# Cognitive Stigmergy: Towards a Framework Based on Agents and Artifacts

Alessandro Ricci, Andrea Omicini, Mirko Viroli, Luca Gardelli, and Enrico Oliva

Alma Mater Studiorum—Università di Bologna
via Venezia 52, 47023 Cesena, Italy
a.ricci@unibo.it, andrea.omicini@unibo.it, mirko.viroli@deis.unibo.it, luca.gardelli@unibo.it, enrico.oliva@unibo.it

**Abstract.** *Stigmergy* has been adopted in MAS (multi-agent systems) and in other fields as a technique for realising forms of emergent coordination in societies composed by a large amount of ant-like, non-rational agents. In this paper we discuss a conceptual (and engineering) framework for exploring the use of stigmergy in the context of societies composed by cognitive / rational agents, as a means for supporting high-level, knowledge-based social activities.multi-agent We refer to this kind of stigmergy as *cognitive stigmergy*. Cognitive stigmergy is based on the use of *artifacts* as tools populating and structuring the agent working environment, and which agents perceive, share and rationally use for their individual goals. Artifacts are environment abstractions that mediate agent interaction and enable emergent coordination: as such, they can be used to encapsulate and enact the stigmergic mechanisms and the shared knowledge upon which emergent coordination processes are based.

In this paper, we start exploring this scenario introducing an agent-based framework for cognitive stigmergy based on artifacts. After discussing the main conceptual issues—the notion of cognitive stigmergy and the role of artifacts—, we sketch an abstract architecture for cognitive stigmergy, and outline its implementation upon the TuCSoN agent coordination infrastructure.

## 1 Introduction

In the last years, the study of *stigmergy* has influenced a number of different research fields, including MAS (multi-agent systems). In general, and in MAS research in particular, stigmergy is mostly used as the source of simple yet effective coordination metaphors and mechanisms, to be exploited for building robust and reliable systems in unpredictable settings. The main source of inspiration is obviously represented by the studies on insects and ant societies [1], which have led to a basic meta-model based on (ant-like) *simple* and homogeneous *agents* possessing no relevant cognitive abilities. Such agents interact with each other through *local* modifications to the *environment*, eventually originating *global* structures and behaviours [2].

D. Weyns, H.V.D. Parunak, and F. Michel (Eds.): E4MAS 2006, LNAI 4389, pp. 124–140, 2007.

While this stream of research has produced a number of very interesting approaches in MAS (see [3,4] among the many others), it has also brought on two main biases: *(i)* the agent model is very simple—ant-like agents do not exploit any cognitive ability of theirs—and *(ii)* the environment model is often quite elementary, featuring pheromone-like signs/signals with simple mechanisms for diffusion, aggregation and evaporation—at most extended to force fields [5].

By contrast, a number of relevant works in the field of cognitive sciences put in evidence how stigmergy—as the social mechanism of coordination based on interaction through local modifications to a shared environment—is a fundamental coordination mechanism also e.g. in the context of human societies and organisations [6,7]. In this context:

- modifications to the environment are often amenable of an interpretation in the context of a shared, conventional system of signs;
- the interacting agents feature cognitive abilities that can be used in the stigmergy-based interaction;
- the environment is articulated, and typically composed of *artifacts*, which build up the social workspace, or field of work;
- artifacts can be suitably engineered in order to process information cognitively shared by agents.

Starting from this consideration, in this paper we start exploring what we call *cognitive stigmergy*, that is, the generalisation of stigmergic coordination to enable social activities of cognitive agents. We hence consider a wider notion of agency, which includes high-level knowledge representation capabilities, explicit representation of agent goals, inferential / planning / deliberation abilities, and so on. Our goal is to promote the idea that the general notion of stigmergy can suggest new models for interaction, coordination, and organization within MAS including cognitive agents.

Following the approaches in cognitive sciences, and Computer Supported Cooperative Work (CSCW) in particular [7], we stick to a meta-model featuring *artifacts*, which are instruments and tools that make up and constitute the agent environment, and which agents can select and use for their own purposes. Artifacts are environment abstractions [8] that are *(i)* the subject of cognitive agent activity, *(ii)* the enabler and rulers of agent interaction, and *(iii)* the natural *loci* for cognitive stigmergy processes.

The main aim of this line of research is to propose a reference conceptual framework for cognitive stigmergy, which can also serve as a basis for engineering practical experimentation in the field of MAS. We identify at least three different objectives:

- from a scientific-synthetic viewpoint, we aim at constructing a model for stigmergic coordination going beyond ant-like metaphors: agents are not only ants, and signs for stigmergy are not only pheromones. The cognitive abilities of agents, and the articulation of the environment through artifacts are the essential ingredients to move from stigmergy to cognitive stigmergy;

– from a scientific-analytic viewpoint, the proposed framework should be combined with agent-based and simulation technologies in order to provide predictive models for systems based on cognitive stigmergy, such as human organisations and societies;
– from an engineering viewpoint, we aim at devising out a framework for the construction of MAS stigmergic mechanisms to coordinate complex activities of any sort within articulated operating contexts. Coordinated MAS behaviour should then emerge as the result of both cognitive and non-cognitive activities by the agents, and by their local interaction mediated by suitably engineered artifacts.

