Available online at www.sciencedirect.com
ScienceDirect
Cognitive Systems Research 38 (2016) 4–13
Cognitive Systems RESEARCH
www.elsevier.com/locate/cogsys

Stigmergy as a universal coordination mechanism I: Definition and components

Francis Heylighen

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

Abstract

The concept of stigmergy has been used to analyze self-organizing activities in an ever-widening range of domains, including social insects, robotics, web communities and human society. Yet, it is still poorly understood and as such its full power remains underappreciated. The present paper clarifies the issue by defining stigmergy as a mechanism of indirect coordination in which the trace left by an action in a medium stimulates subsequent actions. It then analyses the fundamental concepts used in the definition: action, agent, medium, trace and coordination. It clarifies how stigmergy enables complex, coordinated activity without any need for planning, control, communication, simultaneous presence, or even mutual awareness. The resulting self-organization is driven by a combination of positive and negative feedbacks, amplifying beneficial developments while suppressing errors. Thus, stigmergy is applicable to a very broad variety of cases, from chemical reactions to bodily coordination and Internet-supported collaboration in Wikipedia.

Keywords: Stigmergy; Coordination; Actions; Agents; Self-organization; Feedback

1. Past, present and future of the "stigmergy" concept

The concept of *stigmergy* was proposed by the French entomologist Pierre-Paul Grassé (Grassé, 1959) to describe a mechanism of coordination used by insects. The principle is that work performed by an agent leaves a trace in the environment that stimulates the performance of subsequent work—by the same or other agents. This mediation via the environment ensures that tasks are executed in the right order, without any need for planning, control, or direct interaction between the agents. The notion of stigmergy allowed Grassé to solve the "coordination paradox" (Theraulaz & Bonabeau, 1999), i.e. the question of how insects of very limited intelligence, without apparent communication, manage to collaboratively tackle complex projects, such as building a nest.

The insight came from Grassé's observation of how termites repair their nest. He noted that initially termites wander around more or less randomly, carrying mud and depositing it here or there. However, the deposits that are created in this haphazard way then stimulate the insects to add more mud in the same place. Thus, the small heaps quickly grow into columns that eventually come together to form an intricate cathedral of interlocking arches. The only communication between the termites is indirect: the partially executed work of the ones provides information to the others about where to make their own contribution.

Another classic example of stigmergy can be found in the pheromone trails left by ants that come back from a food source (Sumpter & Beekman, 2003). The pheromone stimulates other ants to follow the same path. When they find food, they too will reinforce the pheromone trail while following the trail back to the nest. This mechanism leads to the emergence of an efficient network of trails

---

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

http://dx.doi.org/10.1016/j.cogsys.2015.12.002
1389-0417/

connecting the nest via the shortest routes to all the major food sources.

Up to about 1990, the notion of stigmergy appears to have remained limited to a small circle of researchers studying the behavior of social insects. However, one of these insect specialists, Jean-Louis Deneubourg, was also a member of the "Brussels School" of complex systems, headed by the late Nobel Prize in chemistry, Ilya Prigogine. In this interdisciplinary environment, it became clear that stigmergy was a prime example of spontaneous ordering or *self-organization* (Camazine et al., 2003; Deneubourg, 1977) and as such potentially applicable to complex systems other than insect societies.

With the advent of the agent-based paradigm in computer simulation, insect societies were conceptualized as *swarms* of simple agents that are able to perform complex tasks using various forms of self-organization and especially stigmergy (Deneubourg, Theraulaz, & Beckers, 1992). The general ability to tackle complex problems exhibited by such self-organizing multi-agent collectives became known as *swarm intelligence* (Bonabeau, Dorigo, & Theraulaz, 1999; Kennedy, 2006). One class of stigmergic mechanisms in particular, so-called *ant algorithms*, turned out to be surprisingly powerful in tackling a variety of computational problems, including the notorious traveling salesman problem (Dorigo, Bonabeau, & Theraulaz, 2000) and the optimization of packet routing along communication networks (Kassabalidis, El-Sharkawi, Marks, Arabshahi, & Gray, 2001).

