Digital Semiochemical Coordination

Holger Kasinger [1], Jörg Denzinger [2] and Bernhard Bauer [1]

[1] Institute of Computer Science, University of Augsburg, Universitaetsstr. 14, 86135 Augsburg, Germany
[2] Department of Computer Science, University of Calgary, 2500 University Drive NW, Calgary (AB), Canada T2N 1N4
Email: {kasinger, bauer}@informatik.uni-augsburg.de, denzinge@cpsc.ucalgary.ca

Abstract: Indirect interactions by olfactory stimuli between living organisms are a powerful mechanism for self-organizing coordination in biology. Various adoptions of this paradigm for computer systems however are mainly based on the usage of digital pheromones, although these chemical substances are only one type that mediate indirect interactions. Biology provides an ingenious diversity of such substances, all grouped by the term semiochemicals. In this paper we adopt the principles behind semiochemical coordination in biology and present a model that defines a coarse-grained architecture of self-organizing computer systems based on indirect interactions. This model allows for any combination of semiochemical coordination mechanisms within one single system architecture, which will pave the way for an easier engineering of self-organizing solutions better adapted to complex problems. We further demonstrate how to efficiently combine different types of semiochemical coordination into one mechanism, based on pollination in biology, and evaluate its application to instances of pickup and delivery problems.

Keywords: biologically-inspired, decentralized coordination, engineering self-organizing emergent systems.

1. Introduction

An essential key leading to self-organizing emergent systems is the capability of self-organizing coordination. Beside other sources of inspiration for decentralized coordination mechanisms (DCMs), such as physics, economics, human societies, or social science that inspired mechanisms as e. g. presented in [5][8][10][18], biology serves as the major one. Living organisms use different forms of communication to manage coordination in a self-organizing manner, in which beside tactile, visual, and acoustic in particular olfactory stimuli play a major role. This latter form of indirect, stigmergic interaction by chemical substances through the environment has inspired various useful decentralized solutions particularly for multi-agent systems (MAS). In these solutions, stigmergic interaction is primarily based on the usage of *pheromones*, e. g. see [2][6][13]. Pheromones are chemical substances secreted by an animal to the outside, that causes a specific reaction in other members of the same species. The term pheromone is also often used for the general description of any digital information carrier involved in indirect interactions, e. g. see [14][17].

However, literature in biology [12] goes beyond the sole use of pheromones for stigmergic interactions, as it pays attention to the different species involved in this form of communication. It divides the corresponding chemical substances into two separate groups: on the one hand side *pheromones*, on the other side *allelochemics* as enabler of interactions between members of different species, which can be further divided into *allomones*, *kairomones*, *synomones*, and *apneumones*. All substances are encompassed by the term *semiochemical*.

Because one of the major problems in engineering self-organizing emergent systems currently is the identification of appropriate mechanisms enabling the required self-organizing coordination [1][4][7], in this paper we adopt the principles behind semiochemical coordination in biology. This more fine-grained differentiation of stigmergic interactions will help software engineers in choosing the most appropriate mechanism(s) of semiochemical coordination to design self-organizing emergent solutions better adapted to complex problems. In order to support this design process, we present a conceptual model for digital semiochemical coordination (DSC) that defines a coarse-grained architecture of self-organizing computer systems based on indirect interactions (Section 2). This high-level model allows for any combination of semiochemical coordination mechanisms within one single system architecture, according to the specific needs of different application areas. Thus, the concepts of indirect interaction become applicable for a wide field of problem classes.

One mechanism of semiochemical coordination is the coordination by synomones, where both species involved benefit from the created coordination. An example of this symbiosis is the pollination of flowers by honey bees. A computational adoption of this paradigm, as proposed in [9], enables robust and flexible solutions to individual problems in various areas, such as production and logistics, traffic and mobility, as well as health care. Experimental results however revealed that in some cases the sole coordination by digital synomones does not provide all the required coordination for solving certain problems efficiently. But the biological classification of semiochemicals guided us to the idea to combine the coordination by synomones with other appropriate types of semiochemicals, which improved the efficiency of solutions significantly.

Based on these experiences, we present a more general conceptual model of pollination-inspired coordination (PIC) (Section 3). This model instantiates the DSC model (considered as a meta model) and demonstrates, how to combine different semiochemicals in an efficient manner. We instantiate our PIC model to solve instances of the General Pickup and Delivery Problem (GPDP) [16] (Section 4) and our experimental results show the usefulness of both the DSC model and the extended PIC model (Section 5).

2. Semiochemical Coordination

2.1 Terminology

The term *semiochemical* (from greek *simeon* - "a mark or signal") is used to describe the chemicals involved in the indirect interactions between organisms. Semiochemicals are subdivided into two major groups, *pheromones* and *allelochemics*.

The term *pheromone* (from greek *phereum* - "to carry" and *horman* - "to stimulate") is used to describe the chemicals involved in intraspecific interactions, i. e. between members of the same species. The perception of a pheromone may result in an immediate behavioral response or a complex set of physiological responses that are simply set in motion by the initial perception (which applies to allelochemics as well). Some common types are territory, dispersal, sex, trail, aggregation, and alarm pheromones.

