Design Pattern for Self-Organizing Emergent Systems Based on Digital Infochemicals

Holger Kasinger, Bernhard Bauer
Institute for Software and Systems Engineering
University of Augsburg
Augsburg, Germany
Email: {kasinger, bauer}@informatik.uni-augsburg.de

Jörg Denzinger
Department of Computer Science
University of Calgary
Calgary, Canada
Email: denzinge@cpsc.ucalgary.ca

Abstract

An essential element in the engineering of computer systems are design patterns that capture current best practice and knowledge about recurring solutions for standard problems. In case of decentralized autonomic computing systems, also known as self-organizing emergent systems, appropriate design patterns have to structurally describe decentralized coordination mechanisms along with information on which kind of macroscopic effects, the self- properties, can be achieved in which situations. In this paper we present a design pattern for self-organizing emergent systems coordinating by means of digital infochemicals. Infochemicals, in the natural context, are chemical substances that convey information in the interaction between two individuals. Because infochemical coordination is the most universally employed mechanism of communication in nature, there exists plenty of inspiring examples of decentralized coordination usable for the solution of complex problems in need of certain self-* properties. The presented design pattern captures the general biological principles behind infochemical coordination, which simplifies a systematical systems engineering. It extends existing coordination models, in particular pheromone-based coordination and digital semiochemical coordination, in terms of terminology, functionality, as well as generality, and thus becomes applicable to a much wider set of problem domains.*

1. Introduction

Future computer systems will be characterized by context-awareness, openness, locality in control, and locality in interactions [1]. They will contain a multiplicity of elements (which can be autonomous software entities such as agents as well as autonomous real-world entities with computing and networking capabilities such as servers, mobile devices, robots, or modern cars) whose actions and interactions cannot be controlled by a single element anymore. Thus, in contrast to centrally controlled Autonomic Computing (AC) systems [2], the desired self-managing behavior of these systems along with the required self-* properties have to emerge solely from the interactions between their constitutive elements. In literature, this class of systems is commonly referred to as Decentralized Autonomic Computing (DAC) systems [3] or self-organizing emergent systems [4]. An essential key for the autonomy of these systems is their capability of effective, spatial and/or temporal, decentralized coordination for purposes such as resource[1] allocation, load balancing, spatial distribution, dynamic clustering, group formation, or re-organization.

Spatial and temporal coordination between individuals is an omnipresent problem pertinent not only to computational individuals but also to natural individuals such as animals and plants. There, the purpose of coordination is rather the selection of food, the selection of mates, competition, or the avoidance of predators. However, living organisms rely on different, well-elaborated mechanisms of communication and coordination based either on radiational (light perception or visual), mechanical (tactile or auditory), or chemical (gustatory or olfactory) stimuli. These mechanisms are purely decentralized and consequently enable a high degree of flexibility, robustness, and scalability of the natural systems. Thus, the understanding of the principles behind these coordination mechanisms as well as their computational adaptation is of high importance to engineer effective, self-organizing emergent systems.

The most universally employed communication mechanism in nature with a plethora of inspiring examples is based on chemical stimuli [5], more precisely *infochemicals* [6]. In the natural context, infochemicals convey information in interactions between individuals. They are divided into two categories: *pheromones* that mediate intraspecific interactions, i. e. between organisms of the same species, and *allelochemicals* (*allomones*, *kairomones*, and *synomones*) that mediate interspecific interactions, i. e. between organisms of different species. Whereas the principles behind intraspecific coordination by means of pheromones are already well

1. A resource can be a task, power, bandwidth, space, (CPU) time, a device, a machine, etc.

understood and extensively used in the computational world (see e. g. [7], [8]), the adaptation and usage of interspecific coordination by means of allelochemicals, e. g. between plants and insects, are yet scarce, although they provide more appropriate solutions to complex problems in need of certain self-* properties. Thus, as biological foundation, Section 2 describes the general biological principles behind decentralized coordination based on infochemicals.

For the systematical engineering of self-organizing emergent systems based on these principles, however, it is important to capture these principles in a structured design pattern that can be instantiated and used according to specific needs. The use of design patterns in software engineering offers several advantages such as reducing design-time by exploiting off-the-shelf solutions and promoting collaboration by providing a shared ontology. Thus, Section 3 presents the addressed design pattern along with current best practice and knowledge about exemplary instantiations of this pattern. The pattern makes essential improvements to pheromone-based coordination [7] and digital semiochemical coordination [9], in terms of terminology, functionality, and generality. As a result, the pattern becomes applicable to a much wider set of problem domains, defining the basic architecture and coordination principles of the later self-managing system. To facilitate an easier engineering, the pattern additionally includes information on which kind of desired macroscopic effects, the self-* properties [3], can be achieved, and in which situations the pattern is appropriate. Finally, the remainder of this paper presents related work (Section 4) and concluding remarks (Section 5).

2. Principles of Infochemical Coordination

Chemical communication and coordination is a universal feature of life that occurs at all levels of biological organization, including the movement of cells or bacteria (called chemotaxis), the regulation of organs within an individual's body by hormones, as well as social behavior and ecological interactions among individuals by so-called semiochemicals or infochemicals. So far, several attempts have been undertaken to create a unique terminology for this latter type of chemicals. Reviews can be found in [10] and [11], for instance. We will adhere to the terminology of infochemicals [6], which provides an unambiguous classification schema compared to the terminology of semiochemicals [12].

2.1. Terminology

According to [6], an *infochemical* is defined as *"a chemical that, in the natural context, conveys information in an interaction between two individuals, evoking in the receiver a behavioral or physiological response that is adaptive to either one of the interactants or to both"*. Infochemicals are subdivided into pheromones and allelochemicals.

