Available online at www.sciencedirect.com

ScienceDirect

Cognitive Systems Research 38 (2016) 50–59

Cognitive Systems RESEARCH

www.elsevier.com/locate/cogsys

Stigmergy as a universal coordination mechanism II: Varieties and evolution

Francis Heylighen

Evolution, Complexity and Cognition Group, Center Leo Apostel, Vrije Universiteit Brussel, Krijgskundestraat 33, B-1160 Brussels, Belgium

Abstract

The concept of stigmergy, a mechanism for the coordination of actions via the trace they leave in a medium, can explain self-organizing activities in a broad range of domains, including social insects, collaborative websites, and human institutions. The present paper attempts to bring some order to these diverse applications by classifying varieties of stigmergy according to general aspects: the number of agents involved, the relative persistence or transience of the trace, the use of sematectonic or marker-based traces, and the quantitative or qualitative characteristics of traces. The resulting cases are essentially continuous, as more complex cases can be understood as having evolved out of simpler ones. One application is cognition, which can be viewed as an interiorization of the individual stigmergy that helps an agent to perform a complex project by registering the state of the work in the trace, thus providing an external memory. Another application is the evolution of cooperation, in which agents learn to profit from the synergy produced by the spontaneous stigmergic coordination between their initially independent actions, thus bypassing the problem of "free riders" that exploit the cooperators' efforts.

Keywords: Stigmergy; Coordination; Sematectonic; Cognition; Extended mind; Evolution of cooperation

1. Introduction

In a preceding paper (Heylighen, 2015), stigmergy was defined as a mechanism for the coordination of actions, in which the trace left by an action on some medium stimulates the performance of a subsequent action. This generic definition is applicable to a very broad variety of cases, including the pheromone traces used by ants to find food, the self-organization of chemical reactions, and the implicit collaboration between people via the edits they make in a shared website, such as Wikipedia.

To bring some order to these phenomena, the present paper will develop a classification scheme for the different varieties of stigmergy. We will do this by defining fundamental dimensions or *aspects*, i.e. independent parameters along which stigmergic systems can vary. The fact that these aspects are continuous ("more or less") rather than dichotomous ("present or absent") may serve to remind us that the domain of stigmergic mechanisms is essentially connected: however different its instances may appear, it is not a collection of distinct classes, but a space of continuous variations on a single theme—the stimulation of actions by their prior results.

This continuity will further help us to understand the evolution of stigmergic mechanisms, from rudimentary to more sophisticated applications. The paper will in particular focus on the evolution of two applications of stigmergy that are particularly important for humans, cognition and cooperation, arguing that these phenomena, which are traditionally viewed as difficult to explain via conventional evolutionary mechanisms, actually seem to emerge rather naturally via stigmergy.

E-mail address: fheyligh@vub.ac.be

http://dx.doi.org/10.1016/j.cogsys.2015.12.007
1389-0417/© 2015 Elsevier B.V. All rights reserved.

2. Individual vs. collective stigmergy

Perhaps the most intuitive aspect along which stigmergic systems can vary is the number of agents involved. In the limit, a single agent can coordinate its different actions via stigmergic interaction with the medium in which it acts.

An elegant example discussed by (Theraulaz & Bonabeau, 1999) is the solitary wasp *Paralastor* sp. building its nest in the shape of a mud funnel. The nest emerges in qualitatively different stages $S_1, S_2, \ldots, S_5$. These subsequently perceived conditions or stimuli each trigger a fitting action or response: $S_1 \rightarrow R_1, \ldots, S_5 \rightarrow R_5$. Each building action R_i produces as a result a new condition S_{i+1} that triggers the next action R_{i+1}. The wasp does not need to have a plan for building such a nest, or to remember what it already did, because the present stage of the activity is directly visible in the trace left by the work already done.

However, the underlying rule structure becomes apparent when the sequence is disturbed so that stages are mixed up. For example, the wasp's initial building activity is triggered by the stimulus S_1, a spherical hole. When at stage S_5 (almost complete funnel) the observer makes such a hole on top of the funnel, the wasp "forgets" that its work is nearly finished, and starts anew from the first stage, building a second funnel on top of the first one. This little experiment shows that the activity is truly stigmergic, and can only run its course when the medium (the mud) reacts as expected to the different actions performed on it, thus registering the information needed to guide the subsequent actions.

As (Theraulaz & Bonabeau, 1999) suggest, it is likely that collaborative stigmergy evolved from the simpler case of individual stigmergy. Imagine that a second wasp encounters the partially finished nest of the first wasp. It too will be stimulated to act by the perception of the present state of work. It does not matter that this state was achieved by another individual: the wasp anyway has no memory of previous actions—its own or someone else's. Assume further that the resulting structure is big enough to house the two wasps. In this case, the wasps will have collaboratively built a nest for both, without need for any additional coordination between their genetically programmed building instructions. Assume that the structure is modular, like the nests of social wasps, so that an unlimited number of modules can be added. In that case, the number of wasps that may start working together simply by joining the on-going activity on an existing nest can grow without limit.

This example illustrates how the number of agents collaborating on a stigmergic project is actually much less fundamental than it may seem. The essence of the activity is always the same. Assuming that the agents have the same skills, adding more agents merely increases manpower and therefore the size of the problem that can be tackled, the speed of advance, or the eventual magnitude of the achievement. Only when the agents are diverse can an increase in their number produce a qualitative improvement in the solution via a division of labor, where differently skilled agents contribute different solutions.

