

The Best from Ants and Humans: Synergy in Agent-Based Systems

Boldur E. Bărbat

Constantin B. Zamfirescu

Sorin C. Negulescu

"Lucian Blaga" University of Sibiu

Department of Computer Science and Automatic Control

4 Emil Cioran, 2400 Sibiu

ROMANIA

*The unpleasantness of a statement is hardly
to be considered a proof of its falsehood.*

Henry Buckle "History of Civilization"

Abstract: Albeit agent-orientation became a well-established course in artificial intelligence at the engineering stage its effectiveness is rather poor for affordable agent-based systems. Despite the increasing number of biologically inspired models, the newer paradigms are in a syncretic stage. Thus, although the fidelity towards the biological model is sometimes quite low, inter-paradigmatic synergy is not manifest enough. The paper aims at: a) boosting such synergy at two levels (ant-like entities and symbolic processing); b) validating its path by testing it on a relevant problem in the field of operational research; c) proposing mechanisms with synergistic potential. Specific mechanisms are designed to graft symbolic components onto the sub-symbolic foundation (the filtered biological model), are tailored to manufacturing control, and are tested with usual benchmarks on an experimental model. The paper concludes that combining stigmergic coordination with symbolic processing components has significant synergistic potential. The most useful mechanism proved to be "user-driven heuristics".

Keywords: Stigmergic Coordination (SC), Agent-Based Systems, Synergy, Manufacturing Control (MC), Travelling Salesperson Problem (TSP), User-Driven Heuristics (UDH)

Boldur E. Bărbat has a M.Sc. degree in Electronic Engineering, postgraduate specializing in Programming, Ph.D. in Digital Computers ("Politehnica" University Bucharest). Currently he is full professor at the "Lucian Blaga" University, Sibiu, Department of Computer Science and Automatic Control. Author/co-author of nine books, 11 textbooks, over 60 papers/articles (IPCs: eight), 30 RD technical reports. Main interests: agent-orientation, stigmergic coordination, synergy, anthropocentric interfaces and real-time systems. He is a member of the NoE AgentLink.

Constantin-Bala Zamfirescu received the computer science engineering degree and M.Sc. degree in distributed systems from the "Lucian Blaga" University of Sibiu in 1996 and 1997, respectively. Currently, he is with the Computer Science and Automatic Control Department of the "Lucian Blaga" University of Sibiu. His research interests include multi-agent systems with applications in manufacturing systems, group decision support systems and information retrieval systems. He is a member of the IFACTC on Large Scale Systems.

Sorin C. Negulescu graduated in Computer Science in 2003 ("Lucian Blaga" University, Sibiu). Currently he is enrolled as a master student in Quality Management. His research interests include multi-agent systems with applications in manufacturing systems using stigmergic coordination, factory automation and the holonic paradigm.

1. Introduction

Albeit agent-orientation (AO) is already a well-established course in artificial intelligence (AI) – and even in information technology (IT) as a whole – at the engineering stage its effectiveness is rather deficient, especially for affordable agent-based systems (ABS), no matter what paradigms are applied. On the other hand, despite the increasing number of biologically inspired models [1], the newer paradigms they are founded on are in a yet syncretic stage, embodying a promising (but too little exploited) niche in itself – for both applied research and effective implementation. Inter-paradigmatic synergy is rather not manifest enough, although the fidelity towards the biological model is sometimes quite poor, letting place for components sticking to older paradigms (mainly the symbolic one).

Applying (at a less general level) Prigogine's idea that the most interesting scientific activities seem to occur at domain interfaces, the paper tries to take advantage of the niche between several (quasi-) mono-paradigmatic approaches to reach effectiveness in affordable ABS, aiming at:

- a) Boosting synergy at two levels, selecting a paradigm with intrinsic synergistic potential to start with and combining it with other paradigms (for engineering reasons, this is firstly the symbolic one).
- b) Validating the approach by testing it on a relevant problem in the field of operational research (the familiar (TSP) was chosen to allow easy assessment against related work).
- c) Deepening the approach by investigating and combining mechanisms with synergistic potential (tested easily via UDH [3]).
- d) Extending it to more difficult application fields, as the nowadays very significant domain of MC [5].

Thus, in Section 2 the paper presents its *rationale* and the *related work* – including its own *history*. Section 3 delineates the *approach*. Section 4 asserts the philosophy, investigating the possible synergy sources and claiming that synergy emerges easier in sub-symbolic context. Focusing on the *engineering perspective*, Section 5 filters the biological model of SC, investigates the alternatives for combining it with symbolic processing (depending on the application sub-fields) and explores *mechanisms* to modify the model in line with the architectural requirements. On this groundwork, Section 6 reviews an *experimental model* in the area of MC, following the results of [6] for operational research. The last section draws *conclusions and intentions*.

2. Rationale, Related Work, and History

Since this is an undergoing research the rationale is presented in detail, but related work is abridged to impair redundancy.

2.1. Rationale

The first motive is obvious: any search for synergy is promising for both research and applications. The likelihood to reach synergy increases in the context described in Section 1, because of:

- a) The niche effect;
- b) The intrinsic synergistic potential of the (new) paradigms based on biological models;
- c) Their syncretic stage.

Specific reasons emerge when clarifying the title:

Why ants? At least for six reasons:

- a) Ant colonies perform remarkably, even in very dynamic environments [7].
- b) As component of the system, the ant itself is extremely simple (having a brain of about 0.1 mg.).
- c) The behaviour rules too: the ant travels from the anthill to the food source and back guided only by pheromones (at least in the case of common ants; desert ants (*cataglyphis fortis*) manifest a much more complex behaviour: they do not use pheromones to mark their path, navigating by path integration and by visual landmarks [7] [8] [9]).
- d) Almost any biologically-inspired model was useful to applied artificial intelligence (AI) [1].
- e) The synergistic effect is evident and huge: although ants behave apparently rather as robots than as beings, the colony they belong to is highly functional and effective.
- f) Even if synergy is impressive, the trouble to understand what is in fact going on at system level, is less upsetting than in the case of more familiar sub-symbolic paradigms (as artificial neural networks or evolutionary algorithms) since ant behaviour is easier to follow due to its simplicity.