Stigmergy was applied not only to software agents, but to their hardware analogues: autonomous robots. Groups of very primitive robots proved able to tackle non-trivial tasks, such as clustering items in different groupings, in a way similar to ants (Beckers, Holland, & Deneubourg, 1994; Deneubourg et al., 1991; Holland & Melhuish, 1999). These robotic implementations inspired the application of stigmergic models to problems of coordination and control in manufacturing (Valckenaers, Van Brussel, Kollingbaum, & Bochmann, 2006). After this expansion of the stigmergy concept from social insects to the domains of artificial life, artificial intelligence and behavior-based robotics, a perhaps obvious next step was computer-supported collaboration between human agents, in particular via the world-wide web (Bolici, Howison, & Crowston, 2009; Dron, Boyne, & Mitchell, 2001; Elliott, 2007; Heylighen, 1999).

A prototypical example is Wikipedia, the free web encyclopedia which has grown to become the largest one in existence thanks to the fact that every reader is stimulated to improve and expand the writings of previous contributors (Heylighen, 2007). A similar dynamics of contributions building further on previous contributions characterizes open-source software development (Bolici et al., 2009; Robles, Merelo, & Gonzalez-Barahona, 2005). But it quickly became clear that human collaboration does not need computer support to profit from stigmergy (Elliott, 2007; Parunak, 2006). Probably the best-known example of stigmergic self-organization is the "invisible hand" of the market: the actions of buying and selling leave a trace by affecting the price of the transacted commodities. This price in turn stimulates further transactions. Via the related conceptions of distributed cognition and the extended mind, stigmergy has now also started to make its mark on theories of cognition and epistemology (Marsh & Onof, 2008; Ricci, Omicini, Viroli, Gardelli, & Oliva, 2006; Susi & Ziemke, 2001).

It is clear that since 1990, the concept of stigmergy has undergone a rapid diffusion across an ever-growing number of application domains. I contend in this paper that the potential for theoretical explanation and practical application of the stigmergy concept is much larger still than hitherto assumed. What (Parunak, 2006) noted about human institutions, that the more difficult issue is to find examples where stigmergy does *not* apply, extends to complex systems in general. When properly defined, the mechanism of stigmergy appears to be nearly ubiquitous and able to illuminate a variety of conceptual problems in a non-trivial manner.

The matter of definition, however, is crucial to a proper understanding and application. Definitions in the literature tend to be vague, restricted in scope and mutually incoherent (Dipple, Raymond, & Docherty, 2014; Shell & Mataric, 2003). Misunderstandings have arisen particularly because of a confusion between the general notion of stigmergy and its specific instantiation in ant algorithms, i.e. the reinforcement with pheromones of frequently traveled paths by virtual "ants". The depositing of pheromone traces is an example of what can be called quantitative, marker-based stigmergy (Parunak, 2006). Stigmergy in the most general sense does not require either markers or quantities. Another, even more common misunderstanding is that stigmergy only concerns groups or swarms consisting of many agents. As we will show, stigmergy is just as important for understanding the behavior of a single individual. The present paper will clarify the meaning of stigmergy, propose an unambiguous definition and summarize its benefits in explaining spontaneous forms of coordination.

Perhaps a last question to conclude this introductory section: if stigmergy is so fundamental, ubiquitous and explanatorily powerful, then why has it taken so long for it to be recognized? The more obvious answer is that the study of termites that gave rise to this conception is a very specific discipline with no evident applications to other sciences. Moreover, the defining publication (Grassé, 1959), appearing in French in a specialized journal, only reached a limited audience. Its reach appears to have widened significantly only after Deneubourg, a researcher active in both French and English and the domains of both social insects and self-organizing systems, started applying the concept outside of its original context.

But why did not someone else come up with this simple and elegant notion? A more fundamental answer is that stigmergic interaction is by definition indirect, while our mind is biased to look for direct causes of the phenomena we observe. If we note that agents act in a coordinated way, our natural inclination is to seek the cause of one agent's behavior directly in another agent's behavior, assuming that there is an immediate communication from the one to other. Failing to find this link, we assume that the agents are driven by the same cause, such as a shared instinct, plan, or leader that controls their behavior. We do not spontaneously consider the option that one agent may drive another agent's behavior only via the indirect route of an unintentional trace left in a passive environment.