In this paper, in particular, we focus on the first issue, and also sketch a possible approach to the third one. In particular, in Sect. 2 we recapitulate some of the multidisciplinary pillars that a theory of cognitive stigmergy should be based upon, then in Sect. 3 we first sketch our conceptual background. In Sect. 4 we provide some remarkable examples of artifacts for cognitive stigmergy, and finally, in Sect. 5, we shortly outline a possible methodological and technological framework for engineering MAS with cognitive stigmergy based on the TuCSoN infrastructure for MAS coordination, adopting tuple centres as artifacts. Conclusion and future work are provided in Sect. 6.

## 2 Trans-disciplinary Background

The notions of stigmergy, interaction through artifacts, and the many sorts of structures and behaviours that emerge in complex societies, are strictly interrelated concepts that have been the subjects of investigation in a multiplicity of heterogeneous research areas. Adopting a multi-disciplinary view is then rather mandatory—but in some sense quite usual in the field of MAS, given the generality and expressive power of abstractions like agent, society and environment.

Even more, a trans-disciplinary approach is potentially very fertile: taking examples and definitions of stigmergic coordination from both ethology and social sciences, bringing them to the MAS field, and building a general model for cognitive stigmergy, is a fascinating perspective indeed, which could induce novel interpretations and metaphors.

### 2.1 Definition and (Mis)Use of the Notion of Stigmergy

The original notion of *stigmergy* was introduced by Grassé in the late 50s while studying and trying to explain the behaviour of social insects. In its first formulation, stigmergy was defined as a "class of mechanisms that mediate animal-animal interactions which is fundamental for achieving emergent forms of coordinated behaviour at the society level". Originally, the concept of stigmergy was used to build up a coherent explanation of the so-called *coordination paradox* between the individual and the societal level: on the one hand, groups of social insects seem to be cooperating in an organised, coordinated way; on the other hand, each individual seems to be working as if it were alone, neither interacting with others nor involved in any collective behaviour [1]. The explanation to the coordination paradox provided by stigmergy is that insects interact *indirectly*: each insect (ants, bees, termites) affects the behaviour of other insects by indirect communication through the use of the environment, which is made of objects and artifacts such as material for the nest, or chemical traces.

From the original formulation of the notion of stigmergy, the key-role of the *environment* firstly emerges, which acts not merely as a passive landscape against which all the interactions occur, but rather as a mediator and a ruler of interactions. Secondly, stigmergic interaction is always *mediated*: it occurs locally to the interacting entity, and directly affects a portion of the environment. Finally, the environment is seen as confined / bounded to well-defined elements, such as a pheromone or a chunk of material for nest construction: so, objects, tools, instruments, and *artifacts* encapsulate the logic of local interaction, and are therefore prominent actors in the process of stigmergic coordination.

In the context of computer science, in general, and in the field of MAS, in particular, stigmergy has been widely used as a technique for complex problem solving, as well (more recently) as an approach to the design and development of systems [3,4]. This of course is mainly motivated by the need for system reliability and robustness in complex and unpredictable environments, which could in principle be addressed by mechanisms for self-organisation like stigmergy. On the other hand, however, ants and pheromones provide for a simple, easy-to-reproduce mechanism for stigmergy: as a consequence, stigmergy is often implicitly reduced to an ant-like phenomenon. This is not to say that ant-based mechanisms, models and technologies do not obtain significant outcomes: instead, a large number of remarkable results were indeed achieved in computer science [9], robotics [10], and MAS [11,12].

What is missing, we believe, is instead a wide and coherent view on stigmergy that while sticking to the general principles of the original Grassé's definition of stigmergy, would also account for the facts that *(i)* agents are possibly cognitive entities—agents are not always ant-like entities—, and *(ii)* environment is possibly more articulated than a mere pheromone container, but is rather composed of suitably engineered artifacts. This is exactly what we find in research from cognitive sciences.

### 2.2 Artifacts, Workspaces, and Stigmergic Coordination

Forms of indirect, mediated interaction are pervasive in complex systems, in particular in contexts where systems take the form of structured societies. In such contexts, in order to scale with activity complexity, sorts of *mediating artifacts* are shared and exploited to enable and ease interaction among the components. Mediating artifacts of different sorts can be easily identified in human society, which are designed and exploited to support coordination in social activities, and in particular in the context of cooperative work. Well-known examples are blackboards, form sheets, post-it notes, and archival tools. Mediation is well-focused by some theories such as Activity Theory [13] and Distributed Cognition [14] adopted in the context of CSCW and HCI (Human Computer Interaction), which explore how the environment can be shaped in terms of mediating artifacts and in order to better support cooperative work among individuals.

Among the most interesting references, the work by Susi [6] represents one of the most coherent efforts toward a theory of artifacts in social interactions, putting together HCI and cognitive sciences. From this work a picture clearly emerges where the activities within complex (human) organisations occur in the context of structured *workspaces*: workspaces are made of artifacts, which are subjects of the human cognitive activity, work as mediators of interaction, and encapsulate coordination functions. The notion of workspace (media spaces, virtual rooms, virtual workspaces in CSCW [7]) clearly exemplifies the idea of a non-trivial, non-passive, articulated environment—where artifacts represent the environment *articulation*. Also, artifacts are mostly a cognitive concept: intelligent activity is required to enact them, make them work, and understand their meaning as coordinating entities—as happens e.g. with triggers, placeholders or entry-points [6].