A *pheromone* mediates an interaction between organisms of the same species whereby the benefit is to the origin-related organism ([+,-] pheromone), to the receiver ([-,+] pheromone) or to both ([+,+] pheromone), whereas an *allelochemical* mediates an interaction between two individuals that belong to different species. Allelochemicals are subdivided into allomones, kairomones, and synomones.

An *allomone* is (a [+,-] allelochemical[2] that is) pertinent to the biology of an organism (organism 1) and that, when it contacts an individual of another species (organism 2), evokes in the receiver a behavioral and/or physiological response that is adaptively favorable to organism 1 but not to organism 2. Examples of allomone-mediated interactions are given by plants that emit chemicals (toxins) to deter herbivores as well as skunks that emit chemicals to keep putative predators away.

A *kairomone* is (a [-,+] allelochemical that is) pertinent to the biology of an organism (organism 1) and that, when it contacts an individual of another species (organism 2), evokes in the receiver a behavioral and/or physiological response that is adaptively favorable to organism 2 but not to organism 1. An example of kairomone-mediated interaction is given by mammals unconsciously releasing chemical cues that attract hungry mosquitoes from a far distance.

A *synomone* is (a [+,+] allelochemical that is) pertinent to the biology of an organism (organism 1) and that, when it contacts an individual of another species (organism 2), evokes in the receiver a behavioral and/or physiological response that is adaptively favorable to both organism 1 and 2. An example of synomone-mediated interaction is given by the pollination of plants by insects. Plants emit floral scents that attract insects and other pollinators to its location. While plants benefit in this interaction by the receipt of pollen grains from other plants, pollinators benefit by the collection of nectar or oils as a reward for their visit.

2.2. Functions and Effects

Apart from their benefits to origin-related or receiving organisms, infochemicals are also often divided by their function in the interactions between organisms, which is not trivial, because their functions are not mutually exclusive [13]. Thus, a given chemical can have several biological functions within a complex network of interactions.

Pheromones mostly function by influencing other members of the same species, not the individual that produced them. A non-exhaustive list of partly overlapping functional pheromone classes includes sex pheromones, alarm pheromones, aggregation pheromones, spacing pheromones,

2. To indicate more obviously the +/- relations between origin-related organism and receiving organism interacting by an allelochemical subtype, we will additionally annotate the respective relation to each allelochemical subtype, comparable to [+,-], [-,+], and [+,+] pheromones. This annotation is not part of the original terminology, probably due to legacy reasons.

home range pheromones, trail pheromones, and surface pheromones. The perception of a pheromone may result in an immediate behavioral response (releaser effect) or a complex set of physiological responses that are simply set in motion by the initial perception (primer effect) (cf. [14]). A primer effect is generally long-term, without an obvious immediate response. A releaser effect in contrast causes an immediate and reversible change in behavior mediated directly by the central nervous system.

Allelochemicals can be also differentiated according to their function. The list of functional classes of kairomones according to the function for the benefiting organism, i. e. the receiver, comprises foraging kairomones, sexual kairomones, aggregation kairomones, and enemy-avoidance kairomones. While all classes can be regarded having a releaser effect, sexual and enemy-avoidance kairomones may also have a primer effect. This classification can be transferred on allomones as well. In the case of synomones, the transfer of the classification is difficult. In many cases, a synomone will have different ecological functions for the emitter and for the receiver. Hence, the criterion for the classification can not be defined unambiguously (cf. [15]).

2.3. Communication

From a chemical point of view, infochemicals are chemical compounds, which can range from highly volatile to non-volatile. Slight genetic, dietary, and environmental differences make it improbable that any two organisms produce the same blend of volatile compounds. The production of an infochemical is regulated through hormones and signal transduction pathways. In contrast to hormones, which are produced in the endocrine glands, infochemicals are produced and discharged from exocrine glands. They are either secreted onto a surface area (e. g. for trial marking) or in most cases released into the surrounding air forming a cloud of vapor about the releasing organism. In a few cases, infochemicals are released into aqueous systems (cf. [16]).

In general, the distance through which an infochemical may transmit information is a function of the volatility of the compound, its stability in air, its rate of diffusion, olfactory efficiency of the receiver, and wind currents. In [16] a mathematical model is derived, which predicts the diffusion behavior of a volatile chemical in still air on the basis of these parameters. An essential factor of this model is the diffusion coefficient specific to each infochemical and an environment. Long-distance communication of a mile or more can be mediated by the use of stable compounds with high vapor pressures.

The perception of an infochemical is usually considered to be an olfactory process, although in some cases it may be gustatory, for example with infochemicals, which are transmitted in aqueous media or which are non-volatile. The perception of an infochemical then triggers an immediate behavioral response (releaser effect) or a delayed physiological response (primer effect) (see 2.2). However, a behavioral response requires the concentration of an infochemical to be high enough. This concentration is called the behavioral threshold concentration [16].

2.4. Benefits

Although infochemicals on their journey to a receiver have to pass through a highly variable environment affected by wind, temperature, moisture, and physical obstructions, the widespread use of chemical stimuli is indicative of the many advantages of this way of interaction (cf. [17]):

1) It is obvious that infochemicals can be used in situations where visual and auditory stimuli are absent or difficult to discern, e. g. at night, in dark burrows, or near loud sound sources.
2) Infochemicals can be easily distributed in both space and time, providing contextual or spatial information. Thus, the distribution, concentration, and qualitative aspects of infochemicals presumably provide organisms with key information about the resident(s) of an area, e. g. their physical size, reproductive state(s), mobility, motivational tendencies, or group size. The temporal aspects of infochemicals allow for the sending of 'time-coded' messages, e. g. the period of time since a given area has been visited or occupied.
3) Relative to other sorts of sensory stimuli, infochemicals can remain in the environment for rather long periods of time without jeopardizing the immediate safety of the signaling individual. If an animal released a continuous noise or visual signal in a manner analogous to leaving a long-lasting infochemical, not only would an inordinate amount of energy be expended, but predators would have a indication for locating it.
4) The sender and receiver need not be in close proximity for the communication to take place. This permits a resident to communicate to an intruder or rival that a given space is occupied, or that he or she is reproductively active, even though they are outside the range of hearing or sight.