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 (indeed, the title suggests to take the best from two quite dissimilar worlds to get the most from intersecting them; here the AND should be read almost as the similar boolean operator).
- 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.

Why agent-based systems? Since defending the use of agents is nowadays pointless, just two nuances (both related to affordability) need explanation:

- a) Although many such systems are multi-agent systems (MAS) – or labelled as such – it is not mandatory to exclude applications with a few agents.
- b) Agent-based, not agent-oriented, to cover also very simple agents in small applications.

2.2. Related Work

Here is surveyed only recent work about swarm intelligence applied to manufacturing control and some essential ideas about self-organization in the context of this paper. Related work regarding SC applied to operational research was reviewed newly in [6].

Decentralized MC systems have been already proposed as a feasible way to overcome the limitations that hierarchical and centralized control exhibit in a highly dynamic environment [10] [11]. As in the original "Ant System" [13] [14] most of the work is focusing either on the algorithms that derive an optimal schedule [13], but using predefined plant configurations, or in contrast, on the architectural features that enable the required agility in facing the integration aspects [12]. Other contributions are dealing with the controller capability to manage the production flow [5]. Most of these implementations are MAS, claiming to pose self-organization capabilities to cope with exogenous disturbances inherited in any manufacturing process. The essence of self-organization is that the system attains a spatial, temporal or functional structure without specific interference from outside [15] (i.e., its structure and functionality are not imposed on the system, but the system is importing energy from outside in an undefined manner [16]). The organisation can evolve in time or space, can maintain a stable form or can show transient phenomena [17].

2.3. History

In [18] the *Elitist Ant System* (EAS) was fine-tuned, obtaining for TSP (slightly) better results. Next, in [6], besides other minor enhancements in the same direction, less quantifiable synergistic effects were achieved deviating from the biological model by adding symbolic processing components (firstly adapting the environment and secondly instituting limited central coordination). On the other hand, from an engineering perspective, the construction of a symbolic navigation map is treated in [19] for multi-objective problems. On a wider-ranging level (regarding inter-paradigmatic synergy), [20] suggests a triangular synergistic approach, considering that MAS become a common means to design and implement holonic software systems, and that threads have a significant unused potential. On this base, a similar approach for robotic teams was proposed in the RoboCup context [20] [21], and extended for any anthropocentric systems in [22]. From a different perspective, a generic architectural outline to support developing flexible interfaces for industrial applications, based on synergetic correlation between persuasive technologies and polymorphic agents was described in [23].

3. Approach

Given the critical non-linearity of synergistic processes, "the best from ants" is very dependent on the application sub-field and sometimes even on the problem. To get also "the best from humans", it is necessary to exploit their symbolic reasoning (it is supposed that humans are still the best in this matter). Hence, the premises and criteria are stated having in mind applied SC and MC, while the major features of ant-inspired models are analysed using the double filter of agents and real-time context. (Fortunately, since agents are real-time beings [3] [4], industrial environments do not impose fundamental new restrictions.)

3.1. Terminology

At least three concepts need explanation, since they are used here with specific connotations:

Stigmergy. The concept is used in its initial meaning proposed by Grasse [24], to characterise the type of interaction taking place in insect societies. Stigmergy is the 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 SC related to MAS, it "describes a form of asynchronous interaction and information exchange between agents mediated by an "active" environment" [25], or "the production of certain behaviour in agents as a consequence of the effects produced in the local environment by previous behaviour" [26].

Agents. Since ants are so simple, here suffices the weak – and generally unchallenged – notion of agent [3] [22], based on the authority of **Error! Reference source not found.** In line with that concept, ants – natural or artificial, alike – should have at least five features to be seen as agents. Still, they show only three of them:

- a) *Autonomy*. Ants operate without direct external involvement and maintain some kind of control over their actions and internal state like any biological being.
- b) *Reactivity*. Its main connotation relates to action in response to stimuli. Vital for even the most undemanding living being (or IT application), reactivity is the key element in the very definition of ant-based models.
- c) *Continuity in time*. Agents are intransient. In both biology and IT, that does not mean immortality but a lasting process, able to achieve lasting activities.

However, ants are too trivial to be genuine agents since they lack the other two fundamental agency features:

- a) *Social ability*. It seems paradoxical that beings famous for their complex social life and cooperative work, lack this feature. The problem is that ants are unable to interact with their peers *directly*, i.e. *communicating* with each other – at least in the reductionist models based on their stigmergic behaviour.
- b) *Pro-activeness*. Ants do not plan their actions, do not take initiative – since they are not supposed to have specialized knowledge (the source of goals, planning and initiatives). Once more, that is assumed in *modelling* based on *stigmergy*– indeed, from an engineering perspective it is irrelevant if ants reveal teleological behaviour or if they do not. Moreover, it is debatable whether artificial ants shall be able to take some initiative or not. In the research related to this paper, the aspect was yet avoided but the authors do not exclude a priori such a prospect.