From a psychologist perspective, the work by Castelfranchi [15,16] points out another key issue: independently of the intentions motivating activities on artifacts (intention to communicate or not, for instance), any behaviour in a workspace is anyway amenable to an interpretation by the observers, which could bring meaningful information, and affect their subsequent behaviour. For instance, when Bianca takes one of the two glasses on the table to drink, she is not explicitly telling Bernie on the other side of the table "take the other glass"—she is just taking her glass plain and simple. However, Bernie is going to interpret Bianca's action on the shared workspace (the table with the glasses) as an implicit communication from her, and take the other glass anyway.

This is also quite apparent in some of the most well-known examples of shared knowledge-based human-oriented artifacts, such as platforms for cooperative work like Wiki (and the Wikipedia [17]), and even platforms for e-commerce (which are also huge sources of information) like Amazon [18]. For instance, one of the most obvious but effective ways of interaction in the Wikipedia is by annotating a page. When looked from an ant-like perspectives, this resembles the release of a pheromone on a shared environment articulated in pages: more (pheromones-)annotations "deposited" on the same page may "aggregate" to indicate a higher level of interest, then attract the interest of other (ants-)readers.

However, the cognitive nature of both page artifacts and annotations, along with the cognitive abilities of human agents, allows for less trivial forms of "stigmergic" processes. For instance, ranking a page based on its perceived utility enables more articulated forms of aggregation (like global average ranking), and may consequently lead to different evolution histories of the whole knowledge base.

Even mediated implicit communication is easy to be observed, for instance in Amazon. For instance, Lilo does not buy book A and then book B to say anything to anyone—just to read them both. However, logging and aggregating this sort of actions allow Amazon to say Stitch, who is buying book B, that "customers who bought this book also bought book A"—which quite often turns to be very informative in practise, and tends to influence both the individual and the overall behaviours. In other terms, individual cognitive actions (read a book presentation, decide to buy that book) in a local context (the view from the browser) upon a cognitive artifact (the purchase page) change the state of the environment (the Amazon portal) and then the behaviour of other individuals, such that in the overall the global behaviour of the system is affected.

Evidence of stigmergic processes involving cognitive features could not be clearer around us—in the scientific arena, as well as in our everyday life. The point is now how to use this evidence in MAS, so that both traditional results from the ant-biased interpretation of stigmergy and the cognitive interpretation drawn from CSCW, HCI, Activity Theory and cognitive sciences could be subsumed, coherently modelled, and then be used in order to build complex, robust and intelligent MAS.

## 3 Cognitive Stigmergy in MAS

Our objective in this work is the investigation of stigmergy principles in the context of cognitive MAS, i.e. societies of goal/task-oriented/driven agents interacting at the *cognitive level*. Such agents are therefore not necessarily simple and reactive ones, as in the ant case, but can typically be rational, heterogeneous, adaptive, and capable of learning. We adopt the term *cognitive stigmergy* to denote this approach, so as to remark the differences with respect to existing approaches to stigmergy in MAS, which are typically based on societies of agents whose capabilities and behaviour resemble those of insect-like entities.

As in the case of classic stigmergy, the environment is a central concept for cognitive stigmergy, as an enabler and mediator of the agent work and interaction. The general picture—reflecting a certain complexity in the corresponding engineering of applications—is given by a (possibly open) set of agents with their own specific tasks and goals, which perform their individual as well as social activities in the same working environment, sharing the same *field of work*. The interaction among the agents is indirect, uncoupled in time and space. From a modelling and engineering point of view, it is natural to model such a working environment as a first-class entity: agents are aware *(i)* of their field of work, *(ii)* of it being shared with other agents, and *(iii)* of its functionality, and of the opportunities to use it so as to achieve their objectives (*affordance* of the environment). Such opportunities are exploited by properly *using* the working environment, that is, by executing the operations that the environment makes available to agents and by observing its state.

As in the case of classic stigmergy, a main point here is that the environment is not a mere passive "container", but it embeds mechanisms and (reactive) processes that promote the emergence of local and global coordinated behaviours. Not only it has a state that can be observed and modified by agents, but also it encapsulates some laws that can be triggered by agent actions (or, by events such as a change in location, or the passing of time) and alter the environment state independently of agent intentions.

Under a cognitive perspective, the working environment in cognitive stigmergy can be framed as a set of shared stateful tools providing specific functionalities that are useful for agents performing their *individual* work. At the same time, such tools are designed to be collectively shared and used by agents, and are generally implemented so as to effectively and efficiently support their shared functionalities, thus largely impacting on the *social* level.

In the rest of this section, we focus on the answer to the following key question: how could this kind of working environment be modelled as a first-class entity in MAS? To this end, in the following we elaborate on the notion of *artifact* as a means to explicitly and directly design and build such a working environment.

### 3.1 Exploiting the Notion of Artifact

The notion of artifact (and the related conceptual framework) has been introduced recently in MAS as a first-class abstraction representing tools or objects (devices) that agents can either individually or collectively *use* to support their activities, and that can be designed to encapsulate and provide different kinds of functionalities or services [19,20]. If agents are meant to be first-class abstractions to model goal/task-oriented/driven pro-active entities, artifacts are those entities modelling systems (or parts of a system) that are better characterised as resources or tools *used* by agents for their own aims. In particular, and unlike agents, artifacts neither have internal goals, nor do they exhibit a pro-active behaviour; instead, they simply provide agents with some kind of functionality they could be suitably exploit, typically in the form of a service—in other words, while agents *communicate with* other agents, agents *use* artifacts.