The term *allelochemic* (from greek *allelon* - "one another") is used to describe the chemicals involved in interspecific interactions. It is defined as a chemical significant to organisms of a species different from its source. Allelochemics are divided into four subgroups, depending on whether the emitter, the receiver, or both benefit in the interaction. An *allomone* (from greek *allos* - "another", horman - "to stimulate") is defined as a chemical substance, produced or acquired by an organism, which evokes in the receiver a reaction adaptively favorable to the emitter, e. g. a plant emits allomones to deter herbivores. A *kairomone* (from greek *kairos* - "opportunistic") is defined as a chemical substance, produced or acquired by an organism, which evokes in the receiver a reaction adaptively favorable to the receiver but not to the emitter, e. g. secondary plant compounds help herbivores in finding plants to feed. A *synomone* (from greek *syn* - "with or jointly") is defined as a chemical substance, produced or acquired by an organism, which evokes in the receiver a reaction adaptively favorable to both the emitter and the receiver. This group of allelochemics includes floral scents and nectars that attract insects and other pollinators and substances that play an important role in symbiotic relationships. An *apneumone* (from greek *a-pneum* - "breathless or lifeless") is defined as a substance, emitted by a nonliving material which evokes a reaction adaptively favorable to a receiving organism, but detrimental to another organism, that may be found in or on the nonliving material. For example, parasites or predators are attracted to nonliving substances in which they may find another organism, their host or prey, by apneumones released from the nonliving substance.

In many cases not a single semiochemical takes effect on its own but different connections in a precisely defined mixture act as an effectively combined information carrier.

2.2 Model of Digital Semiochemical Coordination

A general model for DCMs based on indirect interactions simplifies the engineering of self-organizing emergent systems significantly. Such a meta model supports software engineers in constructing the coarse-grained architectural design of the later system, in particular when using a software engineering methodology [3]. Fig. 1 shows an UML 2 class diagram of our DSC model representing such a general model.

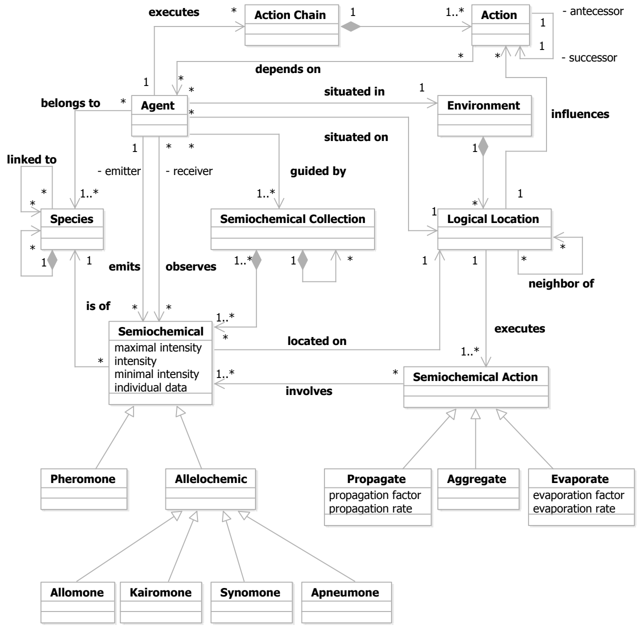

Fig. 1. The conceptual model of DSC

Because DSC is inspired by living organisms, in our computer world such an organism is usually seen as an agent in the systems. Any `Agent` belongs to at least one `Species`, which in turn may be composed of multiple species itself and be linked to other species, and is situated in an `Environment` consisting of multiple `Logical Locations` an agent may be situated on. Although logical locations will be very often instantiated by real locations, the model in this case remains general, so that the logical locations may span any space, for example it could also be a solution space of an optimization. Any agent acting as `emitter` further is able to emit digital `Semiochemicals` at its current location. We assign to each semiochemical also a species, which for better traceability is the same species as its emitter. A semiochemical can be further specialized into a `Pheromone` or an `Allelochemical`, which itself can be further specialized into an `Allomone`, a `Kairomone`, a `Synomone`, or an `Apneumone`. Every semiochemical is attributed with a `maximal intensity`, which is in general its initial intensity when emitted, a current `intensity`, as well as a `minimal intensity`, under which the semiochemical will disappear from the environment. In addition, semiochemicals may hold application specific individual data.