This counter-intuitive aspect may explain why there is still so much confusion surrounding the concept of stigmergy and why the breadth and power of its applications have not yet been fully appreciated. I hope that the present paper will help to truly clarify the mechanism and convince researchers to further broaden their search for applications.

2. From etymology to definition

The term "stigmergy" was derived by Grassé from the Greek roots, *stigma*, which means "mark or puncture" (typically referring to the tattoo used to mark slaves) and *ergon* which can mean "work, action, or the product of work". Grassé motivated this derivation by interpreting *stigma* as a goad, prod or spur, i.e. a stinging movement ("piqure" in the original French text) that incites activity. *Ergon* is then the result of previous work responsible for this stimulus or incitement. Thus, (Grassé, 1959) defined stigmergy as "the stimulation of workers by the very performances they have achieved" (from the original English abstract).

However, in a more recent review paper, (Parunak, 2006) proposes a different reading of the Greek etymology that is at least as compelling: if we interpret *stigma* as "mark" or "sign" and *ergon* as "action", then stigmergy is "the notion that an agent's actions leave signs in the environment, signs that it and other agents sense and that determine their subsequent actions". Summarizing, in Grassé's interpretation the product of work (*ergon*) functions as a stimulus (*stigma*) for action; in Parunak's interpretation, action (*ergon*) leaves a mark (*stigma*). While this double interpretation may seem to add to the confusion, it actually provides an elegant illustration of the bidirectional nature of stigmergy. The process described by both Grassé and Parunak is a *feedback loop*, where an action produces a mark which in turn incites an action, which produces another mark and so on (see Fig. 1). In other words, actions stimulate their own continued execution via the intermediary of the marks they make—where a mark is a perceivable effect, trace or product of an action.

This brings me to my own definition:

> *stigmergy is an indirect, mediated mechanism of coordination between actions, in which the trace of an action left on a medium stimulates the performance of a subsequent action.*

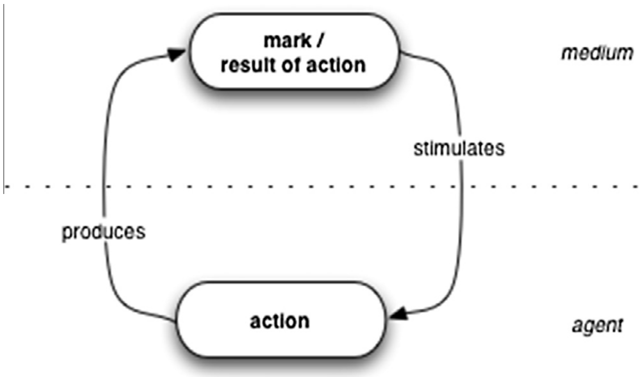

Fig. 1. The stigmergic feedback loop.

3. Basic components of stigmergy

Let us analyze the different terms in this definition and from thereon the conceptual components necessary to build a stigmergic process.

Most primitive is the concept of *action*, which I interpret as a causal process that produces a change in the state of the world. Normally, we assume that an action is performed by an *agent*, which is typically seen as an autonomous, goal-directed system. However, the concept of agent does not appear to be necessary for a definition of stigmergy: the mechanism applies perfectly well to the coordination of actions performed by a single, unspecified agent, in which case there is no need to identify different agents. Moreover, further extensions of the stigmergy concept can even do away with the notion of agent altogether and consider the coordination of "agentless" actions that are merely events or physical processes—such as chemical reactions (e.g. Tabony, 2006). This views fits in with the ontology of action (Heylighen, 2011; Turchin, 1993), which sees action as the primitive element from which all other concepts are derived. The concept of agent remains useful, though, in cases where we wish to distinguish different agents able to perform different actions.