According to the abstract model defined [19], artifacts in cognitive MAS can be characterised by (see Fig.1): a *function*, as its intended purpose, i.e. the purpose established by the designer / programmer / builder of the artifact—in other words the intended functionalities the artifact is meant to provide; a *usage interface*, as the set of the operations that agents can invoke to use the artifact and exploit its functionality; some kind of *operating instructions*, as a description of how to use the artifact to access its functionality; a *structure* and *behaviour*, concerning the internal aspects of the artifact, that is, how the artifact is structured and implemented in order to provide its function.

Unlike agents, artifacts are not meant to be autonomous or exhibit a pro-active behaviour, neither are they expected to have social capabilities. Among the main properties that are useful according to the artifact purpose and nature, one could list [20]: *(i) inspectability* and *controllability*, i.e. the capability of observing and controlling artifact structure, state and behaviour at run-time, and of supporting their on-line management, in terms of diagnosing, debugging, testing; *(ii) malleability* (or, *forgeability*), i.e. the capability of changing / adapting artifact function at run-time (on-the-fly) according to new requirements or unpredictable events occurring in the open environment,[1] and *(iii) linkability*,

[1] Such adaptation is not meant to be realised autonomously by the artifacts themselves, but by MAS agents & engineers acting on the artifacts. Mechanisms for non-intentional self-adaptation of artifacts are not excluded a priori, but they are not directly related with malleability, and are not necessarily a desirable property of artifacts.

i.e. the capability of composing distinct artifacts at run-time as a form of composition, as a means to scale up with complexity of the function provided, and to support dynamic reuse. It is worth to be remarked that these artifact features are not agent features: typically, agents are not inspectable, do not provide means for malleability, do not provide operations for their change, and do not compose with each other through operational links.

Also, artifacts can have a *spatial extension*, i.e. given a MAS with a topology, the same artifact could cover different nodes: in other words, a single artifact can be both conceptually and physically distributed. For instance, a blackboard artifact can cover different Internet nodes, where agents may use it by exploiting a local interface. Technically, agents could be distributed, too—for instance, having the knowledge base in some node and the deliberation engine hosted by some other node: most often, an agent is situated within a specific location, at least by considering the agent models and architectures that are most diffused (an example is the FIPA model).

Given this notion of artifact, we can reformulate the context of cognitive stigmergy in terms of a set of agents sharing a set of artifacts representing their working environment. This set can be split along two different levels:

– a *domain level*, with artifacts that represent the target of the agent work, or an *objectification* of such a target.
– a *tool level*, with artifacts that represent the working tools which can help agents in doing their work.

Our objective is to instrument the tool level with a web of linked artifacts which can be used to improve the work of the collectivity of agents sharing the same working environment. At the systemic level, these artifacts are meant to be used both to improve the knowledge about the *practises* in using the artifacts at the domain level, and to possibly support *social construction* and *evolution / adaptation* of such artifacts, toward directions that are useful for the collectivity of agents in the overall. In order to support this functionality, the artifacts belonging to the tool level should encapsulate stigmergic mechanisms partially similar to the mechanisms found in ant-based systems and pheromone infrastructures: such mechanisms are described in Sect. 4.

### 3.2 Re-framing the Notion of Locality: Workspaces

In classic approaches to stigmergy the notion of topology (and related notion of locality) is mostly physical, defining from the viewpoint of agents—which are typically mobile—the portion of the environment which can be directly affected by their actions or can be perceived. In the case of cognitive stigmergy, this crucial notion could be formulated in a natural way with the notion of *workspace*, as the set of artifacts directly available (usable) for an agent. Workspaces can cross each other sharing agents and artifacts, can be nested, and so on: in synthesis they are a way to define the topology in a rigorous way.

Actually, the topology induced by this characterisation is more abstract and could be articulated along different dimensions. An important one is for instance organisation: the same artifacts could be accessible and usable in different ways according to the roles and permissions assigned to agents by the organisation they belong to.

It is worth noting that the nature and functionality of the artifacts could bring in situation where—to some extent—the principle of physical locality is violated. This is evident in our society, where artifacts (for humans) such as cell phones, televisions, or the Internet itself can be used to observe and interact in a direct way with entities—e.g. humans—located at completely different places of the world. Conceptually, the action of an agent executing an operation on an artifact of its workspace (its locality) can have "instant" effects on a completely different workspace. This happens because artifacts can be either shared among workspaces or linked together across workspaces.[2] Actually, the principle of locality still holds, since agents can only use the artifacts belonging to their workspaces.

### 3.3 From Pheromones to Annotations

In every stigmergic system, the effects of agent actions on the environment are understood as signs. Once created, signs persist independently of their creator and are observable by the other agents, and are subject of manipulation by the environment itself according to the laws which characterise the stigmergic processes—e.g. diffusion and evaporation. Differently from pheromones in the case of ant-based stigmergy, in the case of cognitive stigmergy signs typically hold a *symbolic* value, embodying information of some sort, with a formal or informal semantics, referring to some ontology. We refer to such a symbolic information in cognitive stigmergy as *annotations*.

Coming back to the two levels previously introduced, annotations are useful first of all for expressing some kind of comment or knowledge about the artifacts (and about the practise of use of artifacts) belonging to the domain level, which are targets of the agent work. Then, annotations are useful to objectify also comments or reflections that do not concern a specific artifact, but more generally a working practise, and could possibly refer to multiple artifacts. Finally, annotations can be used for expressing a comment on the annotations themselves, typically about their utility, effectiveness, and so on.