Any agent acting as `receiver` is able to observe semiochemicals located on its current location emitted by itself, by members of its species, or by another species, in case that the relevant species are linked together. Together with `Semiochemical Collections` these semiochemicals may guide the agent through the environment. Semiochemical collections, are a combination of multiple single semiochemicals, e. g. pheromone paths, or a combination of multiple semiochemical collections itself. Based on these means, an agent executes `Action Chains` composed of appropriate `Actions`. On the one hand side the actions are influenced by the agent's current logical location, on the other hand may depend on other agents the agent is interacting with. In general a semiochemical is a dynamical information carrier that is managed by its current logical location. Any of these logical locations is able to execute three different `Semiochemical Actions` in order to enable stigmergy: (1) It may `propagate` semiochemicals to its neighboring locations according to a `propagation rate`. The amount that is propagated may be governed by a `propagation factor`. Propagation as such supports information diffusion and spreading. (2) It may `aggregate` semiochemicals, so that separate quantities of semiochemicals are perceived as one with a greater intensity. Aggregation in general is a mechanism of reinforcement and supports information fusion. (3) It may `evaporate` semiochemicals in order to forget old information that is not refreshed or reinforced by new semiochemicals, which supports truth maintenance of information in the environment. The decreasing of a semiochemical's intensity i (at a certain `evaporation rate`) is governed by a dispersion function of the form $i(t+1) = i(t) \times ef$, where ef is a constant `evaporation factor` between zero and one.

Note, due to the potentially large number of (different) semiochemicals generated by certain coordination mechanisms we have chosen the logical locations for the processing and maintenance of semiochemicals, although the information carrier itself would have been available for these tasks too, as e. g. realized in the TOTA approach [11] for tuples.

3. Pollination-inspired Coordination

With regard to self-organizing computer systems, a promising biological paradigm is the coordination by synomones, as e. g. performed by prokaryotes, invertebrates, algae, or plants. Thereby, two completely different species interact indirectly with each other, both benefiting from this symbiosis. One appealing example is the pollination of flowers by honey bees.

3.1 Pollination by Honey Bees

The interest and benefit of plants (as pollenizers) in pollination is the accomplishment of their reproduction. For this purpose, an important prerequisite is the transfer of pollen grains from the male reproduction organ of a plant's flower to the female reproduction organ of another plant's flower. This usually requires a pollinator, which is primarily a living organism such as diverse mammals, bats, birds, or insects (biotic pollination). In order to attract and guide biotic pollinators, plants provide diverse signals that are beside tactile or optical mainly olfactory in form of fragrances (synomones).

The interest and benefit of biotic pollinators in pollination in turn is the maximization of the success in their own reproduction by optimal foraging. The latter is measured by (1) energy gains per time unit, which requires the detection of essential nutrients, (2) minimal time for flower detection, and (3) short handling time during the flower visits. Plants facilitate optimal foraging of their pollinators, if those behaviors attend to the transfer of their pollen grains. Therefore, plants provide food such as pollen or nectar as a form of reward for their visitors, but they are not that generous, as that pollinators can renounce the pollen grain transfer to other plants.

The honey bee is classified as the most ecological biotic pollinator species, as it "sticks" to a chosen flower species during one trip out of the hive. Although honey bees are guided to the right food area by the waggle dances they perform in the hive, current research results [15] prove that upon their arriving in this area they start to search for the fragrances of the flowers. These fragrances are the key to locate and approach their specific targets exactly and to transport pollen grains between flowers of the same species effectively and efficiently.

3.2 General Model of Pollination-inspired Coordination

The computational adoption of this paradigm, as proposed in [9], enables robust and flexible solutions to problems that require the self-organizing coordination between multiple autonomous components of heterogeneous types. Experimental results however revealed (see later Section 5) that for the efficient solution of certain problems the sole coordination by synomones does not provide all the required coordination. We identified two aspects that affect the efficiency adversely:

- **Misdirection of pollinators by outdated fragrances:** Not until an unrefreshed fragrance is evaporated, it will cease guiding pollinators to its emitting pollenizer, even if the pollenizer is already pollinated. The higher the evaporation factor, the longer unrefreshed fragrances remain in the environment resulting in a deterioration of the overall efficiency of the system.
- **Attraction of multiple pollinators to one pollenizer:** Often crowds of pollinators converge on a single pollenizer although only one pollinator would suffice. Only the pollinator arriving first is successful in pollinating, while the rest has to go away empty-handed. This kind of undesired swarm movement greatly reduces the overall efficiency, as the unsuccessful pollinators are missed elsewhere.

Whereas in nature these aspects usually carry no weight, for computational problems we expect more efficient solutions. But the diversity of semiochemicals guided us to the idea to extend the coordination mechanism by two additional types of semiochemicals, which, as we will see, improved the efficiency of solutions significantly: (1) Because pollinated pollenizers have to keep further pollinators off from visiting, they have to emit **allomones**, which are in form of fragrances immediately propagated through the environment. Thus, an approaching pollinator observing the outdated synomones of this pollenizer will be kept away due to the additionally observed allomones and look for other pollenizers. (2) Because crowds are a result of missing coordination between pollinators, the latter have to emit **territory pheromones** diffusing in a small area including the information which pollenizer they intend to visit. Thus, pollinators following in a similar direction to the same target can observe these pheromones and use their existence to switch to another pollenizer.