The only complication added by increasing the number of agents is that agents may get in each other's way, in the sense that similar individuals perceiving the same stimulus are likely to move to the same place at the same time, thus obstructing each other's actions. This problem is easily tackled by an additional rule, which is already implicit in individual work but likely to become reinforced during collaborative work: *keep a minimum distance from obstacles*—including other agents. This rule is a well-known ingredient in the many successful simulations of collectively moving animals, such as flocks, schools or swarms (Okubo, 1986), allowing densely packed groups of agents to follow complex, synchronized trajectories without ever bumping into each other. In combination with the basic stimulation by the stimulus object, this leads to what may look like a carefully thought-out strategy of coordinated movement. An example are group hunting strategies, as used e.g. by lions or wolves (Parunak, 2006). Each wolf is attracted to move toward the prey (basic stimulus). On the other hand, each wolf is stimulated to stay as far away as possible from the other wolves. The result is an efficient encirclement of the prey, which is attacked simultaneously from all sides with no opening left for escape.

3. Sematectonic vs. marker-based stigmergy

Grassé's original definition of stigmergy (Grassé, 1959) concerned stimulation by the performed work itself: in his observation, termites are stimulated to add mud by the mud heaps that are already there. Wilson (1975), in his monumental "Sociobiology", called such stimulation *sematectonic*. However, in many cases social insects appear to be stimulated by pheromone traces that are left expressly as a signal for others, not as a contribution to the work itself. In fact, it later turned out that termites are actually stimulated by the pheromones mixed in with the mud by co-workers rather than by the mud itself (which was to be expected given that termites are blind). The situation is even clearer for ants laying trails. In principle, ants could be guided by the perceivable results of their activity—the way humans and large animals are guided by the trails of flattened vegetation and sand eroded by the movement of previously passing individuals. However, the effect of an ant's movement on its surrounding is so small as to make it practically undetectable. Therefore, ants appear to have evolved a special type of chemical markers—phero mones—that make the traces of their activity much more salient. This type of indirect stimulation, not by the work itself but by a specially evolved "side-effect", has been called *marker-based* stigmergy (Parunak, 2006).

The evolution of markers is an obvious method to make stigmergy more efficient, by more reliably focusing the agents' attention on the most relevant aspects of the work that needs to be done. However, it entails an extra cost in that individuals need to perform the task of manufacturing markers in addition to the work itself. A human example can be found in the Wikipedia encyclopedia on the web. Readers are stimulated to improve existing pages either directly, by reading the text and noticing its shortcomings, or indirectly, by reading comments that summarize the tasks that still need to be done—such as adding references, clarifying ambiguous sections, or checking facts (Heylighen, 2007). The direct method exemplifies sematectonic stigmergy, the indirect one marker-based stigmergy. The "markers" in this case are the various "to do" notes that attract the attention to the problems that still require work.

A marker can be seen as an abstract, conventional sign, intentionally representing the work to be done instead of mechanically registering its effects. In Peirce's semiotic taxonomy of signs (Burks, 1949; Merrell, 2001), a marker is a *symbol*, while a sematectonic trace is an *index*. As such, a marker may seem to belong to a higher-order semiotic or communicative category of phenomena—a "meta-level" compared to the "object level" of the work itself. However, as in all phenomena produced by evolution, there is an essential continuity between the more primitive and the more "advanced" versions, as we can illustrate with a well-known example.

Many animals mark their territory by leaving traces of urine all around it. Obviously, excreting urine was not initially intended as a communicative signal, but merely as a way to get rid of liquid waste products. But since urine is easily perceived because of its smell, while its presence is causally connected to the presence of its producer, animals quickly learned to interpret it as a sign ("index") of the presence of another animal in the vicinity. Such a signal constitutes potentially vital information, which is useful, both for the receiver, who is warned of a rival that may be dangerous, and for the emitter, who can use it to frighten away newcomers from its territory. Thus, both parties are taught by evolution to communicate more reliably by means of this signal, turning it into a conventional marker of territory. As a result, animals have learned to deposit a little urine at regular intervals around their territory rather than simply emptying their bladder in a random place when it is full. This marker now supports stigmergic coordination between foraging activities, by clearly delimiting each individual's hunting grounds, and thus minimizing the risks of encounters ending in conflict.

The effect is equivalent to the human institution of "property rights"—the formal establishment of what is owned by whom, which economists consider essential for dependable transactions (Martens, 2004). The simplest way to establish a property right is to put a fence around the territory that you consider to be your property. Like the urine trace, this provides a clearly perceivable signal to others that they should not trespass there, obviating the need for individual communication with each of those others.

In the animal case, we see how a phenomenon (smell) that was merely a side effect of a primary action (getting rid of waste products) turned into an intentional, communicative signal—even though its primary function of waste disposal is still essential. In the case of pheromones, this original function, whatever it may have been, seems to have been lost, leaving only the communicative function. But in the most general case, both functions, primary and communicative, are likely to play a part. The fence, for example, not only warns people not to trespass, but keeps cattle from getting out. Another human example is an artist making a sketch. The sketch functions both as a first step toward performing the intended work (e.g. drawing someone's portrait) and as a representation of what the finished work may look like—which can be used to convince a sponsor who may be interested to order the finished work. The first function is sematectonic, the second one marker-based.

4. Transient vs. persistent traces

After discussing basic aspects of stigmergy that are recognized in the literature (e.g. Parunak, 2006), I wish to suggest a new dimension of variation. Parunak, in his attempt at classification, proposed the dynamics of the environment (what I call medium) as a crucial factor in stigmergy. However, there exists an infinite variety of potential dynamics of different degrees of complexity, thus making classification practically impossible. Moreover, a non-trivial dynamics seems better captured by causal rules, and as such by a system of (agentless) actions transforming the state of the world (Heylighen, 2015). For example, a collectively edited website, like Wikipedia, may have some in-built procedures that automatically correct formatting errors, add links, or signal incoherencies. The fact that these actions are performed by computer programs (e.g. "bots") does not fundamentally distinguish them from the actions of human contributors, since they all undergo the same stigmergic coordination. We have conceptualized the medium as the *passive* component of the stigmergic system, which undergoes shaping by the actions, but does not participate in the activity itself.