Knowledge provided by annotation is both explicit—the content of the annotation—and implicit—the "shape" and the context of the annotation, including for instance the possible intention of the agent or group of agents that created the annotation. The concept of shape for annotations could be considered as analogous to the concept of force in the case of speech acts: it modulates annotation content according to the information that could be of some use when reasoning on annotations.

Some of the artifacts defining working spaces in cognitive stigmergy are possibly devoted to the management of annotations, providing agents with operations

[2] One should remember that, from a physical point of view, an artifact could be distributed across multiple sites.

for creating and observing annotations, and embedding mechanisms for automatically manipulating annotations (with forms of aggregation, diffusion, selection, ordering) in order to implement the functionality required for cognitive stigmergy. Accordingly, we deal with two basic kinds of annotations:

– annotations explicitly and intentionally created by agents. These include, for instance, agent feedback (evaluation) about a specific artifact belonging to the domain level; agent feedback about a specific annotation on one such artifact; agent annotations about a set of artifacts, or a usage practise during a working session.
– annotations automatically created by the artifacts supporting their working activities. Examples include annotations reporting about how much an artifact has been used, how many agents exploited an artifact for their purposes, how many agents considered an annotation as useful for their purposes, which other artifacts have been used (and how) by agents using a given artifact.

## 4 Artifacts for Cognitive Stigmergy

Generally speaking, artifacts in cognitive stigmergy should first of all promote *awareness*, that is, making agents seamlessly aware of the work and practises of other agents, which could in turn be effective to drive or improve their own activities. Awareness is a key aspect to support emergent forms of coordination, where there is no pre-established plan defining exactly which are the dependencies and interactions among ongoing activities (involving agents and artifacts) and how to manage them—instead, such a plan emerges along with the activities themselves.

A simple but effective example of stigmergic mechanism promoting awareness can be found—for instance—in Amazon: a user consulting the page of a book is provided with a list of other books, bought by users that purchased the same book. This kind of mechanism in Wikipedia could be realised through a page annotation of the kind: "people consulting this page have also consulted pages X, Y, Z". In our framework, such a mechanism can be generalised by supporting the automatic creation of annotations on artifacts of the domain level, reporting information about which other artifacts have been used by agents using the same artifact.

In the remainder of this section we describe a basic set of artifacts which could constitute a simple example of an architecture supporting some form of awareness and other features characterising cognitive stigmergy. On the background of this architecture there is the notion of *working session*, as a temporal scope for an agent activities. An agent starts a working session with an objective in mind, which is supposed to persist for all the duration of the session. Knowing the (either explicit or implicit) objective of an agent during a working session is important to provide a context—in terms of the problem to be solved, the goal to be achieved, the task to be executed—to the annotations (evaluations, comments, ...) made by the agents, and to the practise of the agents using the artifact of the domain level. For instance, in Wikipedia, agent feedbacks about the utility of a page would be better understood and evaluated by taking into account the problem the agent is facing (i.e. what it is looking for).

### 4.1 Promoting Awareness: Dashboards, Logs, Diaries and Note-Boards

A first and necessary step toward awareness is to keep track of both the actions and the annotations made by individual agents during a working session. For this purpose, we identify three basic kinds of artifacts, corresponding to three different kinds of functionalities: dashboards, logs and diaries (see Fig.1):

- A *dashboard* provides the functionalities of a panel (interface) used to *focus* on a specific artifact (or a set of artifacts belonging to the domain level) to interact with the artifact and to take / observe / manage annotations. The concept of focus aims at representing the intention of using an artifact.
- A *log* is used to keep track of events, providing operation for their inspections and ordering.
- A *diary* is an artifact used to keep track of annotations intentionally made by an agent. The diary typically keeps the annotations organised by working sessions.

The dashboard is linked to the log so as to trace all the operations executed by the agent during a working session. Actually, the log of the operations executed by an agent is interesting also for analysing paths as sequences of executed operations, which can be important to identify and evaluate practises in using one or a set of artifacts belonging to the domain level. The stigmergic system could be instrumented so as to make agents aware of such practises and of the possibility to provide an evaluation, so as to augment the common awareness about good (and bad) practises.

Besides tracing individual agent actions and annotations, it is necessary to introduce artifacts that actually make it possible to effectively share annotations about specific artifacts of the domain level. For this purpose, the *note-board* artifact is introduced. A note-board is useful for keeping and managing all the annotations about a specific artifact (or set of artifacts) of the domain level. For instance, in the Wikipedia example we could have a note-board for each page (or group of pages) of the system.