Based on these extensions, Fig. 2 depicts a conceptual model of the new general PIC mechanism, representing an instance of the DSC model. Both a `Pollenizer Agent` and a `Pollinator Agent` instantiate an *agent* of the DSC model and are situated on a `Location` in a common `Environment`. Each of them belongs to a separate *species* (`Pollenizer Type` and `Pollinator Type` respectively), linked in such a manner that pollinator agents can observe `Synomones` as well as `Allomones` emitted by pollenizer agents. By an instantiation of *semiochemical collection*s, `Allelochemic Collections` composed of multiple of these `Allelochemics` guide a pollinator from its current location to suitable locations of a pollenizer. Additional `Pheromones` emitted by pollinators themselves and possibly composed in `Pheromone Collections` now additionally support this guidance.

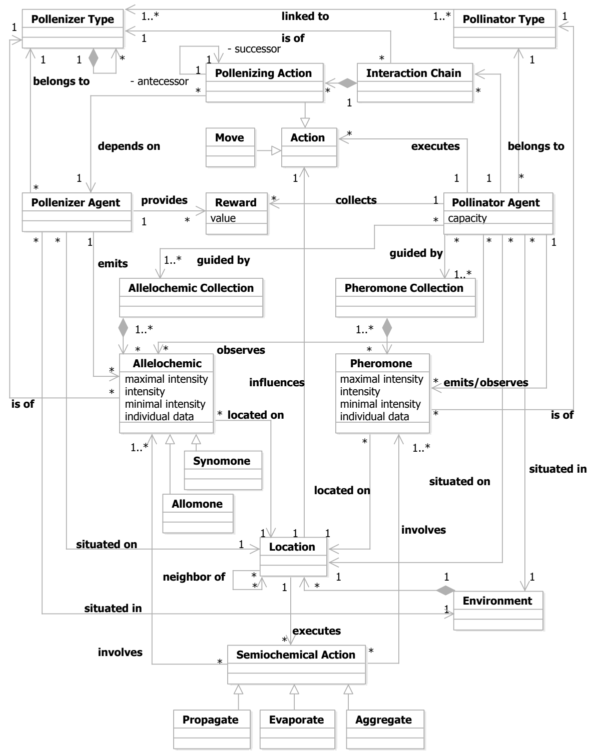

Fig. 2. Conceptual model of PIC

Due to this bouquet of different semiochemicals managed by specialized `Semiochemical Actions` of the locations in the environment, pollinator agents now are able to `Move` efficiently through the environment according to the specific needs of pollenizer agents. The biological pollination shows us the need for having a chain of `Pollenizing Actions` in order to fulfill these needs, namely a pollinator needs to visit at least one more pollenizer agent after the initial visit to the first pollenizer. In our instantiation of the DSC model to PIC, we have the concept of an `Interaction Chain` that instantiates *action chain* and represents an intended sequence of interactions with pollenizers that all need to be performed to fulfill a task the system developer wants to achieve. As already stated, an agent can be involved in several `Interaction Chains`, but we usually limit the number of ”open” chains by a `capacity`.

A pollenizer agent provides a `Reward` with a variable `value` (representing the estimated value of a reward for a pollinator in biology) for each pollenizing action performed. The value of a reward r provided by a pollenizer agent e at time t varies according to Eq. (1), where t_{init} is the initialization time of e and δ is a constant adjustment factor with $0 <= \delta <= 1$. This means that the longer a pollenizer agent waits for its pollination, the higher is the value of its reward, which makes the pollenizer more attractive to visit. The value of a reward hence is comparable to market prices: the lower the supply of pollinator agents, the higher the price a pollenizer agent pays for pollination. This value is encapsulated in the `individual data` of a synomone and thus made indirectly available to potential pollinator agents, which take it into account in their decisions. This competitiveness, while necessary for many applications, naturally enhances the problem of multiple pollinators attracted to one pollenizer. Similar to the increase of the reward value, the pollenizer agent also increases the initial intensity of synomones it emits, which enables a greater propagation area possibly covering more receivers.

$$r^e(t) := r^e(t_{init}) \cdot w \cdot \delta, w = \begin{cases} t - t_{init}, & \text{if } t > t_{init} \\ 1, & \text{if } t = t_{init} \end{cases} \tag{1}$$

Propagation Algorithm The guidance of pollinator agents to the reward providing pollenizer agents is supported by a kind of gradient field the synomones span around its emitters due to the propagation algorithm. Generally, if an arbitrary semiochemical sem^e of an emitting agent e is propagated from a location l_{dep} to a location l_{dest}, its intensity sem^e_i is decreased according to the following function:

$$sem^e_i(l_{dest}) := sem^e_i(l_{dep}) - d(l_{dep}, l_{dest}) \tag{2}$$

where $d(l_{dep}, l_{dest})$ is the distance between these two locations, measured e. g. in meters, hops, etc. If at l_{dest} no other semiochemical sem^f with $e = f$ of the same type (synomone, allomone, or pheromone) is currently present, sem^e will be stored at the end of the incoming connection to l_{dep}, cloned, and its clones propagated on all outgoing connections to other neighboring locations dependent on the applicable semiochemical action for this semiochemical type. If in contrast there is a sem^f, $e = f$, stored at l_{dest}, the intensities of both instances will be compared. If $sem^e_i > sem^f_i$, then sem^e will replace sem^f and will be cloned as well as propagated, otherwise sem^e will be discarded. Hence there is always at most one instance of a semiochemical type of a certain emitter stored at one location. Semiochemicals of different emitters as well as different semiochemical types of the same emitter, however, can be stored in parallel on one location. A pollinator agent following a certain synomone thus will always be guided to the synomone’s emitter on the shortest path, without the comparison of identical synomones.