But even a passive medium is subjected to dissipation, as entailed by the second law of thermodynamics. That means that structures and markers tend to decay—unless they are actively maintained and reconstructed. Examples are the evaporation of pheromones and the wearing down of termite hills by rain, wind and gravity. This decay is not a priori negative. The traces left in the medium function as instructions for further work. It is obvious that without continuing updates this information will little by little become obsolete as the situation changes. For example, pheromone trails that point to exhausted food sources have become not just irrelevant, but misleading, since they incite ants to make useless journeys. Happily, pheromone trails that are no longer reinforced—because ants following them do not return with food—will gradually diffuse, and thus lose their attractiveness relative to trails that continue to receive reinforcement.

This is the same phenomenon of selective “forgetting” that characterizes memory in the brain: neural connections that are no longer reinforced will gradually lose their strength relative to recently reinforced ones. The speed of this forgetting depends on the *learning parameter*, as defined in neural networks (Heskes & Kappen, 1992). A large value of the parameter means that new changes in connection strength are large relative to previous ones, thus promoting the speedy establishment of new memory traces—but also the quick obsolescence of older traces. A small value, on the other hand, means that older learning episodes continue to exert a strong effect.

A similar parameter probably controls the external memory of ants as laid down in pheromone trails: newly added pheromone should be strong enough to allow trails toward newly found food sources to eventually become more attractive than previously found ones; yet, it should not be so strong that some recent journeys by ants carrying food from a new, unproven source can overpower the signals pointing to an older source whose reliability is evidenced by hundreds of successful journeys (Heylighen, 1999). The optimal value of this parameter will depend on the speed with which information becomes obsolete. This will depend on the variability in the environmental diversions and the measures that are taken to control them. For example, the location of a particular pillar in a termite hill is unlikely to become obsolete quickly, since the disturbances and affordances that it regulates, such as protection against sun, cold and predators or the creation of a comfortable interior microclimate, generally do not change position. Abundant food sources for ants, on the other hand, tend to change location every few days or hours.

Some diversions, such as the sudden appearance of a predator or prey animal, are even more short-lived. In this case, a trace inciting the appropriate action should be as quick to appear as to disappear. Typical stigmergic signals will be acoustic (e.g. the warning cry uttered by a monkey that spots a snake—which is marker-based) or visual (e.g. the visible movement of a wolf toward a deer—which is sematectonic). The reason is that sound and light spread and decay almost immediately. An intermediate decay speed is typical for chemical traces in a liquid environment, where concentrations of molecules may change within minutes. An example of such kind of stigmergic coordination are the chemical signals broadcasted by bacteria (Gloag, Turnbull, & Whitchurch, 2015) when they encounter either an affordance, such as food, or a disturbance, such as a concentration of toxins. The first type of diffusing signal will incite bacteria of the same colony to swim toward the food source, so that they too can profit from it. In the second case, it will incite them to move away from the danger.

These examples illustrate once again that no sharp distinction can be made between *persistent* and *transient* traces used in stigmergy: these are merely the opposite ends of a continuum. Yet, the distinction may be useful for conceptual clarification. Persistent traces lead to what may be called *asynchronous* stigmergy: the different agents do not need to be present at the same time, since the trace remains to guide them at any later time. Asynchronous communication (Cristian, 1996) can be illustrated by media such as print, email, or websites. Its advantage is that information remains available, so that it can be processed at the most appropriate occasion, and can accumulate and mature over the longer term. Transient traces lead to *synchronous* stigmergy: the agents need to be simultaneously present for the coordination to succeed. Synchronous communication may be exemplified by media such as telephone and Internet “chat”. Its advantage is that interaction, and therefore feedback, is instantaneous, so that urgent problems can be tackled immediately.

Synchronous action is rarely conceived as stigmergic, since it is typically used for direct communication, such as conversation or discussion. Yet, a warning cry or a chemical signal exemplify indirect communication: they are targeted at no one in particular, but merely “released” in the medium. Examples of stigmergy in synchronous interaction are even clearer when the signal is sematectonic. For example, a bird spotting a danger (condition) will start to fly (action), and by this example (transient trace) set off the whole flock to fly away (subsequent action). Synchronous stigmergy may be best exemplified by the collective movement in herds, flocks or swarms (Moussaid, Garnier, Theraulaz, & Helbing, 2009; Okubo, 1986), where the agents are continually adjusting their trajectory on the basis of real-time perceptions of the movements of other agents.

A human example would be the self-organization of traffic, where drivers continuously react to the traffic conditions they perceive, by e.g. stopping, accelerating, or changing lanes, thus affecting these very conditions and the subsequent actions of other drivers. Roads, lanes, road markings and traffic signs, on the other hand, function like a persistent trace developed over decades in order to stimulate the drivers to move in a coordinated manner. The continuity between the two is demonstrated by the fact that in sufficiently dense traffic lanes tend to self-organize and acquire some form of stability, even when they leave no permanent trace (Helbing, 2001; Moussaid et al., 2009). Nevertheless, when the surface is soft enough to show signs of erosion, like in dirt roads, traces persist after the traffic stops, thus maintaining a memory of the self-organized traffic pattern. This persistent trace reduces the time necessary to rebuild a coordinated movement pattern when the traffic starts up again. It seems likely that many roads have emerged in this manner across historical time.