3. Design Pattern: Digital Infochemical Coordination (DIC)

Although there exist different formats for describing patterns, it is generally agreed that the following sections are mandatory [18]: A *pattern name* (here the section heading), providing a clear, distinguishable identifier for the pattern. A *context* section, describing a situation when the pattern would apply. A *problem* section, giving a precise statement of the problem to be solved, in this case several problem characteristics. A *forces* section, describing items that influence the decision for the pattern, indicating trade-offs that might be made. A *solution* section, describing how the problem is being solved, balancing the forces. We additionally add three optional sections: A *rationale* section, explaining why the solution is appropriate for the problem along with its achieved (self-*) properties. An *examples/known uses* section, presenting a non-exhaustive list of examples/references that illustrate the application of the pattern. A *related patterns* section, mentioning other decentralized coordination mechanisms that may be also of interest for the solution of the problem. Because this format is well known in software engineering, e. g. [19] uses a similar format, the usage of the pattern is promoted.

3.1. Context

The problem in hand demands an autonomous solution that requires the decentralized coordination of multiple homogeneous and/or heterogeneous, more or less autonomous elements in order to achieve a common and globally coherent goal. The elements are situated in a physical or logical environment, which can be extended with an appropriate infrastructure, whereas the environment structure may represent a part of or even the entire problem to be solved. Some kind of spatial movement of the elements may be required or information about the spatial location of the elements has to be exchanged. The only possible way to coordinate are local estimates of global information. The desired solution has to be robust, flexible, and scalable in the face of frequent dynamic changes in the environment or the system.

3.2. Problem

- **Spatial routing:** Autonomous elements have to move or route themselves adaptively and as optimal as possible through the environment or problem structure. Elements may have to be attracted to certain locations or in a certain direction and be deterred from certain locations or directions, respectively.
- **Spatial awareness:** Autonomous elements have to be provided with abstract, simple yet effective contextual information, i. e. spatial information such as distance and/or direction to a location, facilitating the coordination process.
- **Homogeneity and heterogeneity:** Homogeneous and heterogeneous autonomous elements with different capabilities regarding their mobility, ability to communicate, or functionality have to be taken into account and are part of the problem or the solution.
- **Robustness and adaptiveness:** Autonomous elements have to move appropriately and achieve or maintain the globally coherent goal in face of dynamic changes in the environment, e. g. obstacles, failures, emerging/vanishing locations, emerging/vanishing pathways.
- **Openness and scalability:** Autonomous elements may leave or join the coordination process at any time and any location without affecting the overall performance negatively. In case of leaving, a graceful degradation is expected. In case of joining, a smooth and seamless integration is expected.
- **Various information sources:** Various sources can produce various types of information that have to be considered. The information has to be processed in a completely distributed and decentralized environment.

3.3. Forces

- **Centralization vs. decentralization:** In relation to a centralized approach, a decentralized approach usually causes a communication as well as coordination overhead, except the information to control the system is intrinsically distributed or every element has almost global knowledge about the system state. However, the global state usually can not be obtained without any further assumptions or restrictions. In return, in very dynamic environments a decentralized approach has no bottleneck or single point of failure.
- **Optimality vs. robustness/flexibility:** In an adaptive approach without central means to optimize its efficiency an optimal solution to a problem can not be guaranteed. On the other side, in face of frequent dynamic changes in the environment or in the system itself, a durable optimal solution does not exist at all. In these instances, a robust and flexible approach may be preferable to an optimal but inflexible approach.
- **Exploration vs. exploitation:** In contrast to only exploit already known information, new information has to be explored sufficiently in order to have an adaptive solution. This prevents the autonomous elements from trapping in local optima and supports finding new pathways or sources. On the other side, a too high level of exploration may result in very insufficient solutions.
- **Responsibility of the environment vs. the elements:** Effective coordination often requires intensive information processing and communication, which can be accomplished by the elements themselves or the environment they are situated in. In the former case, the elements may explicitly reason about the information and control which and when information is distributed. However, this may require complex reasoning algorithms and communication capabilities and is not recommended in dynamic environments. If in contrast the environment itself represents the needed coordination information by transparently processing and distributing it toward the elements, the elements will be able to use that information as a kind of "red carpet" which, when followed, achieves the global goals and avoids complex processing within the elements.

- **Greediness vs. purposefulness:** In decentralized approaches, the need for adaptive and flexible coordination usually rules out globally informed and purposeful decisions by the autonomous elements. Thus, the elements act "greedily" and try to exploit any information immediately, instead of disregarding some information in order to receive a greater benefit later.

3.4. Solution

For a more detailed structure, this section is subdivided into a *conceptual description* of the solution, a *parameter tuning* subsection describing the essential parameters that can be tuned in this solution, and an *infrastructure* subsection that describes the functionality that is required from an infrastructure to realize the solution.

3.4.1. Conceptual Description. The solution is inspired by the principles behind infochemical coordination in nature (see Section 2). To make these principles usable for decentralized coordination in computer systems, they have to be meaningfully adopted into a computational model. Figure 1 illustrates the description of this model by an UML diagram. In this model, a living organism is seen as an autonomous `Element`, situated in a spatial `Environment` consisting of multiple `Locations` the element may be situated on. Connections between the locations define the possible ways an element may choose from in order to move between the locations, whereupon the connections may be directed or undirected as well as of different length, depending on the physical conditions.