Synergy. Although *ants* behave rather as robots than as agents, the system they belong to is not a “multi-robot” system, but a “multi-agent” one. This wonder is due to the synergistic effect of their interaction: beyond the *individuals* (ants or ant-like entities), the *team* (colony, society, system) comes out [3]. That is the pre-terminological meaning of “synergy”, due to Aristotle: the whole is stronger than the sum of its parts. On the other hand, here it is not necessary to consider all of Haken’s principles [28]:

- a) “subsystems slaved by the system” is rather irrelevant since the system remains unexplained;
- b) “cooperating subsystems” does not apply directly since ants do not communicate;
- c) the “threshold” principle is of unquestionable importance in both nature and IT. A few ants are surely unable to run a colony and, as well, a few artificial ones are unable to solve TSP no matter how long they try. However, from an engineering perspective, this principle is here beside the point (because the initial number of entities is always much above the threshold);
- d) in contrast, the “self-organization” principle is crucial (see next Section). Since not all commented upon above is suggested by the more usual term used in similar contexts (“emergent synthesis”), at least in this paper, the term “synergy” is preferable.

3.2. Premises and Criteria

In the context above they are:

- a) To be usable, the system should be of reduced *cognitive complexity*. *Corollary*: the interface should be ergonomic and extensible to other applications subfields.
- b) To be affordable, the system has to be of manageable *structural complexity*. *Corollary*: the problem-classes chosen have to admit approximate solutions, to prevent failure under combinatorial explosion (for instance, exhaustive search algorithms are outside the scope of this paper and should be used only as elements for generating a performance metrics).
- c) To be workable, as well as to allow evaluation, the undertaking has to avoid starting from scratch (a “tabula rasa” stance impairs any genuine comparative assessment). *Corollary*: for both theoretical and practical reasons the approach should be based on “micro-continuity” [23] (for instance, the test-bed problem should be a familiar “workhorse” – this is why TSP advanced from a toy problem to a benchmark generator).
- d) Beyond the paper’s aim, practicality itself entails that the behaviour of any biological entity 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 become compliant problem-solving tools. *Corollary*: after “squeezing” the standard algorithms by fine-tuning their parameters (as in [18]), in order to reach the target, symbolic processing has to be directly brought in.

First *synergy* itself must be deeper investigated, in its relationship to both complexity and symbols. The results achieved will highlight the directions for selecting and combining paradigms. On this groundwork will be set the design in three stages:

- a) proposing specific mechanisms to graft symbolic components onto the sub-symbolic foundation;
- b) tailoring the mechanisms to sub-fields (here, primarily, manufacturing control);
- c) building an experimental model as a test-bench for this sub-field but, considering also future extensions.

4. Swarm, Synergy, Stigmergy, Symbols

The best-known source of *synergy* is *swarm*. Ants achieve it through *stigmergy*. However, engineering requires also *symbols*.

4.1. Three Sources Synergy Springs from

The sources are ordered according to increasing individual complexity:

Homogeneous Amassing of Many Simple Entities. Synergy is intrinsically linked to *multiplicity*. Obviously, if it would be only “one part”, the “whole” could not be stronger than itself. However, multiplicity implies

parallelism – not just because the „parts” creating the swarm *coexist* but, because they *interact* incessantly (imposed by the real-world dynamics). This is evident in all sub-symbolic paradigms: artificial neural networks are founded on massive, fine-grain parallelism – as a premise for connectionism – while evolutionary algorithms yield relevant results only with numerous populations.

Heterogeneous Interacting of Few Complex Entities. Here fine-grain parallelism is replaced by coarse-grain one: the entities are dissimilar and their interaction creates “added value” rather through complementarity (as in the well-known case of symbiosis). Widening the perspective, physical or biological entities could be (conceptually) substituted by areas of expertise, sub-fields, or paradigms. For instance, hybrid-architecture agents manifest a “second degree” synergy. In this light in **Error! Reference source not found.** the synergistic effect of combining programming paradigms is explored. Of course, the synergy due to four paradigms flowing together into the stream of soft computing is even more promising.

Trans-Disciplinarity. The third source matches Prigogine’s idea, mentioned in Section 1, regarding domain interfaces. The prefix “trans” (instead of the more usual “inter” or “multi”) emphasise the trend towards osmotic-like confluences based on common ontologies.

4.2. A Common Idea: Synergy Likes Parallelism.

Synergy implies *multiplicity*; multiplicity entails *parallel interaction*. This implies two kinds of complexity: *structural* (regarding the system) and *cognitive* (regarding the way the system is perceived by its users). From another perspective, complexity may lie at two levels: *entity* (the individual parts are complex) or *system* (the whole is complex). In this respect synergy is exclusive (ant colonies are a first-class example): entities should be simple, while complexity should emerge through the huge number of interacting components at system level. That may be intriguing since, sometimes, the way the system works is hidden (ant colonies are extremely effective in long-distance expeditions but where is their high-quality database?). Fortunately, if it works, it is not an engineering problem and, anyhow, there are precedents in other domains (for instance, in physics: where lies the global information about quantic numbers that enables Pauli’s exclusion principle?).

4.3. A Heterodox Idea: Does Synergy Dislike Symbols?

From the above it looks like that not even nature with its huge resources can afford to deal with very many complex entities and that the key principle is self-organization, the “order emerging from disorder”. A self-organising system is working without central control, where the behaviours of its entities lead to an emergent coherent result. According to Heylighen [30], there are at least seven characteristics of self-organising systems. In this paper’s context four of them are relevant: a) global order from local interactions; b) distributed control; c) robustness and resilience; d) non-linearity and feedback.

From a slightly different perspective, the AI paradigms could be listed according to their (decreasing) relationships to symbols (the labels are ad-hoc and are used only to underline the idea in ordering the list; in the parentheses are given examples):

- a) *Symbolic* (the Newell-Simon hypothesis).
- b) *Vague symbolic* (fuzzy logics).
- c) *Crypto-symbolic* (“situated automata”: the symbols exist but are hidden through compilation).
- d) *Pseudo-symbolic* (artificial neural networks: symbols are only in embryo, in configuration patterns, weights, etc.).
- e) *A-symbolic* (the physical-grounding – ethological – paradigm applied in Brooks’ automata challenges the need for symbol-based explicit representation, not its existence).
- f) *Non-symbolic* (evolutionary algorithms: symbols do not appear, intelligent behaviour emerges randomly, supposed there are enough tries).