5. Quantitative vs. qualitative stigmergy

Quantitative stigmergy (Theraulaz & Bonabeau, 1999) refers to perceived conditions that differ in strength or degree, and where stronger traces typically elicit more forceful (intense, frequent, ...) actions. This quantitative variation is perhaps best captured using the definition of stigmergic stimulation in terms of conditional probability (Heylighen, 2015): the stronger the trace, the higher the probability of a certain action given that trace. Over an extended period, higher probability implies more frequent actions by more numerous agents, and therefore more intense overall activity. The two paradigmatic cases of stigmergy, termite nest-building and ant trail-laying, follow this quantitative logic. The higher the emerging heap of mud (stronger trace), the more an individual termite is attracted to it, and therefore the larger the probability or frequency of mud being added. The stronger the scent of pheromone on a trail, the less likely an ant is to deviate from that trail, and therefore the higher the probability that it too will reinforce the trail with additional pheromone. These are typical examples of the positive feedback that efficiently amplifies beneficial developments (Heylighen, 2015).

But quantitative stigmergy is also exemplified by negative feedback, where a stronger trace leads to less activity. A human example can be found in the market mechanism. Extensive buying of a commodity (action) reduces the supply and thus increases the price. The price functions as a quantitative trace left by the collective buying and selling activity. A higher price will normally reduce the probability that someone would buy additional stock of that commodity (negative stimulation). Thus, a higher price reduces demand, which in turn will reduce the price. This mechanism of self-organizing, distributed control (Heylighen, 1997) implements the “invisible hand” of the market. It stabilizes prices and efficiently allocates production capacity to the commodities that are most in demand (Witt, 2006).

Qualitative stigmergy (Theraulaz & Bonabeau, 1999) refers to conditions and actions that differ in kind rather than in degree. In this case, a different trace stimulates a different type of action. An example can be found in the different stages of the building of a funnel-shaped nest by the solitary wasp that we discussed, where each stage requires a particular type of building action. A human example can be found in “wiki” websites that are edited by their own readers. A paragraph that contains a semantic mistake (e.g. in the definition of a word) will elicit a corrective action (e.g. writing a new definition). Different types of errors, vagueness, or lack of information will stimulate different types of additions and corrections.

In practice, there is no clear boundary between quantitative and qualitative cases of stigmergy. All non-trivial activities require a choice from a range of potential actions. Which of the different possibilities will be chosen is typically determined probabilistically: in some conditions one type of action is more likely, in other conditions another type of action, and this probability is a quantitative variable.

6. Extending the mind

Traditionally, cognition has been viewed as the processing of information inside the brain. More recent approaches, however, note that both the information and the processing often reside in the outside world (Clark, 1998; Dror & Harnad, 2008; Hollan, Hutchins, & Kirsh, 2000)—or what we have called the *medium*. For example, documents function as an external memory for storing knowledge and data, while calculations are typically performed on a piece of paper or on a computer. Without such supporting media, most advanced reasoning—as performed e.g. in science and technology—would be simply impossible. Thus, the human mind *extends* into the environment (Clark & Chalmers, 1998), “outsourcing” some of its functions to external support systems. The reason is that our memory and information processing capabilities have rather strict limitations (Heylighen, 2013; Heylighen & Vidal, 2008). Books and computers are relatively recent inventions. However, the use of an external medium for supporting cognition is probably as old as cognition itself.

In fact, our mental capabilities can be seen as an interiorization of what were initially stigmergic interactions with the environment. The perspective of situated and embodied cognition (Aydede & Robbins, 2009; Steels & Brooks, 1995) focuses on the interaction between the agent and its environment: the agent senses the state of the environment via its sensory organs and reacts to it by producing an appropriate action via its muscles or effectors: it reacts in the same way to the returning feedback signal. Such a reaction requires merely a condition-action rule, which is nothing more than a causal process transforming an antecedent into a consequent (Heylighen, 2015). As such, condition-action rules are readily implemented in the simplest systems, such as thermostats. Yet, when the activity of these rules is coordinated by stigmergy, it becomes capable of complex, goal-directed behavior, such as building a wasp’s nest or a spider’s web.

A classic example of the “intelligence” exhibited by such simple rules can be found in Braitenberg vehicles (Braitenberg, 1986; Gershenson, 2004)—rudimentary automata equipped with just two sensors (left and right) for light intensity, two wheels for movement, and connections between them. The rule here is that the speed of a wheel increases in proportion to the intensity of light perceived by its corresponding sensor. When a light shines on the right-hand side of the vehicle, the right wheel will move faster than the left one, thus turning the vehicle to the left, away from the light. The effect of this rule is that the vehicle follows a complex, zigzagging trajectory that evades as much as possible all light sources until it reaches a place of darkness where it comes to a standstill (see Fig. 1). This is similar to the adaptive and goal-directed behavior of a simple animal, such as a cockroach, that is afraid of light and seeks obscurity to hide. The apparent intelligence of this behavior can be understood from the stigmergic coordination between the actions performed, in sequence and in parallel, by the wheels when reacting to the sensed light conditions (Heylighen, 2010). These conditions depend on the vehicle’s position and orientation, which is the accumulated result of its previous series of movements. Thus, the

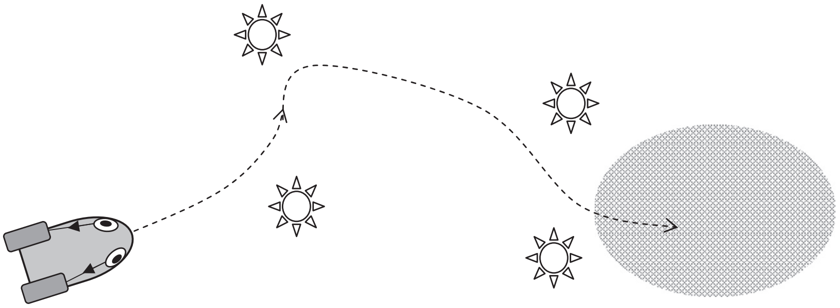

Fig. 1. A Braitenberg vehicle (left) with two light sensors connected to two motorized wheels. Each sensor drives the speed of the corresponding wheel proportionally to the amount of light it senses. The effect is that the vehicle accelerates away from any light source (star-like shapes), while continuing to move until it finds a dark place (shaded area on the right) to "rest".

sensed condition functions like a transient trace of actions that stimulates subsequent actions.