An element belongs to at least one `Type`, which in turn may be hierarchically composed to higher types, reflecting the taxonomic ranks in biology, as well as being linked to other types, reflecting interspecific relationships in biology. This allows for homogeneous as well as heterogeneous elements situated in the same environment interacting with each other. An element acting as `emitter` is able to emit digital `Infochemicals`, i. e. [+,-], [-,+], or [+,+] `Intra-type` infochemicals respectively `Inter-type` infochemicals, according to a specific `emission rate` into the environment, in order to communicate and coordinate with other elements indirectly. The abstraction of the biological terms pheromones and allelochemicals in the model by the terms intra-type and inter-type infochemicals allows on the one hand side still the use of certain pheromone or allelochemical types in biologically-inspired instantiations of this pattern, but on the other side also instantiations apart from biological background knowledge using neutral terms. A digital infochemical basically contains four attributes:

- `Individual information`, which reflects the biological role of an infochemical as a dynamic information carrier. The content of this attribute may vary between different applications, whereas at least the type of the emitting element is included in this information.
- Its `current concentration`, which reflects the dynamically changing concentration of diffusing biological infochemicals.
- A `threshold concentration`, which reflects the behavioral threshold concentration of living organisms regarding specific infochemicals. Admittedly, according to the object-oriented paradigm, this attribute should be ideally modeled as an attribute of an autonomous element. However, whereas in biology the diffusion of infochemicals proceeds up to the last molecule, this has no practical effect in the computational world, in particular not from an object-oriented perspective. Thus, if the current concentration of an infochemical falls below this threshold, it will not be propagated any further but removed immediately.
- Its `diffusion coefficient`, which reflects its biological pendant and thus allows for a very fine-tuned propagation radius and evaporation time specific to each infochemical.

An element emits an infochemical to the environment by handing it over to the location it is currently situated on. All locations in the environment are in charge of providing a stigmergic functionality to the elements. So every location is able to execute three different `Infochemical Actions`, each governed by a respective `Infochemical Policy`:

1) A location may `propagate` an infochemical to its neighboring locations according to an infochemical-specific `Propagation Policy`. The amount that is propagated is governed by a `propagation factor` and the infochemical-specific diffusion coefficient, both affecting the decrease of the infochemical concentration. The rate is governed by a `propagation rate`. Propagation as such supports information diffusion and spreading.
2) A location may `aggregate` different infochemicals of the same type according to an infochemical-specific `Aggregation Policy`, such that separate infochemicals are perceived as one with a greater concentration. Aggregation in general is a mechanism of reinforcement and supports information fusion.
3) A location may `evaporate` infochemicals according to an infochemical-specific `Evaporation Policy`. An individual `evaporation rate` governs the speed of evaporation, whereas the infochemical's diffusion coefficient and an `evaporation factor` govern the amount that is evaporated. Evaporation serves to forget old information that is not refreshed or reinforced by new infochemicals, which supports truth maintenance of information in the environment.

Due to these infochemical actions and policies an infochemical emitted by an element diffuses across the neigh-