A note-board is meant to contain both the annotations intentionally made by agents on the specific artifact, and the annotations automatically created by the artifact itself or by other artifacts by virtue of the stigmergic mechanisms and processes. A simple example can be an annotation reporting how many different kinds of agents used a specific artifact. Such a functionality can be obtained by properly combining the dashboard and note-board: for instance, each time a dashboard focuses for the first time on an artifact $X$, an annotation about this fact can be made on the note-board of artifact $X$. The note-board can then transform the set of such annotations in a single annotation (by means of *aggregation* mechanisms, described in next subsection), reporting the number of agents that used the artifact. Another

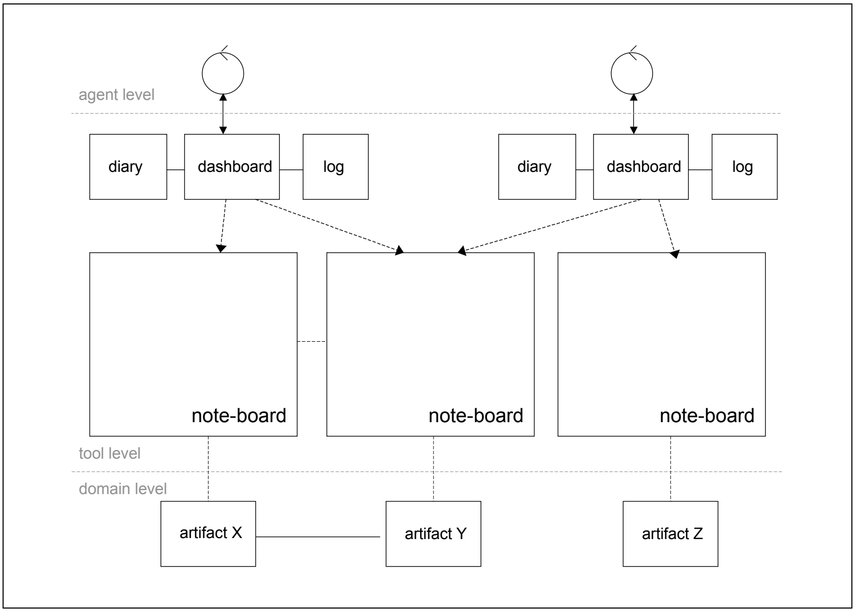


**Fig. 1.** An abstract representation of an architecture for cognitive stigmergy based on dashboard, log, diary and note-board artifacts

example could be an annotation reporting information about which other artifacts have been used by agents using this artifact. In this case, when the focus of an agent switches from an artifact $X$ of the domain level to an artifact $Y$, a suitable annotation can be automatically created on the note-board of artifact $X$ reporting the fact that an agent using artifact $X$ has then used artifact $Y$, and on the note-board of artifact $Y$ with the dual information.

The above examples suggest how the combined use of artifacts with relatively simple functionalities could be effective enough to improve agent awareness about their working practises. Functionalities provided by artifacts are instrumental to realise the forms of reinforcement and positive feedback that typically characterise stigmergic systems as dynamic non-linear systems: the more agents are aware of the usefulness of an artifact, the more they use it, augmenting the overall awareness about the utility of the artifact.

### 4.2 Some Basic Stigmergic Mechanisms

Analogously to the case of ant-based stigmergy, also in the context of cognitive stigmergy it is possible to identify some basic and recurrent mechanisms which can be embedded within artifacts in order to support stigmergic processes:

**Diffusion** — Diffusion is one of the basic mechanism in ant-based stigmergy. In the context of cognitive stigmergy, an analogous principle could be exploited to improve awareness, according to the simple rule that annotations that concern a specific artifact could be useful also for artifacts that are directly linked to that artifact according to some kind of relation explicitly established at the domain level. For instance, in the case of Wikipedia, annotations concerning a specific page could be useful also for pages that are directly linked to or directly link such a page.

Note-boards could be designed to suitably support diffusion capabilities: annotations intentionally made by agents about an artifact could be automatically propagated from the related note-board to all the note-boards of the linked artifacts. Then, among the information that gives shape to an annotation, a *diffusion level* could also be included, indicating whether the annotation has been made directly by an agent or it has been propagated from other artifacts. Different kinds of diffusion policies are possible: for instance, note-boards could support either diffusion of direct annotations only, or propagation of annotations, too, possibly specifying a sort of propagation radius in terms on maximum diffusion level.

**Aggregation** — In our framework, the aggregation mechanism accounts for automatically transforming a set of annotations—related by some criteria—into a single annotation, typically containing an explicit information describing the aggregation in the overall (for instance, a quantity). Note-boards have the fundamental role of aggregators of the annotations concerning a specific artifact of the domain level. For instance, note-boards could automatically aggregate annotations containing agents' feedback (evaluation) on an artifact or on an annotation made on the artifact.

**Selection and Ordering** — Annotations may have a different relevance according to the different kinds of criteria / dimensions, which can be either subjective or objective. Consequently, such annotations could be automatically ordered by artifacts managing them in order to reflect their relevance. An example of ordering criteria is freshness, measuring relevance of an annotation according to its age. Another one is pertinence, measuring the relevance of a propagated annotation according to its diffusion level, as defined previously. A selection mechanism accounts for keeping and making available only a limited set of annotations—typically the most relevant ones according to the selected criteria / dimensions. Selection is often combined with ordering. To some extent, selection and ordering mechanisms could be considered as a generalisation—in the context of cognitive stigmergy—of the evaporation mechanism, as found in ant-based system. Also *dissipation*—a frequent mechanism in stigmergy system—could be considered as a specific case of selection, where all annotations not selected according some criteria are forgotten.

Actually diffusion, aggregation, selection and ordering are general kinds of mechanisms which can be considered as useful for a wide range of artifacts. In the examples we mainly consider note-boards, however it is easy to identify their utility also in diaries, where annotations are typically organised (aggregated) according to working sessions, ordered according to temporal criteria, and possibly diffused to note-boards, in case they concern specific artifacts.