Both the strength and weakness of such stigmergic activity is that it lacks internal memory: information about the state of the process is stored purely in the medium from where it is sensed by the agent. The advantage is that there is no need for the registration, maintenance, and recollection of information in the brain. The disadvantage is that if the medium is disturbed, then the trace and with it the memory may be erased. We saw an example of this problem with the nest-building wasp: when the experimenter creates a misleading trace on the nearly finished nest, the wasp starts building a new nest on top of the old one, thus uselessly duplicating its effort. It is likely that our capability for internal information storage evolved at least in part to avoid this problem: if the state of the activity can be registered and processed internally, complex activities can be planned even when the external medium does not cooperate.

Thanks to this capability, humans are much smarter than insects. Nevertheless, our brain is an energy-intensive, costly organ, whose storage capacities remain quite limited. That is why we continue to use stigmergy to support our memory and reasoning. Let us discuss a few examples. Whenever we have to do a complex job, such as repairing a bicycle, preparing a dinner, or filling out our tax forms, we tend to keep both the objects we work on, and the different tools and resources that support the work at hand, in such a way that they are easy to see and to manipulate.

For example, while taking apart the bicycle we arrange all the screws and pieces in clear view, close to the screwdrivers or pincers we will need to put them back on, so that we are unlikely to forget what must be added when and where. Each tool or piece is a stimulus for performing a particular action. The perceived state of the bicycle is the condition that determines which action is to be performed next. If before we start we had to analyze, plan and memorize all the steps that need to be performed in taking apart, repairing, and then reassembling the bicycle, it is unlikely that we would ever succeed in this task. The arrangement of the physical components in space here plays the role of the activity's trace, which constantly guides the stigmergic coordination of actions.

Ergonomic studies have shown that the spatial arrangement of a workplace is crucial to the efficient performance of work (Hollan et al., 2000; Kirsh, 1995, 1996). One obvious reason is that when tools are positioned near to where they are likely to be used, there will be less need for physical movement. However, stigmergy reminds us that good arrangement saves cognitive effort as well as physical effort, by connecting the right reminders to the right circumstances. For example, one of the reasons why "Post it" notes are popular is that they make it easy to spatially connect a cognitive "call for action" (challenge, stimulus, marker) with the physical resource needed to perform the action. Sticking a "Please photocopy!" note on a document, e.g., makes it obvious for anyone what needs to be passed through the copying machine.

The full power of individual stigmergy is seen with creative work—such as drawing a picture, writing a text, or modeling a piece of clay. Here, the provisional results of the work are fully embodied in the trace, be it a sketch, a draft document, or a clay shape. This preliminary registry of the work performed calls out for more. It challenges the user to add, to enhance or to correct. Each addition changes the trace, thus attracting the attention to further imperfections, or suggesting further additions. It would be extremely difficult, if not impossible, to achieve the same level of sophistication in a design that would only exist inside the creator's brain, where all the planning would take place without any external medium to store it, test it, and be challenged by it.

While painters or writers may have a general idea of the piece they want to create, the actual details will only take shape when that idea is exteriorized in a medium that can be scrutinized and manipulated, so that its structure step-by-step acquires the ideal shape for the purpose. That makes it possible to take into account all the possibly unforeseen properties and side effects of an initially still abstract idea. This principle is at the basis of the method of *stigmergic prototyping* (De Couvreur, Detand, Dejonghe, & Goossens, 2012; Dejonghe, Detand, & De Couvreur, 2011), in which a conceived artefact is immediately given a rudimentary physical shape that can be easily tried out and thus adapted to the circumstances. In contrast, the traditional approach is to first create a detailed, abstract blueprint of the artefact. Unfortunately, it then often turns out that its physical implementation does not work as intended, forcing the designer to go "back to the drawing board".

7. The evolution of cooperation

As we have noted, stigmergy intrinsically does not distinguish between individual and collective activity: the trace left in the medium coordinates actions, while being indifferent as to the specific agent or agents initiating these actions. The only additional requirement for collective action is that the different agents should not work at cross-purposes, so that the one's actions negate or obstruct the other one's (Heylighen, 2015). But even such a conflict tends to remain localized to a small part or aspect of the trace, while allowing the rest of the trace to develop unhindered.

An example can be found in Wikipedia "edit wars" (Sumi, Yasseri, Rung, Kornai, & Kertész, 2011), in which two contributors who disagree about a particular statement in a Wikipedia article repeatedly undo each others' corrections. This does not prevent other contributors from elaborating the rest of the article (and the encyclopedia). Often, the conflict tends to get resolved by a third party who proposes a compromise statement that the conflicting parties no longer object to. Even without third party intervention, the conflict is unlikely to continue, either because the antagonists themselves chance upon a statement that is acceptable to both, or because one of them simply gives up repeating the same ineffectual action, and decides to focus on some more productive task.

From this stigmergic perspective, the emergence of cooperation between selfish individuals seems a much less daunting issue than from a traditional evolutionary or economic perspective (Axelrod, 1997). Traditional models of the evolution of cooperation pit one individual against another one in a Prisoners' Dilemma type of interaction, where it pays to "defect" (i.e. be uncooperative) in the short run, even though everybody would be better off being cooperative in the long run. Another common paradigm is the Tragedy of the Commons, in which selfish individuals ("free riders") exploit—and eventually exhaust—the common good that others try to maintain cooperatively (Feeny, Berkes, McCay, & Acheson, 1990; Hardin, 1968). For example, a person who consumes more than his fair share of a common resource, such as water, grass, or land, will leave less of the resource for the people dividing up the resources more evenly. In such cases, the cooperative arrangement tends to be undermined by selfish agents appropriating more benefit from it than earnest cooperators, thus tempting others away from cooperation.