Of course, SC has its place on the end. The list lets a (rather heterodox and confusing) idea come to light: 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]. In other words, from an engineering point of view, synergy emerges easier in sub-symbolic contexts. Hence:

- a) SC, with its vast intrinsic synergistic potential is a good choice for exploiting the first synergy source mentioned in section 4.1.
- b) To get synergy also from the second source, components based on other paradigms could be graft upon sub-symbolic structures. Due to complementarity, the best choice is the symbolic paradigm.

parallelism – not just because the „parts” creating the swarm *coexist* but, because they *interact* incessantly (imposed by the real-world dynamics). This is evident in all sub-symbolic paradigms: artificial neural networks are founded on massive, fine-grain parallelism – as a premise for connectionism – while evolutionary algorithms yield relevant results only with numerous populations.

Heterogeneous Interacting of Few Complex Entities. Here fine-grain parallelism is replaced by coarse-grain one: the entities are dissimilar and their interaction creates “added value” rather through complementarity (as in the well-known case of symbiosis). Widening the perspective, physical or biological entities could be (conceptually) substituted by areas of expertise, sub-fields, or paradigms. For instance, hybrid-architecture agents manifest a “second degree” synergy. In this light in **Error! Reference source not found.** the synergistic effect of combining programming paradigms is explored. Of course, the synergy due to four paradigms flowing together into the stream of soft computing is even more promising.

Trans-Disciplinarity. The third source matches Prigogine’s idea, mentioned in Section 1, regarding domain interfaces. The prefix “trans” (instead of the more usual “inter” or “multi”) emphasise the trend towards osmotic-like confluences based on common ontologies.

4.2. A Common Idea: Synergy Likes Parallelism.

Synergy implies *multiplicity*; multiplicity entails *parallel interaction*. This implies two kinds of complexity: *structural* (regarding the system) and *cognitive* (regarding the way the system is perceived by its users). From another perspective, complexity may lie at two levels: *entity* (the individual parts are complex) or *system* (the whole is complex). In this respect synergy is exclusive (ant colonies are a first-class example): entities should be simple, while complexity should emerge through the huge number of interacting components at system level. That may be intriguing since, sometimes, the way the system works is hidden (ant colonies are extremely effective in long-distance expeditions but where is their high-quality database?). Fortunately, if it works, it is not an engineering problem and, anyhow, there are precedents in other domains (for instance, in physics: where lies the global information about quantic numbers that enables Pauli’s exclusion principle?).

4.3. A Heterodox Idea: Does Synergy Dislike Symbols?

From the above it looks like that not even nature with its huge resources can afford to deal with very many complex entities and that the key principle is self-organization, the “order emerging from disorder”. A self-organising system is working without central control, where the behaviours of its entities lead to an emergent coherent result. According to Heylighen [30], there are at least seven characteristics of self-organising systems. In this pa-per’s context four of them are relevant: a) global order from local interactions; b) distributed control; c) robustness and resilience; d) non-linearity and feedback.

From a slightly different perspective, the AI paradigms could be listed according to their (decreasing) relationships to symbols (the labels are ad-hoc and are used only to underline the idea in ordering the list; in the parentheses are given examples):

- a) *Symbolic* (the Newell-Simon hypothesis).
- b) *Vague symbolic* (fuzzy logics).
- c) *Crypto-symbolic* (“situated automata”: the symbols exist but are hidden through compilation).
- d) *Pseudo-symbolic* (artificial neural networks: symbols are only in embryo, in configuration patterns, weights, etc.).
- e) *A-symbolic* (the physical-grounding – ethological – paradigm applied in Brooks’ automata challenges the need for symbol-based explicit representation, not its existence).
- f) *Non-symbolic* (evolutionary algorithms: symbols do not appear, intelligent behaviour emerges randomly, supposed there are enough tries).

Of course, SC has its place on the end. The list lets a (rather heterodox and confusing) idea come to light: 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]. In other words, from an engineering point of view, synergy emerges easier in sub-symbolic contexts. Hence:

- a) SC, with its vast intrinsic synergistic potential is a good choice for exploiting the first synergy source mentioned in section 4.1.
- b) To get synergy also from the second source, components based on other paradigms could be graft upon sub-symbolic structures. Due to complementarity, the best choice is the symbolic paradigm.

5. The Engineering Metamorphosis: From Ants to Agents

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 *synergistic potential*, this section aims at creating *problem-solving methods*. From this perspective, symbolic processing is not a (best) *choice*, but a *must*. Indeed, any engineering undertaking involves symbols, since design means to *project*, and any relation to the future implies symbols.

5.1 The Filtered Biological Model.

Why filter it? Because any problem needs its tailored biological model to graft upon its own symbolic-processing requirements. Put bluntly, ants do not care about TSP, whereas salespersons have no reason to smell pheromones. Thus, to get the best from both “worlds” we have to boost the model’s strengths and to reduce its weaknesses.

From an application point of view, “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” [25]. Thus, the major advantages of SC [25] [31]:

- 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 negative feedback provides “forgetting” fruitless paths), and *protection* (without explicitly conveyed information, confidentiality is preserved: a paramount asset for military applications).

The price to pay for those strengths can be high, because of the unavoidable differentiation between artificial ants and their natural counterparts. Moreover, such differences are heavily problem dependent. For instance, for operational research applications, such significant discrepancies 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.