As causal processes, actions have an antecedent or cause and a consequent or effect. In simple agent-based models used in artificial intelligence the antecedent is usually called *condition* and the consequent simply *action*. The condition specifies the state of the world in which the action occurs, while the action specifies the subsequent transformation of that state. The causal relation is represented as a "production rule", which consists of the simple relation:

condition → *action*

It is to be read as:

IF *condition* holds, THEN perform *action*

For example, a thermostat obeys the production rule: IF the temperature is below the goal temperature, THEN switch on the heating (Heylighen & Joslyn, 2003). While this reading seems to imply an elementary cognitive process of sensation or perception, to ascertain whether the condition holds, the notation is equally applicable to agentless, physical processes. For example, consider the following chemical reaction:

$$\mathrm{NaOH} + \mathrm{HCl} \rightarrow \mathrm{NaCl} + \mathrm{H_2O}$$

The first part represents the necessary condition for the reaction to occur: an NaOH molecule and an HCl molecule must be simultaneously present. The second part of the reaction represents the product or result of the reaction: the formation of an NaCl molecule together with an $\mathrm{H_2O}$ molecule. More generally, we can write such reactions in the following form (where the "+" operators represents a conjunction of the conditions or resources necessary for the process to take place):

$$a + b + \cdots \rightarrow x + y + \cdots$$

Chemical Organization Theory (Dittrich & Fenizio, 2007) is a formal framework that shows how collections of such simple reactions tend to become coordinated by acting on a shared medium, the "reaction vessel", where they produce an evolving trace expressed by the concentrations of the different "molecules" (a, b, ...). This coordinated pattern of activity defines an "organization": a self-sustaining, dynamic network of interacting "molecules". While the inspiration for this model comes from chemistry, it is equally applicable to other kinds of abstract actions in which the agent is ignored—such as political and economic interactions (Dittrich & Winter, 2005, 2008).

According to our definition, the action part of a rule produces a change in the state of the world. This means that it creates a new condition, which may activate another condition → action rule and thus a new action. For example, the thermostat, by switching on the heating, will eventually produce the new condition "temperature high enough", which in turn will trigger the new action "switch heating off" (Heylighen & Joslyn, 2003). The NaCL molecule may react with another molecule in the solution and thus produce yet another compound. This triggering of an action by a previous action via the intermediary of its result is precisely what Grassé defined as stigmergy. Yet, the way we arrived at this notion is so simple and general that it merely requires a minimal assumption of causality. In the next sections, we will need to explain how such a simple mechanism can produce such rich and unexpected phenomena.

First, we should note that the causal relation does not need to be deterministic: in general, the condition is neither necessary nor sufficient for the action to occur. According to Grassé's definition, the condition merely *stimulates* the performance of the action. This means that the presence of the condition makes the performance of the action more probable. Formally:

$$P(\text{action}|\text{condition}) > P(\text{action}),$$

where $P(A)$ is the general probability of A occurring and $P(A|B)$ the conditional probability of A occurring given that B is the case.

Note that in some cases, a condition may on the contrary *inhibit* an action, i.e. make it less likely. For example, the presence of a red light makes it less likely that someone would cross the street. In that case, we may still keep to the definition of "stimulation" above, simply by considering the opposite or negation of that condition as the stimulus: e.g. the disappearance of a red light makes it more likely that someone would cross the street.

We must now introduce another core component of stigmergic activity, the *medium*. The medium is that part of the world that undergoes changes through the actions and whose states are sensed as conditions for further actions. The medium is a non-trivial entity, since many aspects of the world are either not affected by actions, or not perceivable as conditions for new actions. For example, while I can clearly see the clouds in the sky, no matter how hard I try, I cannot change their position. Vice-versa, I have the power to throw a rock in the sea, but I cannot see where that rock will end up. In either case, there is no basis for a stigmergic chain of actions triggering further actions. On the other hand, I can both perceive and affect the arrangement of sand on a beach and this allows me to build an intricate sand castle via a coordinated sequence of condition → action pairs. Neither the sea nor the sky is a stigmergic medium, but the beach is.