In the stigmergic paradigm, the common good (e.g. Wikipedia, or a network of trails connecting common destinations) is gradually built up via the cooperation implicit in stigmergically-coordinated actions. Free riders may profit from this common good without putting in any effort in return. However, the benefit derived from a stigmergic trace does not in general reduce the value of that trace. For example, an ant that follows a pheromone trace laid by others without adding pheromone of its own does not by that action make the pheromone trace less useful to the other ants. Similarly, a person who downloads a piece of open source software without contributing to the development of that software does not impose any burden on the software developers (Heylighen, 2007). Thus, in a situation of stigmergy, a free rider or "defector" does not weaken the cooperators, in contrast to situations like the Prisoners' dilemma or Tragedy of the Commons.

In a sense, by not contributing the free riding agent merely weakens its own position, because it passes by the opportunity to adapt the trace to its own preferences. As we saw, the stigmergic trace is the aggregate of many independent actions, each of which helps the agent that performed it to achieve its goals. The ant that finds food but does not leave a pheromone trace on its way back to the nest not only does not help others to get to that food: it also does not help itself, because without the trace it is very unlikely to find the same food source again. The trace is both an individual and a collective "mental map" that indicates effective actions (Heylighen, 1999). Not leaving a trace makes your own future work harder than it needs to be.

Let us analyze the dynamics of free riding in more depth on an example inspired by what may be the simplest type of stigmergy, the creation of a trail across irregular terrain through the flattening of grass, dirt and other obstacles. Here, an easy-to-travel path emerges as a side effect of the regular movement of people or animals, while requiring no special effort from these agents. A more demanding version of this task is the establishment and maintenance of a path through dense vegetation, like in a forest. As quickly growing bushes and trees extend their branches, they eventually obstruct the path. A person following that path will have to either duck around these obstacles, or remove them, e.g. by cutting the twigs that intrude upon the open space. The first option may demand somewhat less effort, but that applies only to the short term, as the underbrush will grow until it becomes impassable. Somebody who regularly uses the path will be motivated to follow the second course of action, and remove any obstruction before it becomes insurmountable.

This preference is independent of the number of hikers actually using the path. Yet, the larger the number of people applying the strategy, the less work any one of them needs to perform. Thus, their actions are cooperative, as they help each other achieve their objectives. But such cooperation is purely stigmergic, because they travel independently of each other, at different times, and thus cannot communicate about their common purpose. People who only use the path occasionally may not contribute to this ongoing clearing activity, and thus "free ride" on the effort of others. But unless the others do enough work, the ones who use the path regularly will eventually have to make the effort for purely selfish purposes, because without that effort they will not be able to use the path anymore.

This example resembles the Prisoners' Dilemma or the Tragedy of the Commons in that there is a temptation to defect by letting the others do the hard work, while profiting from their results. The crucial difference, however, is that such a free riding strategy will eventually hurt the defector more than the "cooperator", because the cooperating agent will continue to clear its own path independently of any others (cooperators or defectors) using that path, while the defecting agent will eventually encounter a path that has become impassable without clearing effort, forcing it to either become a stigmergic cooperator, or give up on passing altogether. Thus, a defector will in the long run collect less benefit than a cooperator. This makes the strategy of non-cooperation self-defeating.

In the short run, the free rider may seem to have the benefit over the cooperator of spending less energy establishing and maintaining the trace. However, the cooperator collects other benefits. First, as we noted with the ant leaving pheromone or the hiker breaking off branches, the cooperating agent helps itself by creating a trace. Second, the stigmergic interaction will boost the benefits of that individual trace by stimulating others to expand on it. For example, an ant creating a trail to a new food source will incite others to explore the neighborhood of that source, potentially discovering even better sources or shorter trails. Similarly, the hiker who partially cleared a path will thus increase the probability that others would follow that same path while performing further clearing themselves. This is the positive feedback of actions eliciting more actions that makes stigmergy so effective. The free rider simply misses out on this potential amplification of its actions.

The full power of such synergetic interaction supported by stigmergy is seen in complex, creative work environments, where different agents contribute different skills, experiences and perspectives. Here, the work done by one individual is enhanced by the work of others with complementary abilities in a way that the single individual never could have achieved. Wikipedia and communities developing open source software development are prime examples, having achieved results that could not even have been reached via hierarchical, command-and-control strategies of coordination (Heylighen, 2007; Heylighen, Kostov, & Kiemen, 2013). Smaller scale examples are people posting photos, ideas, artwork, or essays on their blog, Twitter feed, or Facebook page, and getting feedback from friends, followers, or strangers, which help them to further develop their insights, while inspiring these others to build further on their experiences. In such cases, the benefits that accrue to the "cooperators" are direct, concrete, and stimulating enough to motivate them to produce more of such "public traces" in their medium of choice (Wikipedia, Facebook, ...).

Thanks to the user-friendly electronic medium, the material and human cost of publishing such traces is nearly zero. This combination of strong motivation, minimal cost, and effective stigmergic coordination turns the medium into a powerful system for mobilizing joint action (Heylighen et al., 2013). The result is a rapidly expanding "collaborative commons" (Rifkin, 2014)—a virtual workspace for stigmergic (and more traditional) cooperation that encompasses the planet. This world-wide stigmergic medium is presently developing into the equivalent of a *global brain* able to efficiently tackle the collective challenges of society (Heylighen, 2008, 2014).