Concluding, in comparison to natural ants, artificial ones, must not be so:

- a) *Many* (the number of artificial ants must be reducible in order to save affordability on usual configurations for very complex problems; e.g., in the experimental model, for a 1.5 GHz processor 200 ants needed about 1 second to travel through 200 places).
- b) *Dumb* (if artificial ants are implemented as threads, their functionality can be increased, diminishing the distance towards genuine agents).
- c) *Uniform* (following the example of bees, ants can have one of several behaviours).
- d) *Routine-driven* (corollary: if several behaviours are available, the same ant should be able to switch from one behaviour to another).
- e) *Independent* (the ant behaviour should not be exclusively limited to reaction to stimuli from the stigmergic environment but also influenced through central coordination).
- f) *Time-indifferent* (this the most challenging problem since to give ants a “sense of time” is very difficult even for high-level agents; in the experimental model, this aspect is approached in a still unpretentious way, through thread priorities). To implement this flexibility, several mechanisms are needed.

5.2 Mechanisms to Add Symbolic Processing.

Since technological *constraints* (mainly, processing power) and problem *complexity* are unavoidable, the only way out is to exploit technological *freedoms*. These are based on symbolic processing and must be kept at minimum (to carry on the advantages of SC). They can be grouped in three general mechanism classes [[6]] which have to be instantiated for every specific problem (in parentheses are examples from two fields: TSP since it is the practical test bench, and MC, since it is the research target):

A. Adapting the Environment (the Journey from Pheromones to Maps).

Albeit having to return to their nest (e.g., for TSP that is the departure town and for MC the duration of a technological operation), 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*. Thus, artificial ants can mark their paths with two kinds of pheromone: their own individual pheromone and the colony's one, trying to avoid the paths that have already been marked with their own pheromone at previous tours. The result is a higher map exploration [6]. On the other hand, *the pheromone evaporation rate can be adjusted* depending on the dynamics of the problem context. Thus, for highly dynamic environments it is helpful to have a higher rate in order to avoid local optima [32]. This mechanism has been used extensively in [6].

B. Instituting (Limited) Central Coordination (the Queen of Ants Becomes Supervisor).

In this respect, the distance to the biological model can be large and will surely increase for MC. Indeed, the MAS itself may have a vital role in guiding the search. For instance, 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 UDH [3]. That means to give the user the power to guide dynamically the search process (i.e. to modify the most sensitive parameters), after assessing the existing partial results (see next section). For instance, the positive/negative feedback can be adjusted to the current state (e.g., when a path is clearly un-promising, evaporation can be instantaneous).

This method proved to be very effective despite an important drawback: because of the user interference, the program performance cannot be measured and, as a result, its effectiveness is evaluated subjectively. Fortunately, this problem can be partially 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. (This aspect is not yet dealt with in the experimental model, since it involves deeper investigation regarding control-flow strategies.)

C. Boosting the Agents (Clever Ants Are Expensive).

Artificial ants could be smarter than natural ones, tending to become closer to agents. However, should they be? It depends on the amount of symbolic processing required:

- a) *No*, if complexity can be managed via the two mechanisms described above (as the case in [6]), since the cost of each functional enrichment has to be multiplied by the number of ants, before considering its worth.
- b) *Maybe*, when flexibility is desirable and extensibility affordable, but effectiveness is still the main concern (as imposed now; for the current trade-off, see next section).
- c) *Yes*, in the larger context of investigating inter-paradigmatic synergy sources (the very target of this paper), since it is clear that if an entity changes over time its behaviour will improve.

If the agents have to be enhanced – no matter reasons, extent, or kind of problem – the first step is to give them genuine autonomy, i.e. implementing them as execution threads (see next section). The possibilities to apply those mechanisms in order to tune artificial ants are summarised in Table 1.

Table 1. Tuning Artificial Ants

Mechanisms	Adapting the environment	Limited central coordination (UDH)	Boosting the agents
Not so			
Many		●	
Dumb			●
Uniform	●	●	●
Routine-driven	●	●	●
Independent		●	
Time-indifferent		●	●

Table 1 is relevant suggesting that:

- a) Adapting the environment is a path with almost exhausted potential;
- b) UDH is useful in all cases, except when smarter ants are needed (indeed, intelligence has to be inoculated explicitly);
- c) If the ants are smarter they need less coordination through UDH.

5.3. Adapting the Mechanisms to Manufacturing Control

The sub-section describes briefly how the mechanisms detailed above are reflected in MC [5] [19]. Since effective modelling is very difficult, a substitute is possible through user-driven disturbances, giving the opportunity to get insights about the way the system will react to changes that are context-dependent and cannot be efficiently captured in a model. To cope with disturbances encountered in current manufacturing systems, the ants manifest multiple behaviours in order to react efficiently to uncommon circumstances (e.g., when the solution already built becomes obsolete because of many possible circumstances: machine breakdowns, changes in order requirements, etc.). In contrast with real ants' behaviour, MC poses special constraints. Firstly, in MC the system has to carry out simultaneously multiple goals (which often are not even stable in time). Secondly, there are conflicting objectives in allocating existing resources. Thirdly, some individual goals may differ from the society one – the collaboration terms are not predefined in advance but negotiated. All of these lead to a significant refinement of the biological model.