Decision Mechanism The local decision mechanism of any pollinator agent p is very critical for an efficient coordination on a global level. The agent has to decide very quickly, based only on the locally observable information, which synomone to follow in order to act efficiently for itself but also for the entire system. This is exacerbated by the fact that p is allowed to handle multiple interaction chains Ic in parallel up to its capacity (defined as $Ic(p) < cap$), but only one interaction chain Ic_{pt} of a certain pollenizer type pt in parallel, i. e. $Ic_{pt}(p) = \{0, 1\}$. The pollenizer type of any Ic_{pt} is determined by the type of the pollenizer agent the first pollenizing action was executed with.

The decision mechanism hence is based on the estimated utility $u(s^e_{pt})$ of each observable synomone s^e_{pt} emitted by e of pt on the current location, but now additionally regards any present allomone a or pheromone ph. The calculation of this utility is governed by the following policies:

- *If p has not started an interaction chain Ic_{pt}, yet:*
    - The longer the time p follows s_{pt}^{e}, $t_{s_{pt}^{e}}$, the higher is the utility $u(s_{pt}^{e})$. This is a tribute to the effort that resulted from following the gradient of s_{pt}^{e} up to this location.
    - The nearer p is to the location of e, the higher is the utility $u(s_{pt}^{e})$. Thus, nearer pollenizer agents are privileged.
    - The higher the value of the provided reward by e, r^{e}, the higher is the utility $u(s_{pt}^{e})$. This privileges pollenizer agents that already wait longer for pollenizing actions.
    - If p observes an allomone a^{e} emitted by e, the moving to e will have no more utility for p. Note, a pollenizer agent emits an allomone as soon as a pollenizing action by a visiting pollinator agent fulfills its needs.
    - If p observes a pheromone ph_{e}^{o} of another pollinator agent o that also intends to visit e with $p \neq o$, for a certain amount of time s_{pt}^{e} will have no more utility for p and will be excluded from the utility calculation. If after this time p still observes s_{pt}^{e} but no more ph_{e}^{o}, s_{pt}^{e} will be regarded again for calculation.
- *If p has already started an interaction chain Ic_{pt}:*
    - The utility $u(s_{pt}^{e})$ has to be remarkable higher compared to the previous case. This guarantees that the processing or even closing of interaction chains is prioritized in contrast to the starting of new interaction chains.
    - If p observes an allomone a^{e} emitted by e, the moving to e will also have no more utility for p.
    - If p observes a pheromone ph_{e}^{o} of another pollinator agent o that also intends to visit e with $p \neq o$, the utility $u(s_{pt}^{e})$ depends on the amount of further suitable synomones s_{pt}^{f}, $e \neq f$, that can be observed on this location. If there are no more, the utility $u(s_{pt}^{e})$ is calculated as described by the first policy in this case. If there is at least one more suitable synomone s_{pt}^{f}, again for a certain amount of time s_{pt}^{e} will have no more utility for p.

These policies are incorporated into a utility function which is applied to every synomone observed at p's current location:

$$u(s_{pt}^{e}) := \begin{cases} b, & \text{if } Ic_{pt}(p) = 0 \wedge l(a^{e}) = 0 \wedge l(ph_{e}^{o}) = 0 \\ b \cdot \lambda, & \text{if } Ic_{pt}(p) = 1 \wedge l(a^{e}) = 0 \\ & \wedge \left(l(ph_{e}^{o}) = 0 \vee l(ph_{e}^{o}) > 0 \wedge l(s_{pt}^{f}) > 0\right) \\ 0, & \text{else} \end{cases} \tag{3}$$

where b is defined as $t_{s_{pt}^{e}} \cdot \alpha + \frac{s_{i\,max}^{e}}{s_{i\,max}^{e} - s_{i}^{e}} \cdot \beta + r^{e}(s) \cdot \gamma$, $s_{i\,max}^{e}$ is the initial intensity of s_{pt}^{e}, and α, β, $\gamma, \lambda >= 0$ are a adjustment factors. $l(a^{e})$ returns the amount of a^{e}, $l(ph_{e}^{o})$ returns the amount of ph_{e}^{o}, and $l(s_{pt}^{f})$ returns the amount of s_{pt}^{f} observed at the location.

