*. In [20] less quantifiable synergistic effects were achieved deviating from the biological model applied in EAS by adding symbolic processing components (firstly adapting the environment and secondly instituting limited central coordination). On the other hand, the construction of a symbolic navigation map was treated in [10] for multi-objective problems.

The confluence begun in [8] [2] (on a more abstract level) and got shape for SC in [3] and [4]. In [5] the relationship between stigmergy and synergy was explored based on the threshold principle.

# 3 The Adapted Model

In order to assure autonomy to the paper, the model description is adapted from [4].

The model [11] [16] [6] [13] [22] [9] is in essence 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 syntagm "stigmergic coordination" related to "ant-based systems" (ABS), it "describes a form of asynchronous interaction and information exchange between agents mediated by an "active" environment" [6], or "the production of certain behaviour in agents as a consequence of the effects produced in the local environment by previous behaviour" [9] (cited in [17]). 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" [6].

From the point of view of operations research, the second major advantage of SC, mentioned above, is not vital, since a kind of global coordination is needed anyhow, and the way to achieve it is less relevant. The other two strengths remain crucial; moreover, there are specific features of the ant behaviour that are similar to several requirements of this sub-field (or easily adapted to). The most relevant are:

a) The usually huge number of iterations needed to reach the solution.

b) The likeness between the environment of the ant colony and a graph.

c) The straightforwardness of using pheromones as weights (enhanced by the simple ways to adjust their intensity in order to control dynamic environments).

d) The pheromonic positive reaction used to support promising search paths as well as the negative reaction used to dissuade unpromising ones.

The price to pay for those strengths can be high, because of the unavoidable differences between artificial ants and their natural counterparts. For operational research applications, such significant discrepancy are memory requirement (e.g., to retain partial solutions), communication with the "system", world and time (discrete vs. continue), and a minimal look-ahead ability.

# 4 Implementation Example

The chosen variant for solving TSP was EAS [14] because it was the first relevant attempt to add symbolic-processing components to ant-inspired models. The investigation in this section branches in two directions, both taking advantage of the UDH. The indirect UDH is the heuristics driven by modifying the parameters, and the direct UDH is the one obtained by the express modification of the environment (i.e. the pheromone intensity).

## 4.1 Indirect User-Driven Heuristics

Having many parameters, this algorithm is complex enough and it can be efficiently fine-tuned for specific map configurations at design-time, but especially through UDH at run-time as illustrated in Figure 1.

Considering the double task of the model, these parameters are used in a bi-semantic manner. Thus, the original EAS parameters have the following general meanings:

- $m$ – number of ant-like entities used to solve the problem;
- $d$ – the generalized distance expressing the weight of a graph edge (e.g. the distance between two towns for TSP or the duration of a technological operation for manufacturing control);
- $\tau$ – pheromone intensity on a graph edge;
- $q$ – the amount of pheromone deposited by an ant on a graph edge;

- ρ – evaporation rate of the pheromones (expressing the (in)stability of the environment the system is acting in).