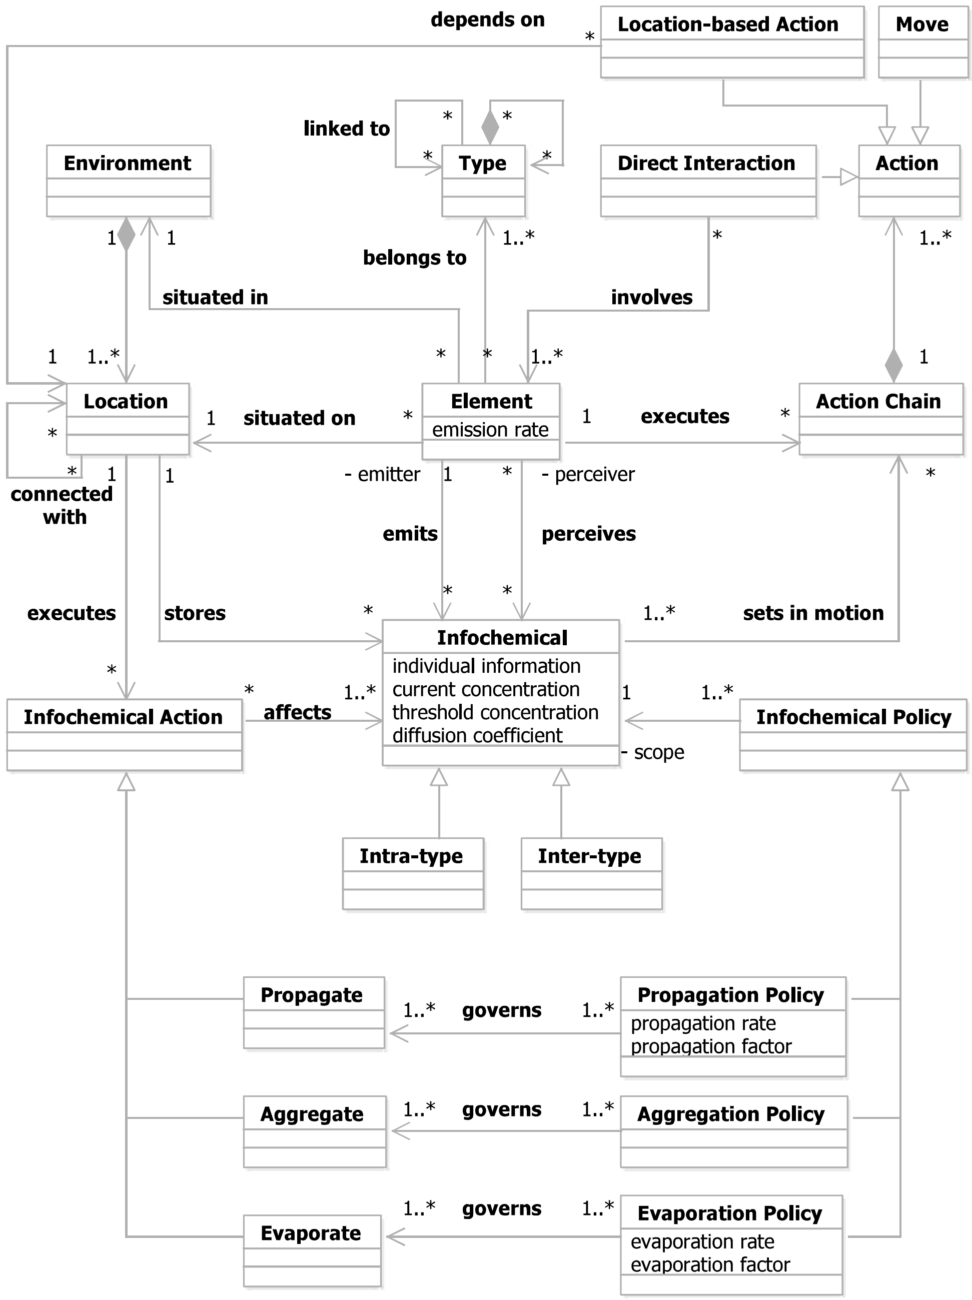

Figure 1. Conceptual model of DIC

boring locations in the environment, where every location affected by this diffusion stores a certain quantity of the infochemical, as long as the infochemical is to be removed. The diffusion in general will produce a kind of infochemical field around the emitting element. The location with the highest infochemical concentration of a field is the one the emitting element is currently situated on.

An element acting as `perceiver` is able to perceive infochemicals contained at its current location, possibly emitted by itself, by elements of its own type, or by elements of another type, in case that the relevant types are linked together. The perception of an infochemical by an element may set an `Action Chain` executed by the element in motion, reflecting the individual function of a given infochemical on a living organism. An action chain consists of at least one `Action`, which can be of the following types:

- `Move`: The element moves from its current location to a neighboring location, depending on the perceived infochemical field. The movement can be in the direction of the perceived infochemical, in the opposite direction, or equal to the concentration of the infochemical field, depending on the desired behavior of the element and the coordination to achieve.
- `Location-specific Action`: This action can have different shapes. On the one hand side, triggered by the perception of e. g. an alarm pheromone, an element may response by emitting alarm pheromones in turn. On the other side, if its current location has a special meaning to the element, the element may also execute a well-defined action at this location. For example, an ant picks up food at its destination location and drops food at its source location, while it emits pheromones at every location in between.
- `Direct Interaction`: The element directly interacts with one or more other elements. Direct interaction is only possible, if the interacting elements are situated on the same location. The reasons for direct interactions can e. g. be information exchange, reflecting the direct exchange of pheromones between ants or between bees in the case of surface pheromones, or resource exchange, reflecting the exchange of resources, e. g. pollen grains between flowers and bees for instance, but also any other act of communication or negotiation.

Independent of a perceived infochemical triggering a reactive action, an element may also execute certain action chains proactively, e. g. if no infochemicals can be perceived on the current location or infochemicals are to be emitted due to other reasons.

3.4.2. Parameter Tuning. Depending on the specific use of the coordination model, a number of parameters may be tuned. Their proper tuning has significant impact on the efficiency of the system.

- **Emission concentration:** The initial concentration of an infochemical when it is created by an element and emitted into the environment. Even though more than one type of infochemical is used for the coordination, the emission concentration value applies to all types.
- **Threshold concentration:** The minimal concentration of an infochemical. If the concentration of an infochemical falls below this value due to its propagation and evaporation, the infochemical will be removed from the environment. If there are different types of infochemicals used for the coordination, every infochemical type may have its own threshold concentration. A higher threshold concentration value narrows the diffusion area of an infochemical, a lower threshold concentration widens the area.
- **Emission rate:** The rate with which an infochemical is emitted by an element. If there are different types of elements participating in the coordination process, every element type may have its own emission rate. A high emission rate should be used, if information changes frequently or other elements have to be noticed of information change, a lower emission rate should be used, if information is rather static.
- **Diffusion coefficient:** If there are different types of infochemicals used for the coordination, every infochemical type may have its own diffusion coefficient, allowing for fine-tuned information diffusion areas. While some information may be required to be spread over a greater distance, other is not. Thus, an unnecessary communication overhead can be reduced. Note, elements may also be allowed to change the diffusion coefficient of their emitted infochemicals dynamically, in order to adapt to possible changed situations.
- **Propagation factor/rate** and **evaporation factor/rate:** Together with the diffusion coefficients, these parameters control the information spreading and truth maintenance individually for each participating type of infochemical. The settings may depend on the application-specific propagation/evaporation functions, which may be linear, degressive, distance-dependent, etc. In general, if evaporation proceeds very fast, information will be forgotten more rapidly. If evaporation is too slow, too many elements might will be attracted into the wrong direction.

Although the coordination model allows for a very fine-granular tuning, it is recommended to keep the amount of variable parameters as small as is necessary. The dynamics of parameter changes have to be taken into account.