After calculating the utilities of all synomones observed on its current location, p emits a pheromone ph_{e}^{p} in order to indicate its intension to visit e, with $u(s_{pt}^{e})$ is maximal, to possibly following pollinators, and moves to the location connected by the incoming direction of the synomone s_{pt}^{e}. If this is the location of e, p then will either start a new interaction chain or execute the next (or even the last) action of any open interaction chains. If no suitable synomones are found at all, the pollinator agent will either remain on its current location or search for nearby synomones by moving randomly, depending on the application needs.

Due to this modeling, PIC now becomes applicable for a much wider field of problem classes that require the self-organizing coordination between multiple autonomous components of homogeneous and heterogeneous types. PIC enables robust and flexible solutions in the face of dynamic changes. The components therefore have to be situated in a logical or physical environment, which may be extended with the needed infrastructure (for propagation, evaporation, etc.) whereas the environment structure may represent a part of or even the entire problem that has to be solved. Spatial movement of the components is supported whereas information about their spatial locations is indirectly exchanged.

4. Application Areas

This section exemplifies the applicability of PIC for the self-organizing solution of certain Pickup and Delivery Problems (PDPs), which are a specialization of GPDPs. Primarily, PDPs comprise the transportation of objects, goods, or persons by vehicles from a set of pickup stations to a set of delivery stations along with the optimization (as best as possible) of an objective function, e. g. minimization of route length. Practical examples are freight transportation systems or courier services.

4.1 Case Study: PDPs in Manufacturing Systems

In manufacturing systems PDPs comprise the transportation of incoming loads, i. e. packets, workpieces, materials, or products, between various production machines. In many cases, the transportation tasks are accomplished by Automated Guided Vehicle (AGV) systems. Usually, highly dynamic operation conditions govern these systems: transportation requests typically emerge irregularly and unpredictably, i. e. they arrive at any moment in time with variable load sizes to be transported. Particular areas may temporarily be closed for maintenance, as well as AGVs may leave and re-enter the system for the same reason. Single AGVs may fail and become obstacles on the way, or multiple AGVs on the same way may produce congestions. Traditionally a central management system receives all transportation requests, schedules them, plans the corresponding transportation routes, and instructs the AGVs to execute the calculated transportation tasks automatically.

AGVs based on an appropriate DCM however can handle these tasks of resource assignment and spatial routing as well as the mentioned dynamics itself, without any type of external or internal central control. Such a mechanism eliminates the bottleneck as well as a single point of failure in the shape of the central management system and reduces spending human and monetary resources.

4.2 Instantiation of PIC

In order to design AGV systems for the solution of such PDPs based on PIC, the concepts of the PIC model have to be instantiated by the concepts of the application domain. Thus, a `Production Machine` of the manufacturing system instantiates a *pollenizer agent* of the *pollenizer type* `Machines`, whereas an `AGV` instantiates a *pollinator agent* of the *pollinator type* `Vehicles`. Both production machines and AGVs are situated in a physical `Environment`, defined by the layout of the physical `Locations` and the connecting segments, on which the AGVs can `Move`. The *reward* of a pollenizer agent and its *value* are instantiated by `Virtual Money` with a `market price` by which the AGVs are paid for their service.

Synomones, allomones, and pheromones are instantiated by their digital counterparts. The digital semiochemicals are managed according to the model by the locations in the environment, that may be extended by an appropriate middleware executing the `Semiochemical Actions`. The *pollenizing actions* an AGV is able to execute are `Pickup Load` and `Delivery Load` that are comprised in an *interaction chain* that defines the load pickup as *antecessor* of a load delivery, which becomes the *successor*. An AGV may execute several such interaction chains in parallel, i. e. transport several loads at the same time, dependent on its `capacity`.

5. Experiments

For the experiments we have developed a simulator, which allows us to define different problem scenarios for arbitrary environment layouts. For ease of analysis, the simulator is based on a time management by iterations representing discrete time steps. Note, we do not simulate real world scenarios, but only simulate PDP under realistic conditions in order to prove the benefits of the DSC model as well as the applicability of PIC.

5.1 Simulation Settings

The experiments were executed on a sample environment layout (Fig. 3, which consists of 63 locations between which AGVs can move, connected by directed segments defining the possible paths that can be taken. A depot (d_0), the AGVs are housed in at the beginning of every simulation, is located in the middle. The production machines, the loads have to be picked up from and delivered to, are located on the left (pickup stations, $s_1 - s_6$) and on the right side (delivery stations, $s_7 - s_{12}$).