Note that most authors (e.g. Parunak, 2006) use the term "environment" for what I call "medium". This term is much less accurate, though. First, as noted, the environment is not in general both perceivable and controllable. Second, the environment normally denotes everything *outside* the system or agent under consideration. However, stigmergy can also make use of an *internal* medium. For example, different physiological processes in the body communicate via the release of hormones in the bloodstream (medium). This communication is indirect: e.g. the liver does not directly send a message to the brain; both merely "read" the hormonal messages deposited in the blood that irrigates both. More generally, many aspects of the agent's own state, such as the agent's position, speed and orientation, belong to the medium, since they are controllable and perceivable by self and others.

Finally, if we conceive the environment as that part of the world that interacts with an agent, then different agents live in different environments or "Umwelts": not all phenomena perceivable or controllable by one agent are similarly perceivable and controllable by another agent. When we consider stigmergic coordination between different agents, we need to define the medium as that part of the world that is controllable and perceivable by all of them. This is necessary to ensure that the different agents can interact via the medium. The role of the medium is to allow interaction or communication between different actions and thus, indirectly, between the agents that perform the actions. It is this *mediating* function that underlies the true power of stigmergy (Heylighen, 2006).

A final component of a stigmergic system is the mark or *trace*, i.e. the perceivable change made in the medium by an action, which may trigger a subsequent action. I prefer the term "trace" because it can denote an unplanned or even undesired side effect of the action, unlike a "mark" which is normally made intentionally. Some forms of stigmergy rely on intentionally made signs or signals ("markers"), but in the most general situation, this is not the case. The trace is a consequence of the action and as such, it carries information about the action that produced it. We might see the trace as a message, deposited in the medium, through which the pattern of activity communicates with itself, while maintaining a continuously updated "memory" of its achievements.

From the point of view of an individual agent, on the other hand, the trace is a *challenge*: a situation that incites action, in order to remedy a perceived problem or shortcoming, or to exploit an opportunity for advancement (Heylighen, 2012). We may assume may here that agents have a minimal form of intentionality. This means that their actions are not random, like the mutations that underlie biological evolution, but directed toward the agent's preferred state or *goal*. Reaching a far-away goal, however, requires more than a minimal intelligence: this will typically necessitate a complex scheme of actions, performed according to a specific order or logic. The difficulties involved in problem solving, planning and project management may remind us that there is no obvious way to go from elementary actions to complex activity schemes. This brings us to the problem of coordination (Crowston, 1997), which stigmergy appears to solve.

4. Coordination

According to the Oxford Dictionary, *coordination* can be defined as:

> the organization of the different elements of a complex body or activity so as to enable them to work together effectively.

In the case of stigmergy, the "elements" are the different actions or agents. "Effectively" means that they achieve an intended effect or goal. "Working together" means that the actions are harmonious or synergetic, the one helping rather than hindering the other. "Organization" can be defined as structure with function (Gershenson & Heylighen, 2005). The function is the achievement of the intended effect. A "structure" consists of distinct elements (the actions or agents) that are connected in such a way as to form a coherent whole. This brings us to focus on the *connections* that integrate the actions into a synergetic, goal-directed whole.

According to coordination theory (Crowston, 1997), we can distinguish the following fundamental dependencies or connections between actions or processes:

(1) one action can be *prerequisite* for the next action: the product or output of the first is a necessary condition or input for the second. This determines the *sequential* organization of the process, or *workflow*, where activity moves step-by-step through a sequence of tasks (what needs to be done next?).

(2) two actions can require the same condition (input) and/or contribute to the same effect or goal (output), i.e. they are performed in *parallel*. This determines the *division of labor* between agents (who is to do what?).

Effective coordination means that the right actions are performed by the right agents at the right time and place. Let us consider the building of a house as an activity that requires coordination between its different tasks. The task of laying electricity obviously can only be performed once the windows and roof are installed. Roofing is therefore prerequisite for laying electricity and the electricians will have to wait until the roofers are finished. Plastering the interior walls, on the other hand, can only be done after the electrical cables and outlets have been dug into the walls. This implies the sequence: roofing → laying electricity → plastering. On the other hand, plumbing and laying electricity can be performed simultaneously or in parallel, since they both require roofing and are prerequisite for plastering, but are otherwise independent of each other. The dependencies or connections between these different processes can be represented in the following "workflow" diagram (Fig. 2).