3.4.3. Infrastructure. The application of this solution requires a kind of *infochemical infrastructure*, comparable to a pheromone infrastructure [7], to be provided by the locations composing the environment. Every location therefore has to provide a certain functionality:

- It has to accept and store infochemicals emitted by an element situated on it.
- It has to propagate, aggregate, and evaporate infochemicals according to the respective infochemical policies.
- It has to provide access to locally stored infochemicals for an element situated on it.

The realization of this infrastructure usually depends on the application domain. In case of a multi-agent system (MAS) running on a single machine, for example, the implementation will be simply a piece of software. In case of a distributed manufacturing system, for example, a kind of distributed middleware will be required.

3.5. Rationale

- **Routing:** In general, following the increasing concentration of an infochemical field is the shortest path to the emitter. Attracting elements to specific locations and to move in a specific direction according to an infochemical's concentration is supported as well as repelling an element from a location or direction. Obstacles are bypassed adaptively.
- **Feedback:** Feedback is given by the fact that infochemicals can change when changes occur in the environment, when the element that emits the infochemical decides to move, or is required to change the individual information, which is included in the infochemical. Other elements can then take the perceived change into account and react on it by, for example, emitting corresponding infochemicals on their own or changing the individual information of their own infochemicals according to the observed information (positive feedback). As outdated information is not refreshed anymore and gradually evaporates, negative feedback occurs. As such feedback cycles are established enabling self-organization.
- **Environment topology:** The structure of the environment reflects a part or even the entire problem to be solved. The distribution of the infochemicals along with their concentration guides the elements to the current solution of that problem.
- **Decentralized control:** Local decisions are made without requiring centralized reasoning or control. This way emerges a global self-organized motion pattern due to the related effects of elements emitting infochemicals and moving according to other observed infochemicals. The goals that are accomplished are not due to single elements, but due to the system as a whole without any central controller.
- **Information diversity:** Various types of information from various sources are supported. Coordination can be based on multiple types of infochemicals, even on multiple types of the same infochemical type, e. g. different types of pheromones.
- **Dynamic situations:** The environment is able to incorporate dynamic changes immediately, enabling the elements to react in a *flexible* way. New information is quickly integrated, while outdated information is quickly forgotten. Concepts such as exploration, information refreshment, and evaporation result in an *adaptive* coordination process. Elements thus can join and leave the system without significant disturbances to the global goal. This *openness* of the coordination process makes the mechanism extremely *robust*. Due to the intrinsically decentralization, the entire mechanism is *scalable* in problem size.
- **Information spreading/distribution:** Infochemicals are dynamic information carriers holding spatial information (direction to or distance to emitter) as well as individual information. The environment represents the distribution mechanism for this information and participates actively in the system's dynamics.
- **Processing complexity:** The elements are responsible for which information is emitted where and when into the environment. The environment is then responsible for storing, propagating, and evaporating this information. As such, the environment makes sure that not too much computational and communication burden is imposed on the elements themselves by automatically providing a dynamically adapting and propagating coordination structure that is immediately usable by elements. The context is represented expressively as infochemical fields, i. e. a kind of "red carpet", which represents how to achieve a coordination task by simply following the field. The coordination is achieved with very little effort and without complex reasoning by the elements. The latter indicates that the problem solving power resides in the local interactions instead of inside the elements' reasoning.
- **Self-* properties:** According to the characteristics of self-* properties in decentralized autonomic computing systems [3], the solution usually achieves smoothly evolving, ongoing, macroscopic, and adaptation-related self-* properties with possible functionalities in resource allocation, group formation, spatial shaping, or load balancing, to name a few. For example, resource allocation problems such as Job Shop Scheduling Problems (JSSPs) usually have self-* properties such as *"Workpiece throughput higher than x"*, *"Production time less than y"*, or *"Equal distribution of operations"*. The inherent adaptiveness, flexibility, and robustness of the coordination process yield to some extend self-configuring and self-healing properties.

3.6. Examples/Known Uses

Examples of the DIC pattern can be found in different problem domains forming various solutions:

- Multiple types of pheromones with varying propagation rates and thresholds have been used in [20] to coordinate agents of two species on a hexagonal grid in a military scenario. Evaluations have verified the performance improvements that were achieved due to the different pheromone configurations. Similarly, in [21] multiple types of pheromones have been used for self-optimizing trail behaviors by agents in domains which have obstacles, dynamically changing target locations, and multiple waypoints.
- In [22] a self-organizing emergent system was developed for the solution of Pickup and Delivery Problems (PDPs) in manufacturing systems. It was demonstrated that a combination of both intraspecific and interspecific interactions by different types of chemicals in the same system yields more powerful and efficient solutions. Experiments executed on the same system in [23] show that infochemical-specific diffusion coefficients along with individual propagation and evaporation policies result in a significant message reduction with a simultaneous performance increase.
- The idea of dropping information on specific locations that are picked up by other elements can even be found in search problems. In [24] a state-space search problem is solved concurrently by multiple cooperative agents, whereby the agents exchange information they found during their search, which can then be used by other agents arriving at these *logical* locations.

3.7. Related Patterns

DIC is naturally related to other mechanisms in charge of coordinating multiple autonomous elements in a self-organizing manner, in particular co-field coordination (CFC) [25], digital pheromone coordination (DPC) [7], and digital semiochemical coordination (DSC) [9].

CFC is similarly to DIC an instantiation of classical gradient field-based coordination, but inspired from physics, more precisely magnetic fields. Analogous to its physical counterpart, in CFC the gradient parts do not evaporate over time and hence – in contrast to DIC – have to be removed explicitly by the environment. The strength of the gradient parts usually increases in CFC with increasing distance to the gradient initiator, which sometimes leads to the problem of local minima when gradient fields are combined for the coordination of elements.