The main aspects of any MC (process plans, internal logistics, and resource scheduling) must be represented by dedicated agents [33]. Consequently, the product agent (PA) reflects the product types that are manufactured in the plant, the order agent (OA) correspond to a customer order and the resource agent (RA) mirrors a processing resource available in the plant. Since RA owns methods to control the corresponding production resource, it has to reflect its production capacity and functionality through pheromones available to neighbour agents. That means that the local pheromone should contain information concerning the resource ability to carry out a particular operation (what and how a service can be provided) and the resource availability to grant that operation (when a particular service can be operational). On the other hand, the PA should contain information about the product life-cycle, design, process plans, quality assurance procedures, etc. Moreover, given the possibility to reconfigure the plant, the products delivered in a plant may change over time. Likewise to RA, the PA exposes in the local pheromone the product range that can be manufactured on the existing plant (its capability) together with the time window when these products may be delivered (their processing availability). Considering the logistical information related to the job, the OA exposes in the locale pheromone the required operations to get the product processed and the temporal restrictions to achieve it. All these pheromones are lying in local blackboards attached to each agent. The pheromones are evaporating at an adjustable rate, fine-tuned through UDH.

To carry out SC the agents create artificial ants (with an adjustable frequency) to construct a possible solution for their own goal. To reduce complexity the pheromones are organized on different layers: a) the “*feasibility layer*” – containing pheromone traces that enable the other ants to avoid routes that fail to guarantee that their goal is achievable; b) the “*exploring layer*” – adding the time window to the previous layer in order to ensure the time constrains regarding resource allocation; c) the “*intention layer*” – adding the agents' preferences to the exploring layer. Each pheromone layer is meant for a different type of ant (“*feasibility ant*”, “*exploring ant*”, and “*intention ant*”). From the engineering perspective this is expressed by distinct ant behaviours even if they exhibit the same behavioural pattern: a) in each node, the ant is observing the pheromone placed on the blackboard attached to an agent, and is executing its assigned task (this task could vary from a simple interpretation of the locally available information to overwriting the existing one); b) once the task is executed, the ant decides to which of the neighbouring nodes to carry on searching.

6. Experimental Model

The main tasks of the experimental model are:

- a) To be the means to apply the three mechanisms described above to MC;
- b) To be the test-bench for quantitative assessments regarding those mechanisms (since benchmarks for MC are still in their infancy (see the “test-bench assistant” from [1]).

On the groundwork described in Section 5.3, the experimental model outlines a plant, organized around an automated storage and retrieval system (ASRS) supported by a transportation device (T) (Figure 1). The parts are stored and transported in containers along the working stations (WS1, WS2, WS3) in order to get processed through the possible operations (O1, O2, O3). The “*feasibility ants*” build up the feasibility pheromone layer located in blackboards attached to the physical outputs of each resource. Following this pheromone trace the “*exploring ants*” build up the exploring layer – a map of possible resources scheduling. In this map the “*intention ants*” try to find the most optimal solution for the whole plant.

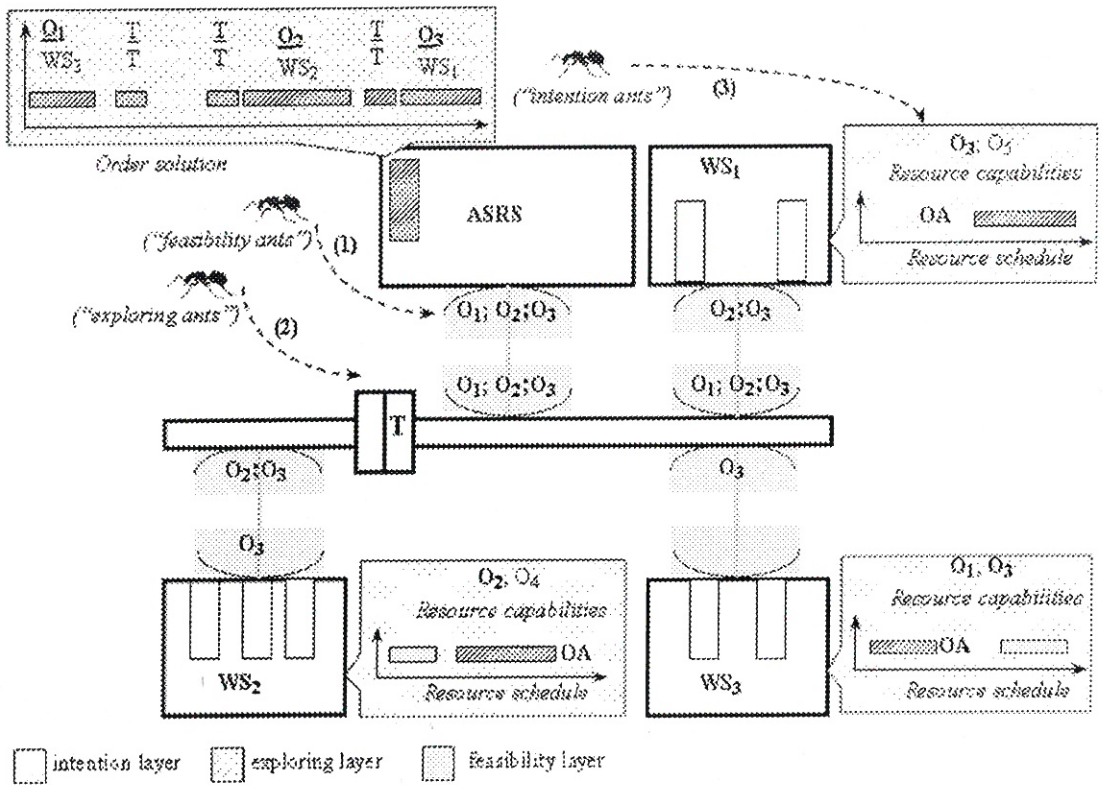


Figure 1: Artificial Phomones for MC (Adapted from [5])

The chosen variant for solving TSP was EAS [34] because it was the first relevant attempt to add symbolic-processing components to ant-inspired models. Having many parameters, this algorithm is complex enough to be efficiently fine-tuned for specific map configurations especially through UDH as illustrated in Figure 2. 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 MC); τ – 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). To express the effect of the above parameters, some ancillary parameters are useful in choosing the route in the graph: α – the pheromone intensity (exploitation); β – the generalized distance (exploration); e – the successful search history.