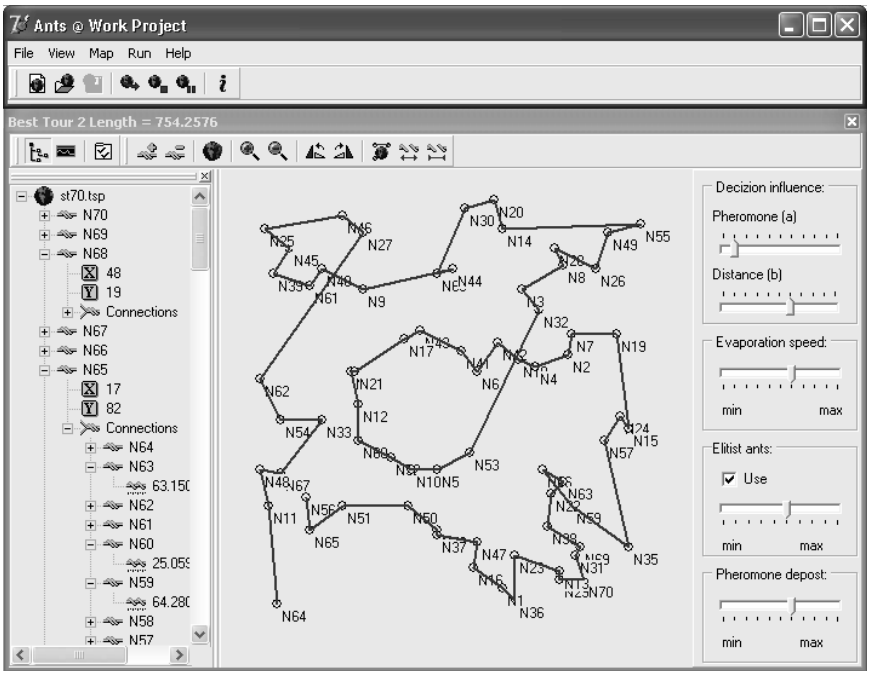


Figure 1

Snapshot of the experimental model, capturing some parameters for UDH

To express the effect of the above parameters, some ancillary variables are useful in choosing the route in the graph: α –the pheromone intensity (exploitation), β – the generalized distance (exploration), *e* – the successful search history.

Although for some of the parameters described above, their influence on the system behaviour could be intuitive, for the most of them, to discover their optimal values, repetitive runs are indispensable. Moreover, in certain dynamic settings finding their optimal values could be unfeasible.

Among these parameters, the most sensitive proved to be m and ρ; as regards the ancillary parameters, their sensitivity related to the sensitivity of ρ is shown in Figure 2.

Usually a trade-off between them has to be found in order to achieve the best quality solution (i.e. near to optimum solution in acceptable time). For real-world problems, this trade-off cannot be made without direct user involvement (e.g. finding in the next 5 minutes a better solution than the current one is more important than finding the best solution in the next 20 minutes).

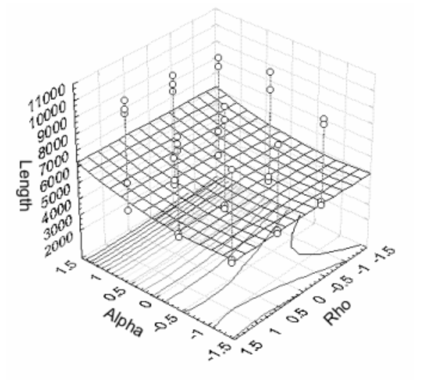


Figure 2a
Solution sensitivity against α and ρ

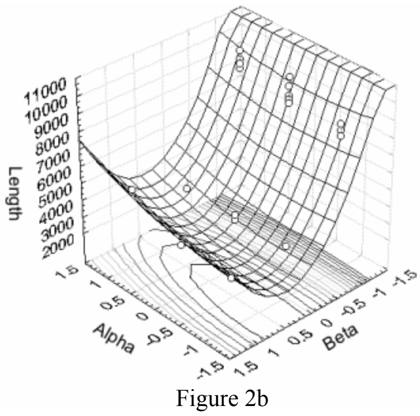


Figure 2b
Solution sensitivity against α and β

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.

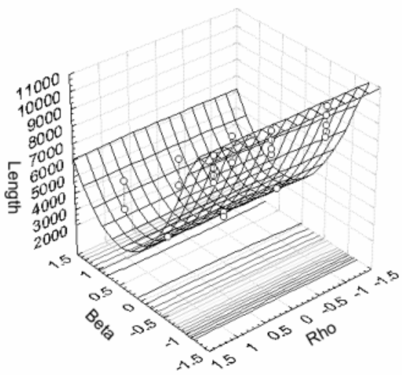


Figure 2c
Solution sensitivity against β and ρ

An example of applying UDH in shown in Figure 3, presenting the system reaction to user-initiated variations of α and β (using the same interface components as in Figure 1). Here, UDH was applied to speed-up convergence. Thus, after 45 iterations expressing a search based on a significant weight of the "look-ahead" ability (distance to the next town) against "natural" ant behaviour (pheromone-driven search), the user remarks that the ants "oscillate" around the near-to-optimum solution. Hence, the search strategy is modified towards a more "ant-like" search, i.e. boosting the weight of pheromone traces and reducing that of map-awareness.

As a result, the search converges better (successive iterations are closest to each other and to the final solution). Thus, in this punctual situation, a bit of "more sub-symbolic processing" proved to be adequate. Of course, in other cases, "more symbolic processing" is better (either to sped-up convergence, or to find a closer-to-optimum solution).

Another kind of approach is the dynamic modification of the parameters through a monitoring thread, but this raises problems that are even more difficult to solve, such as when to modify them and how. To define such rules, they must be firstly extracted from the user experience.