This diagram represents only a small part of the complex of activities that is necessary to construct a building. Construction work and other complex activities are normally planned in detail beforehand, using tools such as project schedules and GANTT charts, to specify the dependencies between the different tasks. This planning is necessary to make sure that the work is efficiently performed, by avoiding situations such as the plasterers turning up when the plumbing is still going on so that they cannot start their work. The plan will normally specify the beginning and end of all the actions as well as the agents that are to perform them and possibly the places or resources that the agents need to access. If everybody keeps to this plan, the plasterers will show up on the exact time and place that the plumbers are supposed to have finished their work.

The problem with planning, of course, is that there will always be unforeseen contingencies, such as the plumbers needing an extra day to finish their work, or, on the contrary, finishing two days early. In both cases, the work is performed less efficiently than it could be, either because the plasterers need to go home because the plumbing is not ready yet, or stay home waiting for the work to finish when they could already have started. Contingencies disturbing carefully laid-out plans can have even worse results, as illustrated by the following joke:

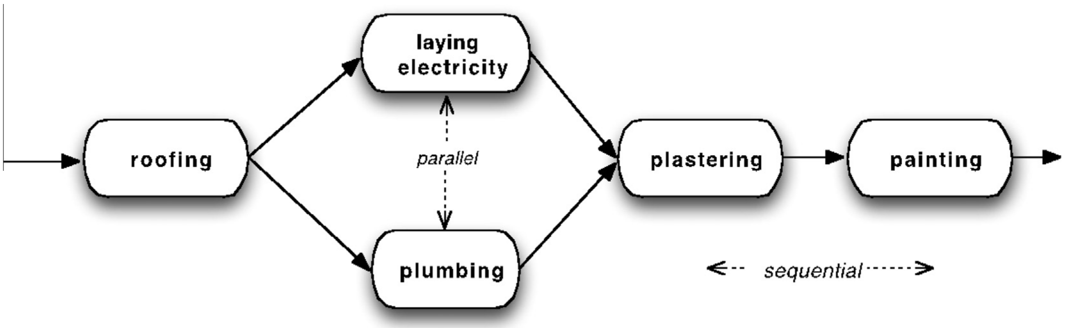

Fig. 2. Parallel and sequential connections between tasks.

> A pensioner watches two city workers busy in the municipal park. The one digs a series of deep holes at regular intervals. The other one then shovels the mounds of earth carefully back into each hole and flattens the soil. The pensioner asks him: “Isn't that a waste of effort what you are doing?”, to which the worker replies: “No, we always work this way and it is very efficient. It is just that the third guy who plants the trees did not show up today.”

One way to deal with such contingencies is to let the agents communicate about their work. For example, the plumbers finishing early or late could call the plasterers to warn them about the different finishing date. However, this assumes that agents know all other agents that depend on their work and have a general notion of what these dependencies are so that they can improvise or reschedule their activity in the light of the new information. With complex activities, this is tricky and can easily lead to misunderstandings or confusions that make things worse. An alternative is that all agents report to a supervisor, who keeps track of the plan, reschedules if need be and warns everyone involved of the changes. However, such central controller becomes a bottleneck that is even more sensitive to disturbances, creating the risk that the whole plan falls apart because the supervisor is not available to pass on reschedules.

5. The benefits of stigmergy

How does stigmergy solve the problem of coordination? In the examples above, the different agents would regularly check the situation at the work site and as soon as they encounter the right conditions, they would start their work. For example, once the plumbers observe that the roof and windows are in place, they would start plumbing. Simultaneously but independently, the electricians would do their job. The plasterers would begin as soon as *both* the plumbing and the electricity are finished. On the other hand, the municipal worker would fill a hole only on the condition that it contains a tree.

While this approach may seem natural for termites, who are anyway all wandering around their nest building site, you might wonder whether it would not be inefficient to demand that specialized workers visit the site every day while they are not yet needed. However, this can be easily tackled with modern technology, by providing a website on which the state of the work is registered in real time. In this way, the plumbers can see immediately whether they are needed, without losing time traveling to the site. The website plays the role of a medium providing special markers to guide the execution of the work—similar to the pheromones used by ants.