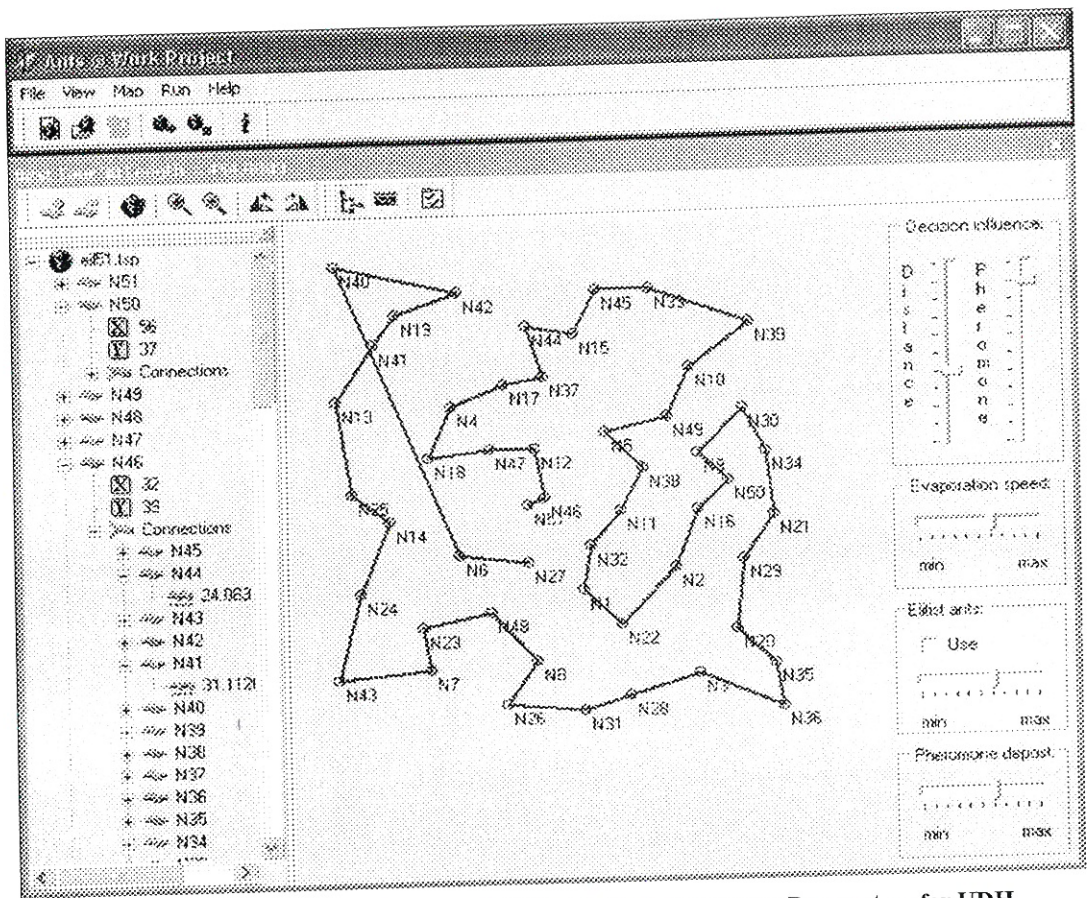


Figure 2: Snapshot of the Experimental Model, Capturing Some Parameters for UDH

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 needed. 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 3. 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 3: Solution Sensitivity Against α , ρ 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.

An example of applying UDH, in shown in Figure 4 presenting the system reaction to user-initiated variations of α and β (using the same interface components as in Figure 2).




Figure 4: Exploration Versus Exploitation Through UDH

Here, UDH was applied to speed-up convergence. Thus, after 104 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 closet 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).

7. Conclusions and Intentions

The conclusions are ordered against decreasing generality.

- a) Combining SC with symbolic processing components proved to have significant synergistic potential.
- b) This potential is expressed in application effectiveness especially for finding near to optimal solutions for NP-complete problems.
- c) To implement those effects in agent-based applications three mechanisms have been proposed. The most useful of them proved to be UDH.
- d) The mechanisms where adapted to MC but, considering the available benchmarks, they have been tested on TSP, with relevant results.
- e) The current version of the experimental model proved to be adequate for testing the mechanisms but is still insufficiently developed for the challenging constraints of MC.

Short-range intentions:

- a) To refine the experimental model in order to reflect adequately MC.
- b) To find a performance metrics able to asses solution quality when applying UDH.
- c) To investigate some open questions (e.g., the way to control non-deterministic ant behaviour through the parameter expressing randomly choosing routes in the graph, i.e. neglecting pheromonic traces).

Mid-range intentions:

- a) Implementing blackboard architecture for clever ants.
- b) Extending UDH for MC.

Long-range intentions:

- a) Going beyond MC towards even more challenging subfields as group-decision support systems (considering that any group-decision is an “emergent synthesis”).
- b) Building an interface agent.