## 5 Building MAS with Cognitive Stigmergy

### 5.1 Toward an Agent Infrastructure for Cognitive Stigmergy

As mentioned in Sect. 1, the conceptual framework of cognitive stigmergy is meant to be useful both for modelling / simulating complex social systems—so as to analyse emergent social behaviours of societies evolving in some specific workspaces—and for engineering complex agent applications, aiming at achieving some sort of fruitful social behaviour in spite of the independent working activities of the individual agents and / or the absence of a global coordination plan to follow. In both cases, in particular for the latter one, it is of foremost importance to have models / infrastructures that make it possible to represent in the most direct and seamless way the main concepts of the framework, in particular artifacts of the kind discussed in the paper. Accordingly, such a middleware would provide a support for cognitive stigmergy as a service, which MAS applications could customise and exploit according to the need.

### 5.2 An Example: TuCSoN as a Middleware for Cognitive Stigmergy

As an example, TuCSoN[3] coordination infrastructure [21] can be used as a middleware to experiment cognitive stigmergy, since it provides on the one side a direct support for cognitive and generative communication, based on the generation and consumption of *tuples* as kind of annotations; on the other side, it provides a natural way to model artifacts as first-class abstractions, with the possibility to define their specific behaviour.

TuCSoN provides *tuple centres* as first-class abstractions that agents can use to support their communication and coordination. Technically, tuple centres are *programmable tuple spaces*—sort of reactive blackboards that agents access associatively by writing, reading, and consuming *tuples*—ordered collections of heterogeneous information chunks—via simple communication primitives (*out*, *rd*, *in*, *inp*, *rdp*) [22]. While the behaviour of a tuple space in response to communication events is fixed, the behaviour of a tuple centre can be tailored to the application needs by defining a set of specification tuples expressed in the ReSpecT language, which define how a tuple centre should react to incoming / outgoing communication events. Basically, ReSpecT primitives make it possible to manipulate the tuples inside the tuple centre and also to establish a link between the tuple centre with other tuple centres—for instance making it possible to insert tuples in other tuple centres directly via reactions. ReSpecT is Turing-equivalent [23], so in principle any kind of tuple manipulation is possible. From the topology point of view, tuple centres are collected in TuCSoN nodes, distributed over the network, organised into articulated domains. A node can contain any number of tuple centres, denoted by a specific name: the full name of a tuple centre consists in its local name plus the Internet address of the hosting TuCSoN node.

[3] The TuCSoN technology is available as an open source project at the TuCSoN web site `http://tucson.sourceforge.net`

Then, it is natural to use TuCSoN tuple centres as general-purpose artifacts that can be programmed according to the need, in order to provide specific functionalities. Annotations can be easily implemented as logic tuples. Interaction between agents and artifacts could be modelled on top of tuple centre basic communication primitives, by choosing a specific format for both tuples and tuple templates. Artifact behaviour could be implemented as a set of ReSpecT reactions implementing the basic stigmergic mechanisms discussed in the paper, by virtue of the Turing-equivalence of ReSpecT. In particular:

- aggregation mechanisms can be implemented as ReSpecT reactions consuming a specific set of tuples and producing a single tuple, according to some specific criteria;
- selection and ordering mechanisms can be implemented as reactions that create and maintain tuples containing a list of other tuples, imposing an order among them;
- both diffusion and artifact composition can be implemented by using the linkability property of tuple centres, with reactions that propagate tuples from a tuple centre to the others.

The basic set of artifacts identified in previous section—dashboards, diaries, logs and note-boards—can then be implemented as suitably programmed ReSpecT tuple centres. There is no space enough here to provide further details about design and implementation of such artifacts—which however are not overwhelmingly complex, indeed. Such details are likely to be discussed in a forthcoming work along with an evaluation of the system performance in supporting cognitive stigmergy.

## 6 Conclusion and Future Works

Stigmergy is a simple and powerful mechanism around which complex coordination patterns can be organised and built. Despite the generality of the original definition by Grassé [1], the full potential of stigmergy has yet to be developed in the area of MAS, as both a modelling and a constructive principle for complex agent-based systems.

In this paper, we proposed an extended interpretation of stigmergy, which we termed as *cognitive stigmergy*, which could on the one hand preserve the benefits of the ant-biased acceptation usually adopted in the MAS field, on the other hand promote the full exploitation of the cognitive abilities of agents and of the environment articulation in artifacts in the stigmergic process. After summarising our main sources of inspiration from a number of different research areas and technology contexts, we proposed a conceptual framework for cognitive stigmergy in MAS, and then sketched a possible engineering approach based on the TuCSoN infrastructure for agent coordination, using tuple centres as artifacts.

Future work will be devoted to further explore both the theoretical framework and the practical perspectives opened by this paper, focusing in particular on scenarios like e-learning systems, and implicit organisations based on overhearing / over-sensing.