DPC can be readily considered as a specialization of DIC, as it supports only the coordination of homogeneous autonomous elements by means of digital pheromones. Also, in DPC all pheromones are propagated and evaporated equally without the possibility to differentiate between various types, which limits its general applicability.

DSC supports the decentralized coordination of homogeneous and heterogeneous agents too, however, DIC exceeds DSC in certain aspects that are essential for the design and an efficient coordination. First, DIC removes terminological ambiguities as it is based on a more precise terminology, which simplifies the engineering in terms of choosing the best chemical types for coordination. Second, DIC additionally reflects the biological multi-functionality of infochemicals by facilitating individual reactions of different elements on a given infochemical, which reduces the amount of unnecessary infochemicals in the environment. Third, DIC reflects also an infochemical-specific diffusion by the integration of a diffusion coefficient that takes effect on the propagation and evaporation of an infochemical. This again reduces the amount of infochemicals in the environment and speeds up the coordination process.

4. Related Work

Although design patterns for the coordination of agents are present for a while (see e. g. [26]), patterns describing decentralized coordination mechanisms enabling self-* properties are yet rare. Gardelli et al. [27] describe basic low level patterns common to various biological systems, such as *replication*, *collective sort*, *evaporation*, *aggregation*, and *diffusion*. On a similar level of abstraction, Babaoglu et al. [28] describe further patterns of biological coordination, including *plain diffusion*, *replication*, *stigmergy*, *chemotaxis*, and *reaction-diffusion*. A proper combination of some of these patterns may produce more complex patterns for self-organizing emergent systems. Even the DIC pattern uses some of these basic patterns, such as evaporation, aggregation, and diffusion, which together enable stigmergy.

In [29] patterns of higher level coordination mechanisms are described, such as pure *pheromone-based coordination*, *co-field coordination*, *market-based coordination*, *tag-based coordination*, or *token-based coordination*. A similar description format is used, what makes these patterns comparable to the DIC pattern. Every described pattern achieves slightly different self-* properties in comparison to DIC.

In [30] a so-called organic design pattern (ODP) is presented for self-* systems that consist of a set of independent agents interacting with each other and where reconfiguration/adaptation can be expressed as a reallocation of roles. However, a structured description of the pattern is missing.

5. Conclusions

Engineering a decentralized autonomic computing system or self-organizing emergent system implies explicitly considering a way to coordinate multiple autonomous elements to achieve desired macroscopic self-* properties. In this paper we have presented a design pattern that facilitates a decentralized coordination by means of digital infochemicals. The pattern captures the principles behind infochemical coordination in nature and describes them structurally to be used in a systematical engineering process. An engineer trying to get a sense of the pattern should first look at the context, problem, and solution sections of the description. Once he has determined that the pattern is of interest, he should look at the forces and rationale sections for guidance on determining whether the pattern is applicable to his particular situation and to its desired self-* properties.

The DIC pattern makes essential improvements to existing coordination mechanisms, particularly DSC, in terms of terminology, functionality, and generality. Due to the abstraction of the terminology to intra-type and inter-type infochemicals respectively, an engineer only has to answer two questions in order to use the right type of infochemical for a desired effect: Does the interaction have to take place between elements of the same type or between different types? Which party will benefit in the interaction? This simplifies engineering, because an engineer is no longer required to be a biological expert aware of the complex meanings of the different infochemical types in order to use their coordination principles for the solution to his problem. The usage of bio-inspired instantiations along with appropriate design guidelines (see e. g. [22]) remains untouched. A meaningful adaptation of the DIC pattern to specific circumstances of a problem is required anyway.

Due to the extension of the functionality, in particular the integration of an infochemical-specific diffusion coefficient as well as the possibility that different element types may react on the same given infochemical, the message overhead in such systems can be minimized whereas the performance of the system can be increased simultaneously, as proved in [23]. However, a multi-interpretation of the same information by different elements has to be used carefully, because a later change of the information may have serious effects on the overall coordination process.

Due to the generalization of the coordination model, the applicability of the mechanism is no longer limited to autonomous software entities or agents. Future applications may also involve autonomous real-world entities and cover a much wider set of problem domains in e. g. power and performance management, traffic management, manufacturing control, robotics, or even astronautics.

Although one expects that a design pattern should have been effectively used quite a few times in different situations to qualify as a design pattern, this is awkward in the case of DAC. Only recently AC solutions became more decentralized and less deterministic (cf. [31]), so the interest in engineering, exploiting, and controlling self-organizing emergent systems to achieve certain autonomic capabilities only starts out (see e. g. [32]). DAC systems are still very often ad hoc solutions, mostly designed and implemented from scratch. Consequently, the known uses of corresponding design patterns in general are sparse yet. The DIC pattern thus can be considered to be more at the revolutionary as opposed to the evolutionary scale of AC systems.

Acknowledgment

The authors would like to thank the reviewers for their useful comments.

References

[1] F. Zambonelli and H. V. D. Parunak, "Towards a paradigm change in computer science and software engineering: a synthesis," *The Knowledge Engineering Review*, vol. 18, no. 4, pp. 329–342, December 2003.

[2] R. Sterritt, M. Parashar, H. Tianfield, and R. Unland, "A concise introduction to autonomic computing," *Advanced Engineering Informatics*, vol. 19, no. 3, pp. 181–187, 2005.

[3] T. De Wolf and T. Holvoet, "A taxonomy for self-* properties in decentralised autonomic computing," in *Autonomic Computing: Concepts, Infrastructure, and Applications*, M. Parashar and S. Hariri, Eds. CRC Press, 2007, pp. 101–120.

[4] M. Jelasity, O. Babaoglu, and R. Laddaga, "Guest editors' introduction: Self-management through self-organization," *IEEE Intelligent Systems*, vol. 21, no. 2, pp. 8–9, March/April 2006.