REFERENCES:

1. ***www.ims-noe.org
2. 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.
3. BĂRBAT, B.E., **Agent-Oriented Intelligent Systems**, ROMANIAN ACADEMY PUBLISHING HOUSE, Bucharest, 2002. (In Romanian.)
4. FILIP, F.G., DONCIULESCU, D.A. and FILIP, CR.I., **Towards Intelligent Real-time Decision Support Systems for Industrial Milieu**, J. Studies In Informatics And Control 11, 4 2002.
5. ZAMFIRESCU, C.B., VALCKENAERS, P., HADELI, VAN BRUSSEL, H. and GERMAIN B.S., **A Case Study for Modular Plant Control. In Holonic and Multi-Agent Systems for Manufacturing**, Springer Verlag, Lecture Notes in Artificial Intelligence, 2744, 2003.
6. NEGULESCU, S.C. and BĂRBAT, B.E., **Enhancing the Effectiveness of Simple Multi-Agent Systems Through Stigmergic Coordination**, Proc. of the 4th International ICSC Symp. on Engineering of Intelligent Systems (EIS 2004), Madeira, 2004.
7. VAUGHAN, R.T., STØY, K., SUKHATME, G.S. and MATARIĆ, M.J., **Whistling in the Dark: Cooperative Trail Following in Uncertain Localization**, Proc. 4th International Conf. on Autonomous agents, Barcelona, 2000.
8. ROUMELIOTIS, S.I., PIRJANIAN, P. and MATARIĆ M.J., **Ant-inspired Navigation in Unknown Environments**, Proc. of the 4th International Conf. on Autonomous agents, Barcelona, 2000.
9. WEHNER, R., **Ant Navigation: Mini Brains – Mega Tasks – Smart Solutions**. (Invited paper.), Proc. of the International NAISO Congress on Information Science Innovations ISI'2001, Dubai, 2001.
10. FILIP, F.G., **Computer Aided Decision-Making: Decisions, Decision Makers and Basic Methods and Tools**. TECHNICAL PUBLISHERS & EXPERT PUBLISHERS, Bucharest, 2002. (In Romanian.)
11. PARUNAK, V.D., **Manufacturing Experience with the Contract Net**, In: Huhns, M.N. (Ed.), Distributed Artificial Intelligence. Pitman, 1987.
12. SHEN, W. and NORRIE, D.H., **Agent-Based Systems for Intelligent Manufacturing: A State-of-the-Art Survey**. Knowledge and Information Systems 1, 1999.
13. DORIGO, M. and DI CARO, G., **The Ant Colony Optimization Meta-Heuristic, New Ideas in Optimization**, MCGRAW- HILL, London, 1999.
14. DI CARO, G. and DORIGO, M., **AntNet: Distributed Stigmergic Control for Communication Networks**, Journal of Artificial Intelligence Research 9, 1998.
15. HAKEN, H., **Information and Self-Organization: A Macroscopic Approach to Complex Systems**, SPRINGER-VERLAG, Germany, 1988.
16. GLANSDORFF, P. and PRIGOGINE, I., **Thermodynamic Study of Structure, Stability and Fluctuations**, WILEY, New York, 1978.
17. SHALIZI, C.R., **Growth, Form, Function, and Crashes**, Santa Fe Institute Bulletin 15, 2.
18. NEGULESCU, S.C., **Applying Stigmergic Coordination in Optimisation**, Diploma work, "Lucian Blaga" University, Sibiu, 2003. (In Romanian.)
19. HADELI, VALCKENAERS, P., ZAMFIRESCU, C.B., VAN BRUSSEL, H. and GERMAIN, B.S., **Self-organising in Multi-agent Coordination and Control using Stigmergy**. In Self-Organising Applications, Springer Verlag, Lecture Notes in Artificial Intelligence 2977, 2004.
20. BĂRBAT, B., CÂNDEA, C. and ZAMFIRESCU, C.B., **Holons and Agents in Robotic Teams. A Synergistic Approach**, Proc. of the 4th International NAISO Congress on Information Science Innovations, 2001.

21. CÂNDEA, C. et al, **Coordination in Multi-Agent RoboCup Teams**, Robotics and Autonomous Systems 36, Elsevier, Amsterdam, 2001.
22. BĂRBAT, B.E., **Holons, Agents, and Threads in Anthropocentric Systems**, Studies in Informatics and Control Journal 9, 3, 2000.
23. BĂRBAT, B.E., **Agent-Oriented Captology for Anthropocentric Systems**, Large Scale Systems: Theory and Applications 2001, F.G., Filip, Dumitrache, I. and Iliescu, S.S. (Eds.), Elsevier, IFAC Publications, 2001.
24. 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, 1959.
25. BONABEAU, E., DORIGO, M., and THERAULAZ, G., **Swarm Intelligence: From Natural to Artificial Systems**, OXFORD UNIVERSITY PRESS, New York, 1999.
26. 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.
27. WOOLDRIDGE, M. and JENNINGS, N.R., **Intelligent Agents: Theory and Practice**, Knowledge Engineering Review 10, 2, 1995.
28. HAKEN, H., **Synergetics - an Introduction**, SPRINGER-VERLAG, Berlin, 1983.
29. FILIP, F.G. and BARBAT, B.E., **Industrial Informatics – New Paradigms and Applications**, TECHNICAL PUBLISHERS, Bucharest, 1999. (In Romanian.)
30. HEYLIGHEN, F., **The Science of Self-Organization and Adaptivity**, The Encyclopedia of Life Support Systems, 2002.
31. PARUNAK, H.V.D. and BRUECKNER, S., **Ant-like missionaries and cannibals: synthetic pheromones for distributed motion control**, Proceedings of the 4th International Conference on Autonomous agents, Barcelona, ACM Press, 2000.
32. EYCKELHOF, C.J. and MARKO, S., **Ant Systems for a Dynamic TSP: Ants Caught in a Traffic Jam**, Proceedings of the ACO 2002 Conference, 2002.
33. VAN BRUSSEL, H., WYNS, J., VALCKENAERS, P., BONGAERTS, L. and PEETERS, P., **Reference architecture for holonic manufacturing systems: PROSA**, Computers In Industry 37, 1998.
34. DORIGO, M., MANIEZZO, V. and COLORNI, A., **The Ant System: Optimisation by a Colony of Cooperating Agents**, IEEE Transactions on Systems, Man, and Cybernetics, Part B, 26, 1, 1996.