## References


1. Grassé, P.P.: La reconstruction du nid et les coordinations inter-individuelles chez bellicositermes natalensis et cubitermes sp. la theorie de la stigmergie: essai d'interpretation des termites constructeurs. Insectes Sociaux **6** (1959) 41–83
2. Di Marzo Serugendo, G., Foukia, N., Hassas, S., Karageorgos, A., Kouadri Mostéfaoui, S., Rana, O.F., Ulieru, M., Valckenaers, P., Van Aart, C.: Self-organisation: Paradigms and applications. [24] 1–19
3. Parunak, H.V.D., Brueckner, S., Sauter, J.: Digital pheromone mechanisms for coordination of unmanned vehicles. In Castelfranchi, C., Johnson, W.L., eds.: 1st International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2002). Volume 1., Bologna, Italy, ACM (2002) 449–450
4. Valckenears, P., Van Brussel, H., Kollingbaum, M., Bochmann, O.: Multi-agent coordination and control using stigmergy applied to manufacturing control. In Luck, M., Mařík, V., Štěpánková, O., Trappl, R., eds.: Multi-Agent Systems and Applications. Volume 2086 of Lecture Notes in Computer Science., Springer (2001) 317–334 9th ECCAI Advanced Course ACAI 2001 and Agent Link's 3rd European Agent Systems Summer School (EASSS 2001), Prague, Czech Republic, 2-13 July 2001, Selected Tutorial Papers.
5. Mamei, M., Zambonelli, F., Leonardi, L.: Co-Fields: Towards a unifying approach to the engineering of swarm intelligent systems. In Petta, P., Tolksdorf, R., Zambonelli, F., eds.: Engineering Societies in the Agents World III. Volume 2577 of LNCS. Springer-Verlag (2003) 68–81
6. Susi, T., Ziemke, T.: Social cognition, artefacts, and stigmergy: A comparative analysis of theoretical frameworks for the understanding of artefact-mediated collaborative activity. Cognitive Systems Research **2**(4) (2001) 273–290
7. Schmidt, K., Wagner, I.: Ordering systems: Coordinative practices and artifacts in architectural design and planning. Computer Supported Cooperative Work **13**(5-6) (2004) 349–408
8. Viroli, M., Holvoet, T., Ricci, A., Schelfthout, K., Zambonelli, F.: Infrastructures for the environment of multiagent systems. Autonomous Agents and Multi-Agent Systems **14**(1) (2007) 49–60 Special Issue: Environment for Multi-Agent Systems.
9. Gambardella, L.M., Dorigo, M., Middendorf, M., Stützle, T.: Special section on ant colony optimization. IEEE Transactions on Evolutionary Computation **6**(4) (2002) 317–365
10. Holland, O., Melhuis, C.: Stigmergy, self-organization, and sorting in collective robotics. Artificial Life **5**(2) (1999) 173–202
11. Brueckner, S.A., Parunak, H.V.D.: Self-organizing MANET management. [24] 20–35
12. Mamei, M., Zambonelli, F.: Programming stigmergic coordination with the tota middleware. In: 4th ACM International Joint Conference on Autonomous Agents and Multiagent Systems, New York, USA, ACM (2005) 415–422
13. Nardi, B., ed.: Context and Consciousness: Activity Theory and Human Computer Interaction. MIT Press, Cambridge, MA (1996)
14. Hutchins, E.: Cognition in the Wild. MIT Press, Cambridge, MA (1995)
15. Castelfranchi, C.: When doing is saying: implicit communication before and without language and gestures. In: International Workshop "Evolving Communication: From Action to Language", Siena, Italy (2004) Communication.

16. Omicini, A., Ricci, A., Viroli, M., Castelfranchi, C., Tummolini, L.: Coordination artifacts: Environment-based coordination for intelligent agents. In Jennings, N.R., Sierra, C., Sonenberg, L., Tambe, M., eds.: 3rd international Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004). Volume 1., New York, USA, ACM (2004) 286–293
17. Wikipedia, The Free Encyclopedia: Home page. http://www.wikipedia.org (2006)
18. Amazon: Home page. http://www.amazon.com (2006)
19. Ricci, A., Viroli, M., Omicini, A.: Programming MAS with artifacts. In Bordini, R.P., Dastani, M., Dix, J., El Fallah Seghrouchni, A., eds.: Programming Multi-Agent Systems III. Volume 3862 of LNAI. Springer (2006) 206–221 3rd International Workshop (PROMAS 2005), AAMAS 2005, Utrecht, The Netherlands, 26 July 2005. Revised and Selected Papers.
20. Omicini, A., Ricci, A., Viroli, M.: *Agens Faber*: Toward a theory of artefacts for MAS. Electronic Notes in Theoretical Computer Sciences **150**(3) (2006) 21–36 1st International Workshop "Coordination and Organization" (CoOrg 2005), COORDINATION 2005, Namur, Belgium, 22 April 2005. Proceedings.
21. Omicini, A., Zambonelli, F.: Coordination for Internet application development. Autonomous Agents and Multi-Agent Systems **2**(3) (1999) 251–269
22. Omicini, A., Denti, E.: From tuple spaces to tuple centres. Science of Computer Programming **41**(3) (2001) 277–294
23. Denti, E., Natali, A., Omicini, A.: On the expressive power of a language for programming coordination media. In: 1998 ACM Symposium on Applied Computing (SAC'98), Atlanta, GA, USA, ACM (1998) 169–177 Special Track on Coordination Models, Languages and Applications.
24. Di Marzo Serugendo, G., Karageorgos, A., Rana, O.F., Zambonelli, F., eds.: Engineering Self-Organising Systems: Nature-Inspired Approaches to Software Engineering. In Di Marzo Serugendo, G., Karageorgos, A., Rana, O.F., Zambonelli, F., eds.: Engineering Self-Organising Systems: Nature-Inspired Approaches to Software Engineering. Volume 2977 of LNAI., Springer (2004)