[5] T. Lewis, "The elements and frontiers of insect communication," in *Insect Communication*, T. Lewis, Ed. London: Academic Press, 1984, pp. 1–27.

[6] M. Dicke and M. W. Sabelis, "Infochemical terminology: Based on cost-benefit analysis rather than origin of compounds?" *Functional Ecology*, vol. 2, no. 2, pp. 131–139, 1988.

[7] S. Brückner, "Return from the ant - synthetic ecosystems for manufacturing control," PhD thesis, Humboldt-Universität, Berlin, 2000.

[8] P. Valckenaers, K. Hadeli, B. S. Germain, P. Verstraete, and H. V. Brussel, "MAS coordination and control based on stigmergy," *Computers in Industry*, vol. 58, no. 7, pp. 621–629, 2007.

[9] H. Kasinger, J. Denzinger, and B. Bauer, "Digital semiochemical coordination," *Communications of SIWN*, vol. 4, pp. 133–139, June 2008.

[10] S. S. Duffey, "Arthropod allomones: chemical effronteries and antagonists," in *Proceedings of the 15th International Congress of Entomology, Washington, D.C., USA*, 1976, pp. 323–394.

[11] D. A. Nordlund, "Semiochemicals: A review of the terminology," in *Semiochemicals: Their Role in Pest Control*, D. A. Nordlund, R. L. Jones, and W. J. Lewis, Eds. New York: John Wiley & Sons, 1981, pp. 13–28.

[12] D. A. Nordlund and W. J. Lewis, "Terminology of chemical releasing stimuli in intraspecific and interspecific interactions," *Journal of Chemical Ecology*, vol. 2, no. 2, pp. 211–220, 1976.

[13] W. L. Brown, Jr., T. Eisner, and R. H. Whittaker, "Allomones and kairomones: Transpecific chemical messengers," *BioScience*, vol. 20, pp. 21–22, 1970.

[14] T. D. Wyatt, *Pheromones and Animal Behaviour: Communication by Smell and Taste*. Cambridge University Press, 2003.

[15] J. Ruther, T. Meiners, and J. L. M. Steidle, "Rich in phenomena-lacking in terms. a classification of kairomones," *Chemoecology*, vol. 12, no. 4, pp. 161–167, November 2002.

[16] W. H. Bossert and E. O. Wilson, "The analysis of olfactory communication among animals," *Journal of Theoretical Biology*, vol. 5, no. 3, pp. 443–469, 1963.

[17] R. L. Doty, "Odor-guided behavior in mammals," *Cellular and Molecular Life Sciences*, vol. 42, no. 3, pp. 257–271, 1986.

[18] G. Meszaros and J. Doble, "A pattern language for pattern writing," in *Pattern languages of program design 3*. Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc., 1997, pp. 529–574.

[19] E. Gamma, R. Helm, R. Johnsona, and J. Vlissides, *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley Professional, October 1994.

[20] S. Brückner and H. V. D. Parunak, "Multiple pheromones for improved guidance," in *Symposium on Advances in Enterprise Control, Minneapolis, USA*, 2000.

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

[22] H. Kasinger, B. Bauer, and J. Denzinger, "The meaning of semiochemicals to the design of self-organizing systems," in *Proceedings of SASO 2008, Venice, Italy*. IEEE Computer Society, 2008, pp. 139–148.

[23] H. Kasinger, J. Denzinger, and B. Bauer, "Decentralized coordination of homogeneous and heterogeneous agents by digital infochemicals," in *Proceedings of SAC 2009, Honolulu, USA*. ACM Press, 2009, pp. 1223–1224.

[24] M. Yokoo and Y. Kitamura, "Multiagent real-time A* with selection: Introducing competition in cooperative search," in *Proceedings of ICMAS 1996, Kyoto, Japan*. AAAI Press, 1996, pp. 409–416.

[25] M. Mamei and F. Zambonelli, "Co-fields: A physically inspired approach to motion coordination," *IEEE Pervasive Computing*, vol. 3, no. 2, pp. 52–61, 2004.

[26] D. Deugo, M. Weiss, and E. Kendall, "Reusable patterns for agent coordination," in *Coordination of Internet agents: models, technologies, and applications*, A. Omicini, F. Zambonelli, M. Klusch, and R. Tolksdorf, Eds. Springer-Verlag, 2001, pp. 347–368.

[27] L. Gardelli, M. Viroli, and A. Omicini, "Design patterns for self-organizing multiagent systems," in *Proceedings of CEEMAS 2007, Leipzig, Germany*, ser. LNCS, vol. 4696. Springer, 2007, pp. 123–132.

[28] 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, pp. 26–66, September 2006.

[29] T. De Wolf and T. Holvoet, "A catalogue of decentralised coordination mechanisms for designing self-organising emergent applications," Department of Computer Science, K.U.Leuven, Leuven, Belgium, Report CW 458, 2006.

[30] M. Güdemann, F. Nafz, F. Ortmeier, H. Seebach, and W. Reif, "A specification and construction paradigm for organic computing systems," in *Proccedings of SASO 2008, Venice, Italy*. IEEE Computer Society, 2008.

[31] M. C. Huebscher and J. A. McCann, "A survey of autonomic computing—degrees, models, and applications," *ACM Computing Surveys*, vol. 40, no. 3, pp. 1–28, 2008.

[32] R. Anthony, A. Butler, and M. Ibrahim, "Exploiting emergence in autonomic systems," in *Autonomic Computing: Concepts, Infrastructure, and Applications*, M. Parashar and S. Hariri, Eds. CRC Press, 2007, pp. 121–148.