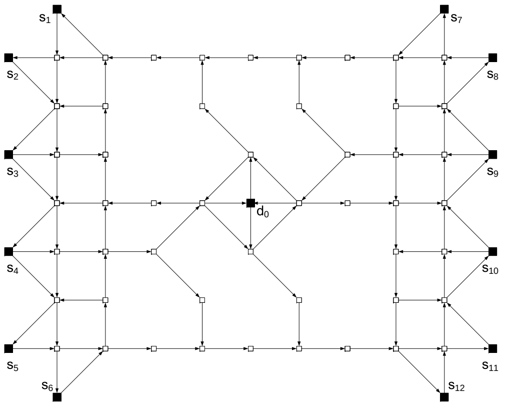

Fig. 3. Layout of the simulated PDP environment

Transportation requests are generated randomly, with a random pickup station and a random delivery station. The load size of each request varies randomly between 1 and 20. The capacity of an AGV is limited to 20 loads at a time, i. e. the processing of a request usually only requires a single AGV. The pickup as well as the delivery actions of a single load take one time step each. The move of an AGV along a horizontal or vertical segment takes three time steps, the move along a diagonal segment takes five time steps. The distances between all locations in the environment however are based on an Euclidean space. Every simulation run is limited to 500 time steps.

5.2 Results

For the evaluation of the performance improvements yield by the combinations of different semiochemicals, we have made four experiments:

- Coordination by means of synomones only (*S*)
- Coordination by means of synomones and allomones (*SA*)
- Coordination by means of synomones and territory pheromones (*SP*)
- Coordination by means of synomones, allomones, and territory pheromones (*SAP*)

In order to measure these improvements in contrast to the model proposed in [9], which is represented by *S*, here we choose two measures: (1) total travel costs (TTC), which is defined as the sum of the distances covered by all AGVs participating in a simulation run, and (2) total loads delivered *(TLD)*, which is defined by the sum of all loads the AGVs successfully drop at the defined delivery stations. The results of the experiments are presented in Fig. 4 and Fig. 5, which display average values over 100 runs of each simulation.

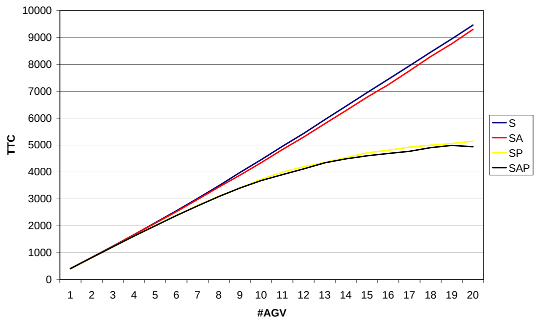

Fig. 4. TTC with increasing number of AGVs

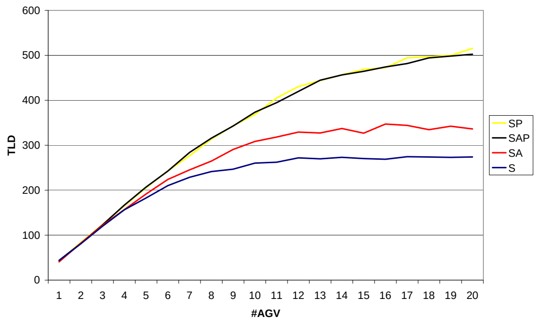

Fig. 5. TLD with increasing number of AGVs

Fig. 4 shows the evolution of TTC with an increasing number of AGVs. The two upper lines belong to *S* and *SA*, the two lower lines belong to *SP* and *SAP*. All combinations of semiochemicals result in a decrease of TTC compared to *S*, whereas the most significant improvements were made by *SP* and *SAP*. Coordination by the help of allomones (*SA* and *SAP*) however here only yields small improvements compared to the respective coordination without allomones (*S* and *SP*), which originates from the described simulation settings: new transportation requests arrive very frequently so every AGV finds fresh synomones to follow. Fig. 5 shows the evolution of TLD with an increasing number of AGVs. All combinations of semiochemicals result in an increase of TLD compared to *S*, which demonstrates their overall efficiency improvements, as less effort in terms of TTC was spent. Although now the improvements made by *SA* (second line from the bottom) become clearer compared to *S* (first line from the bottom), *SP* and *SAP* (the both lines on top) produce the most significant improvements again. Furthermore, provided that the AGV system has to meet specific throughput requirements, e. g. 400 loads per 500 time steps, the system now can be designed more appropriate according to these requirements by using *SP* or *SAP*.

Similar to every other DCM, also PIC requires the specification of certain parameters such as the frequency of semiochemical emissions, propagation factors, etc. However, further experiments not displayed here showed, that variations of these parameters in general only have influence on the scaling of the lines in the diagrams, not on the general appearance.

6. Conclusions and Outlook

In this paper we presented the DSC model, which is based on the principles behind semiochemical coordination in biology. The model defines a coarse-grained architecture of self-organizing computer systems based on any semiochemical coordination mechanism. It contributes to an easier engineering of future self-organizing systems, as it enables the adjustment, enhancement, and replacement of semiochemical coordination mechanisms, without necessary modifications of the architecture of these systems even after their deployment. The DSC model serves as a meta model, also for existing models such as e. g. digital pheromone path coordination [2], that can be integrated into any software engineering methodology. We furthermore demonstrated how a more fine-grained differentiation of indirect interactions helps in designing more appropriate solutions better adapted to complex problems. We therefore extended an existing coordination mechanism with additional semiochemicals, and evaluated the resulting efficiency improvements by a simulation. We described the essential parts of this new general PIC mechanism, presented its model which consequently represents an instantiation of the DSC model, and exemplified how to instantiate this model in turn.