While the ICT applications of human stigmergy most stir the imagination because of their virtually unlimited scale, we need to remember that the same mechanism has been supporting collaboration across human and evolutionary history. A final example may illustrate some of the more down-to-earth applications. People who garden like to show off the fruits of their labor to visitors, guiding them along flower beds, vegetable yards and fruit trees. Visitors with some knowledge of gardening will spontaneously comment on what they see. The resulting exchange of knowledge is triggered by the visible trace of the gardening work, in which visitors e.g. note that certain flowers in the garden are doing better or worse than in their own garden, prompting them to either ask or give advice on how to tend that particular variety. If the garden was communal, this sharing of information would naturally extend into sharing of physical work on the garden, with individual gardeners concentrating on the plants or tasks they feel most competent about or that are most in need of work. Thus, they create a more beautiful garden for all, while reducing the individual workload.

Most forms of human cooperation have this stigmergic dimension, where actions are triggered by the observable results of other people's actions rather than by direct requests or commands. The traffic example may remind us that most people do not require the directions from a policeman in order to cooperatively produce a smooth flow of vehicles. But because the explicit request from a policeman, co-worker or boss to perform a particular action requires conscious processing—if only to decide whether we will honor the request or not—we tend to be much more aware of such direct communication. Therefore, it tends to remain in our memory as a driver of our actions. Our reactions to the implicit challenge left in an evolving piece of work, on the other hand, tend to be subconscious and automatic. Therefore, we assume that we decide to perform a further action purely on our own initiative, ignoring that we are actually being driven by the stigmergic organization of the medium. But the coordinated activity that ensues is a truly ubiquitous mode of human cooperation, albeit one that has hardly received any attention until now.

8. Conclusion

We have examined different variations on the theme of stigmergy by distinguishing basic aspects or dimensions in which the mechanism can vary.

The *number of agents* involved turns out to be less fundamental than is generally assumed. Increasing that number will qualitatively enhance the result only if the agents are sufficiently diverse in the actions they contribute, yet sufficiently aligned in their strategies so that they do not hinder each other. Like the number of agents, the difference between *qualitative* and *quantitative* stigmergy does not seem essential, given that the notion of "stimulation" entails a quantitative aspect of intensity or probability, while the actions that are stimulated more or less intensively differ qualitatively. The difference between *sematectonic* traces—the concrete, observable results of work performed—and *markers*—traces left to guide subsequent actions, but without contributing to the work itself—is important but subtle. The use of markers allows a more fine-grained control of stigmergic coordination, but demands an advanced level of collective evolution, in which certain traces have acquired a conventional meaning among the agents that use them. The *transience* of the trace is crucial in order to ensure that the list of "to do's" remains up to date: in a quickly changing environment, actions need to adapt in time to new circumstances, which means that an outdated trace should decay before it would elicit too much useless activity; in a more stable environment, on the other hand, persistent traces enable the accumulation of a long and detailed memory.

We have then examined two more typically human applications of stigmergy, the support of cognition and cooperation, while emphasizing their evolution out of more primitive mechanisms that can be traced back to the first living organisms. *Cognition* is an application of individual stigmergy: the trace of activity in a medium functions like an external memory that facilitates storage and processing of information, thus reducing the burden on the brain. However, the fact that this mechanism supports some degree of intelligent activity even without a brain suggests that our cognitive abilities may have evolved by simply interiorizing some of this functionality that was initially provided by an external medium.

Cooperation is a side effect of collective stigmergy. Stigmergic coordination arises spontaneously, without need for any cooperative intent from the individuals. Since coordination is beneficial to the agents involved, evolution is likely to strengthen the condition-action rules that make them interact synergetically, while suppressing the rules that produce conflict or friction. Stigmergy moreover bypasses a classic obstacle to the evolution of cooperation: the "tragedy of the commons" where "free riders"—who profit from the fruits of cooperation without contributing to it—do better than the cooperators, thus eroding the cooperative arrangement. This problem is avoided because cooperators (1) do not lose any benefit, since the trace typically does not deteriorate through free rider exploitation; (2) get the additional benefit that the work they do not only helps themselves, but stimulates others to expand on it. Therefore, the free rider benefit of avoiding effort in building the trace does not seem large enough to allow them to outcompete the cooperators.

References

Axelrod, R. (1997). *The complexity of cooperation*. Princeton: Princeton University Press.

Aydede, M., & Robbins, P. (2009). *The Cambridge handbook of situated cognition*. Cambridge University Press, Retrieved from <http://www.cup.cam.ac.uk/aus/catalogue/catalogue.asp?isbn=9780521848329andss=exc>.

Braitenberg, V. (1986). *Vehicles: Experiments in synthetic psychology*. Cambridge, MA: MIT Press.

Burks, A. W. (1949). Icon, index, and symbol. *Philosophy and Phenomenological Research, 9*(4), 673–689.

Clark, A. (1998). Embodied, situated, and distributed cognition. *A Companion to Cognitive Science*, 506–517.

Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis, 58*(1), 7–19.

Cristian, F. (1996). Synchronous and asynchronous group communication. *Communications of the ACM, 39*(4), 88–97.

De Couvreur, L., Detand, J., Dejonghe, W., & Goossens, R. (2012). Expect the unexpected: The co-construction of assistive artifacts. In *Proceedings of the 8th International conference on Design and Emotion*. Central Saint Martins, University of the Arts. Retrieved from <https://biblio.ugent.be/publication/4194683>.

Dejonghe, W., Detand, W. J. & De Couvreur, L. (2011). *Stigmergic prototyping* (ECCO Working Papers No. 2011-06). Retrieved from <http://innowiz.be/Methodologiecursus/Werkelijkheid/Stigmergic_prototyping_2.04.html>.

Dror, I. E., & Harnad, S. R. (2008). *Cognition distributed: How cognitive technology extends our minds*. Amsterdam: John Benjamins Publishing.