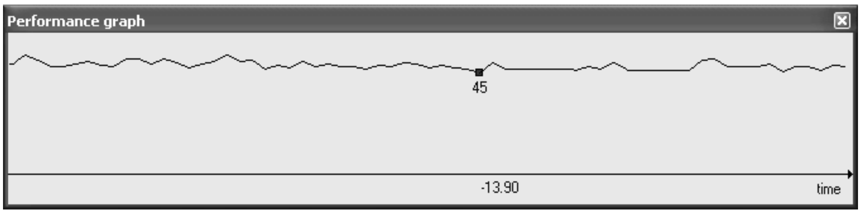


Figure 3
Exploration versus exploitation through UDH

## 4.2 Direct User-Driven Heuristics

Even though the improvement obtained by dynamically modifying the parameters guides the search fine enough, the user must also have the possibility to guide the search in a more explicit manner. Looking for a global-optimum solution, the ants may be deficient in resolving simple local problems produced by the routing between a small number of towns. This kind of problems is illustrated in Figure 4.

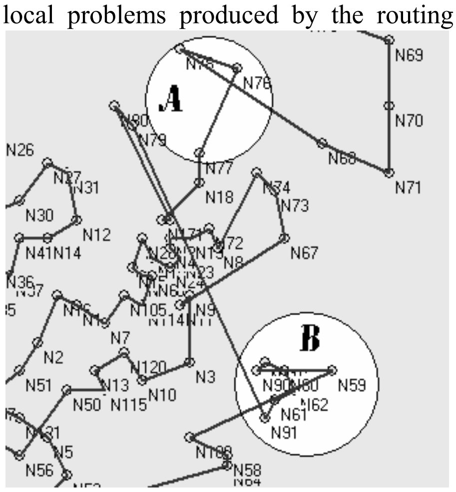


Figure 4
Local (Sub)-Problems of the Global Problem

The “A” zone (formed by towns N77, N76, N78 and N68) and the “B” zone (formed by towns N91, N61, N62, N60, N116, N90, N59 and N100) represent local routing problems for the ants. The reasonable question “why do ants fail to resolve these simple problems?” is well founded especially because the ants would solve them quickly if these problems would not be local-problems of a bigger one. Because of the (very) small distance between the towns, comparing to the distance to other towns in the map, the ants are attracted by choosing the shorter route between them. This leads also to another problem because on the shorter routes the ants deposit a larger amount of pheromone and make these routes even more attractive. This kind of problems can not be easily solved neither by decreasing the amount of pheromones that an ant deposit, nor by “instructing” the ants not to take so much in consideration the distance between towns.

Although the problems described above are hard for the ants to overcome, they are not difficult at all for the user who is guiding the search. For example, the solution for the problem within the “A” zone is the route that goes through the towns N77, N78, N76 and N68 in one direction or N68, N76, N78 and N77 for the other direction, both having the same lengths and both being valid (because the TSP problem described above is symmetrical). Almost in the same manner, the user can easily give the solution for the problem within the “B” zone: N90, N116, N60, N59, N62, N61 and N91, or for the other direction.

The easiest way to implement this kind of UDH is by giving the user the possibility to modify the intensity of the pheromones on any path between towns. For example, in order to solve the problem within the “A” zone, the user must reduce the intensity of the pheromones on the paths that links the towns N77 - N76 and N78 - N68, and to increase the intensity of the pheromones on the paths that links the towns N77 - N78 and N76 - N68.

The comparison between the standard and the direct-UDH is presented in Figure 5. After a number of iterations, the user observes that the ants are blocked in a local-optimum solution and decides to modify the environment to speed-up the search. The lower line in the graphic represents the solution quality after the user intervention.

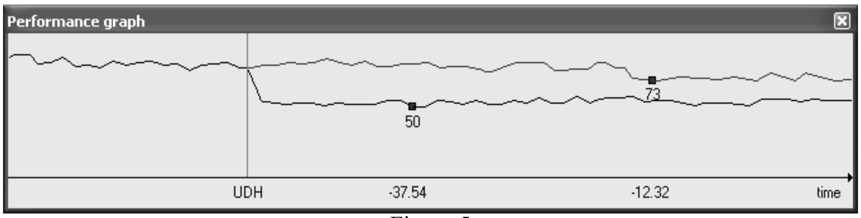


Figure 5

Comparison between direct and indirect-UDH approach

It can be clearly seen that without the direct-UDH, the convergence to the optimum solution would have taken much more time (as represented by the upper line in the graph).