Of course there exist certain similarities between PIC and other existing DCMs in charge of coordinating multiple autonomous components in a self-organizing manner, in particular digital pheromone paths [2], field-based coordination [10], and market-based coordination [5]. However, in contrast to pure pheromone path coordination, the propagation of allelochemics is less limited in space. Furthermore, the pollenizers become immediately ”visible” for the pollinators as soon and as long as they emit synomones and do not have to be discovered in an explorative manner by the latter. In field-based coordination the gradient parts do not evaporate over time whereas their strength usually increases with increasing distance to the gradient initiator, which sometimes leads to the problem of local minima when fields are combined. Market-based mechanisms in contrast usually require direct communication between buyers and sellers. Future experiments of systems based on PIC possibly will reveal potential for further extensions of the model, which may be achieved by the integration of additional semiochemicals with different properties.

References

[1] O Babaoglu, G Canright, A Deutsch, G A D Caro, F Ducatelle, L M Gambardella, N Ganguly, M Jelasity, R Montemanni, A Montresor and T Urnes, Design patterns from biology for distributed computing. ACM Transactions on Autonomous and Adaptive Systems (TAAS), Vol. 1, No. 1, 2006, pp. 26–66.

[2] S Brückner, Return from the Ant - Synthetic Ecosystems for Manufacturing Control, Phd thesis, Humboldt-Universität, Berlin, 2000.

[3] T De Wolf and T Holvoet, Towards a methodolgy for engineering self-organising emergent systems. Self-Organization and Autonomic Informatics (I), Frontiers in Artificial Intelligence and Applications, Vol. 135, 2005, pp. 18–34.

[4] T De Wolf and T Holvoet, Autonomic Computing: Concepts, Infrastructure, and Applications, CRC Press, 2007, ch. A Taxonomy for Self-* Properties in Decentralised Autonomic Computing, pp. 101–120.

[5] M B Dias, R Zlot, N Kalra and A Stentz, Market-based multirobot coordination: A survey and analysis. Proceedings of the IEEE, Vol. 94, No. 7, 2006, pp. 1257–1270.

[6] M Dorigo, V Maniezzo and A Colorni, The ant system: Optimization by a colony of cooperating agents. IEEE Transactions on Systems, Man, and Cybernetics Part B: Cybernetics, Vol. 26, No. 1, 1996, pp. 29–41.

[7] L Gardelli, M Viroli and A Omicini, Design patterns for self-organizing multiagent systems. CEEMAS 2007 Leipzig, Germany, 2007, Vol. 4696 of LNCS, Springer, pp. 123–132.

[8] D Hales, Choose your tribe! - evolution at the next level in a peer-to-peer network. ESOA 2005, Revised Selected Papers Utrecht, The Netherlands, 2006, Vol. 3910 of LNCS, Springer, pp. 61–74.

[9] H Kasinger and B Bauer, Pollination - a biologically inspired paradigm for self-managing systems. Journal of International Transactions on Systems Science and Applications, Vol. 2, No. 2, 2006, pp. 147–156.

[10] M Mamei and F Zambonelli, Field-Based Coordination for Pervasive Multiagent Systems, Springer Series on Agent Technology. Springer-Verlag, 2006.

[11] M Mamei and F Zambonelli, Programming pervasive and mobile computing applications: the tota approach. ACM Transaction on Software Engineering and Methodology, Vol. 2, No. 3, 2008.

[12] D A Nordlund, Semiochemicals: Their Role in Pest Control, John Wiley & Sons, New York, 1981, ch. Semiochemicals: A Review of the Terminology, pp. 13–28.

[13] L Panait and S Luke, A pheromone-based utility model for collaborative foraging. Proceedings of AAMAS 2004, New York, USA 2004, IEEE Computer Society, pp. 36–43.

[14] H V D Parunak, S Brueckner and J A Sauter, Digital pheromones for coordination of unmanned vehicles. E4MAS 2004, Revised Selected Papers New York, USA, 2004, Vol. 3374 of LNCS, Springer, pp. 246–263.

[15] J R Riley, U Greggers, A D Smith, D R Reynolds and R Menzel, The flight paths of honeybees recruited by the waggle dance. Nature, Vol. 435, 2005, pp. 205–207.

[16] M W P Savelsbergh and M Sol, The general pickup and delivery problem. Transporatation Science, Vol. 29, 1995, pp. 17–29.

[17] P Valckenaers, M Kollingbaum, H van Brussel, O Bochmann and C Zamfirescu, The design of multi-agent coordination and control systems using stigmergy. 3rd Int. Workshop on Emergent Synthesis Bled, Slovenia, 2001.

[18] Y Xu, P Scerri, B Yu, S Okamoto, M Lewis and K P Sycara, An integrated token-based algorithm for scalable coordination. AAMAS 2005 Utrecht, The Netherlands, 2005, ACM, pp. 407–414.