Feeny, D., Berkes, F., McCay, B. J., & Acheson, J. M. (1990). The tragedy of the commons: Twenty-two years later. *Human Ecology, 18*(1), 1–19.

Gershenson, C. (2004). Cognitive paradigms: Which one is the best. *Cognitive Systems Research, 5*(2), 135–156.

Gloag, E. S., Turnbull, L., & Whitchurch, C. B. (2015). Bacterial stigmergy: An organising principle of multicellular collective behaviours of bacteria. *Scientifica, 2015*, e387342.

Grassé, P. P. (1959). La reconstruction du nid et les coordinations interindividuelles chez Bellicositermes natalensis et Cubitermes sp. La théorie de la stigmergie: Essai d'interprétation du comportement des termites constructeurs. *Insectes sociaux, 6*(1), 41–80.

Hardin, G. (1968). The tragedy of the commons. *Science, 162*(3859), 1243–1248.

Helbing, D. (2001). Traffic and related self-driven many-particle systems. *Reviews of Modern Physics, 73*(4), 1067.

Heskes, T. M., & Kappen, B. (1992). Learning-parameter adjustment in neural networks. *Physical Review A, 45*, 8885–8893.

Heylighen, F. (1997). The economy as a distributed, learning control system. *Communication and Cognition-AI, 13*(2/3), 207–224.

Heylighen, F. (1999). Collective intelligence and its implementation on the Web: Algorithms to develop a collective mental map. *Computational and Mathematical Organization Theory, 5*(3), 253–280.

Heylighen, F. (2007). Why is open access development so successful? Stigmergic organization and the economics of information. In B. Lutterbeck, M. Baerwolff, & R. A. Gehring (Eds.), *Open source Jahrbuch 2007* (pp. 165–180). Hamburg: Lehmanns Media.

Heylighen, F. (2008). *Accelerating socio-technological evolution: From ephemeralization and stigmergy to the global brain. Globalization as evolutionary process: Modeling global change. Rethinking globalizations* (pp. 284). New York: Routledge.

Heylighen, F. (2010). *Cognitive Systems: A cybernetic approach on the new science of the mind* (Lecture Notes.). ECCO, VUB. Retrieved from <http://pespmc1.vub.ac.be/Papers/CognitiveSystems.pdf>.

Heylighen, F. (2013). *From human computation to the global brain: The self-organization of distributed intelligence. Handbook of human computation* (pp. 897–909). Heidelberg: Springer.

Heylighen, F. (2014). Return to Eden? Promises and perils on the road to a global superintelligence. In B. Goertzel & T. Goertzel (Eds.), *The end* *of the beginning: Life, society and economy on the brink of the singularity*. Humanity+ Press.

Heylighen, F. (2015). Stigmergy as a universal coordination mechanism I: Definition and components. *Cognitive Systems Research*.

Heylighen, F., Kostov, I., & Kiemen, M. (2013). Mobilization systems: Technologies for motivating and coordinating human action. In M. A. Peters, T. Besley, & D. Araya (Eds.), *The new development paradigm: Education, knowledge economy and digital futures*. New York: Routledge.

Heylighen, F., & Vidal, C. (2008). Getting things done: The science behind stress-free productivity. *Long Range Planning, 41*(6), 585–605.

Hollan, J., Hutchins, E., & Kirsh, D. (2000). Distributed cognition: Toward a new foundation for human–computer interaction research. *ACM Transactions on Computer–Human Interaction (TOCHI), 7*(2), 174–196.

Kirsh, D. (1995). The intelligent use of space. *Artificial Intelligence, 73*(1), 31–68.

Kirsh, D. (1996). Adapting the environment instead of oneself. Retrieved from <http://philpapers.org/rec/DAVATE>.

Martens, B. (2004). *The cognitive mechanics of economic development and institutional change*. New York: Routledge.

Merrell, F. (2001). Charles Sanders Peirce's concept of the sign. *The Routledge Companion to Semiotics and Linguistics*, 28–39.

Moussaid, M., Garnier, S., Theraulaz, G., & Helbing, D. (2009). Collective information processing and pattern formation in swarms, flocks, and crowds. *Topics in Cognitive Science, 1*(3), 469–497.

Okubo, A. (1986). Dynamical aspects of animal grouping: Swarms, schools, flocks, and herds. *Advances in Biophysics, 22*, 1–94.

Parunak, H. V. D. (2006). A survey of environments and mechanisms for human–human stigmergy. In D. Weyns, H. V. D. Parunak, & F. Michel (Eds.), *Environments for multi-agent systems II* (pp. 163–186). Heidelberg: Springer.

Rifkin, J. (2014). *The zero marginal cost society: The Internet of things, the collaborative commons, and the eclipse of capitalism*. Palgrave Macmillan.

Steels, L., & Brooks, R. A. (1995). *The artificial life route to artificial intelligence: Building embodied, situated agents*. Mahwah, NJ: Lawrence Erlbaum.

Sumi, R., Yasseri, T., Rung, A., Kornai, A., & Kertész, J. (2011). Characterization and prediction of Wikipedia edit wars. In *Presented at the ACM WebSci'11, Koblenz, Germany* (pp. 1–3). Retrieved from <http://www.websci11.org/fileadmin/websci/Posters/58_paper.pdf>.

Theraulaz, G., & Bonabeau, E. (1999). A brief history of stigmergy. *Artificial Life, 5*(2), 97–116.

Wilson, E. (1975). *Sociobiology: The new synthesis*. Cambridge, MA: Harvard University Press.

Witt, U. (2006). Coordination of individual economic activities as an evolving process of self-organization. *The Evolving Economy: Essays on the Evolutionary Approach to Economics, 62*.