It is also true that ants may overcome this kind of problems by themselves, because of their “free will” (the choice of the next town being a probabilistic one, e.g. an ant can choose to go from the current town to another that is very far, using a path where the intensity of the pheromone is very low), but this can take much more time than a “divine” user-intervention.

**Conclusions and Intentions**

As mentioned in the introduction of the paper this kind of approach suits for a reduced class of problems that do not demand full autonomy and are able to accept user-driven solutions, time being here the key-factor that decide this.

The approach proved to be workable on usual configurations (tested with usual benchmarks on classical problems) and effective in dealing with combinatorial explosion (e.g. the nondeterministic problems) – confirming the conclusions arrived at for TSP in [20].

In comparison with the results from [20], the progress was achieved mainly extending the central coordination, via the "user-driven heuristics" concept, particularly by the direct UDH, where, depending on the user skills the time to reach the best solution can be extremely reduced.

The next step is to implement a monitor thread or an agent that will replace the user: "Agent-Driven Heuristics".

**References**

[1] B. E. Bărbat: Holons, Agents, and Threads in Anthropocentric Systems, Studies in Informatics and Control Journal 9, 3, 2000, pp. 253-268

[2] B. E. Bărbat: Agent-Oriented Intelligent Systems, Romanian Academy Publishing House, Bucharest, 2002 (in Romanian)

[3] B. E. Bărbat, C. B. Zamfirescu and S. C. Negulescu: The best from ants and humans: synergy in agent-based systems, Studies in Informatics and Control Journal, 13, 1, 2004, pp. 47-59

[4] B. E. Bărbat, S. C. Negulescu and C. B. Zamfirescu: Boosting Synergy in Agent-Based Optimisation Tools, Transactions on AUTOMATIC CONTROL and COMPUTER SCIENCE, 49 (63), 4, 2004, pp. 119-124

[5] B. E. Bărbat, S. C. Negulescu, C. B. Zamfirescu: Human-Driven Stigmergic Control. Moving the Threshold, (Invited paper at the 17th World Congress IMACS, Paris, 2005)

[6] E. Bonabeau, M. Dorigo and G. Theraulaz: Swarm intelligence: from natural to artificial systems, Oxford University Press, New York, 1999

[7] C. J. Eyckelhof and S. Marko: Ant systems for a dynamic TSP: ants caught in a traffic jam, ACO 2002 Conference, 2002

[8] F. G. Filip and B. E. Bărbat: Industrial informatics. New paradigms and applications, Technical Publishing House, Bucharest, 1999 (in Romanian)

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

[10] Hadeli, P. Valckenaers, C. B. Zamfirescu, H. Van Brussel, and B. S. Germain: Self-organising in Multi-agent Coordination and Control using Stigmergy, in Self-Organising Applications, Springer Verlag, Lecture Notes in Artificial Intelligence 2977, 2004

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

[12] H. Van Dyke Parunak, S. Brueckner, John A. Sauter, R. Matthews: Global Convergence of Local Agent Behavior, Submitted to The Fourth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'05), 2005. Available at http:// www.erim.org /~vparunak /AAMAS05Converge.pdf

[13] L. M. Gambardella and M. Dorigo: Ant-Q: a reinforcement learning approach to the travelling salesman problem, Proceedings of the Eleventh International Conference on Machine Learning, Morgan Kaufmann, 1995

[14] M. Dorigo, V. Maniezzo and A. Colorni: The ant system: optimisation by a colony of cooperating agents, IEEE Transactions on Systems, Man, and Cybernetics, Part B, 26 (1), 1996, pp. 29 – 42

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

[16] P. Valckenaers, M. Kollingbaum, H. Van Brussel, O. Bochmann, and C. B. Zamfirescu: 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

[17] R. T. Vaughan, K. Støy, G. S. Sukhatme and M. J. Matarić: Whistling in the Dark: cooperative trail following in uncertain localization, Proceedings of the 4th International Conference on Autonomous Agents, Barcelona, ACM Press, 2000, pp. 187 – 194

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

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

[20] S. C. Negulescu and B. E. Bărbat: Enhancing the effectiveness of simple multi-agent systems through stigmergic coordination, Fourth International ICSC Symposium on ENGINEERING OF INTELLIGENT SYSTEMS (EIS 2004), University of Madeira, 2004

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

[22] T. Stützle and H. H. Hoos: MAX-MIN Ant System, Future Generation Computer Systems, 2000