This is in essence how a community of programmers residing in different parts of the world collaboratively develop a complex suite of open-source software: they regularly check their shared website for new modules, updates, requests for features, or postings of bugs. They address these challenges by writing additional code or suggesting solutions, posting these results in turn on the website for others to see and to elaborate (Bolici et al., 2009; Heylighen, 2007; Heylighen, Kostov, & Kiemen, 2013). Such open source development has proven to be at least as effective as the traditional planned software development performed in large corporations (Weber, 2004), without requiring any central supervision or other complicated arrangements (Raymond, 1999).

Perhaps the only disadvantage compared to a perfectly designed and executed plan, is that the stigmergic approach does not guarantee an optimal use of the “workforce”. While the roofers are working, the plumbers must either wait or perform another task. If they are busy with another task, there is no guarantee that it will be finished exactly when the roofers finish their job. This suboptimal use of workers can be minimized by creating a pool of available workers (much) larger than needed for this particular job, so that together they can keep track of several jobs in parallel. Assuming that the tasks do not all start at the same time, there would always be some workers available for any job that opens up, without requiring workers to wait long times in between jobs, or tasks to wait long for workers to execute them. This is the approach underlying the job ticketing systems used in call centers (Heylighen & Vidal, 2008; Orrick, Bauer, & McDuffie, 2000), but also the one used by ants and termites. It explains why the best-known applications of stigmergy typically rely on “massive parallelism”, i.e. many agents active simultaneously

(Manderick & Moyson, 1988). In the case of software development, there is no particular time at which a coding job has to be done, so programmers can be flexible in deciding when to produce an improvement suggested by the stigmergic repository of tasks and products.

With such a stigmergic organization, no conflicts between instructions and reality arise, no needless delays occur and no effort is wasted—whatever the contingencies that may disturb the plan. Moreover, this solution is perfectly robust and independent of any errors in communication or control. It also does not depend on the number of agents, tasks, or dependencies between tasks. This allows it to scale up to tasks of indefinite sizes. The only requirements are that the agents can recognize the right conditions to start their work and that they can all access the medium in which these conditions are registered. In summary, stigmergy provides an extremely simple and reliable solution to a problem that is potentially unlimited in complexity.

Compared to traditional methods of organization, stigmergy makes absolutely minimal demands on the agents. In particular, in stigmergic collaboration there is *no need for*:

- **planning or anticipation**: agents only need to know the present state of the activity; the overall goal, next step or end result is irrelevant for their present work. In Wikipedia, there is no plan specifying which information should be added to the encyclopedia when.
- **memory**: agents do not need to remember their previous activity; no information about the state of the work needs to be stored anywhere except in the medium.
- **communication**: no information needs to be transferred between the agents, except via the work done in the medium; there is in particular no need for the agents to negotiate about who does what.
- **mutual awareness**: each agent works independently; it does not even need to know that others participate. For example, contributors to Wikipedia generally do not know each other or communicate with each other.
- **simultaneous presence**: there is in general no need for the agents to be present at the same time or at the same place; tasks are registered in the medium so that they can be picked up by agents whenever and wherever they are available. That is how worldwide communities can collaborate on a single software project.
- **imposed sequence**: actions are performed automatically in the right order, since an action will not be started until the right condition is in place; the workflow emerges spontaneously, as the completion of one task triggers the initiation of the next task(s).
- **imposed division of labor**: each agent will only perform the actions for which it has the required competence, i.e. for which it possesses adequate condition-action rules; normally, the more "confident" the agent is about the right action (i.e. the stronger the connection between condition and action), the more it will be stimulated by the condition and the quicker it will be to start the job; in this way, tasks are automatically assigned to the most competent agents (Heylighen & Vidal, 2008).
- **commitment**: agents do no need to commit to a particular task (in contradiction to what (Jennings, 1993) claims about multi-agent coordination); an agent decides on the spot what work it should do, depending on opportunity and other contingent conditions; an agent that quits or otherwise becomes unavailable is automatically replaced by another one.
- **centralized control or supervision**: errors or perturbations are automatically corrected, as they merely create a new condition stimulating new actions to deal with the challenge; the activity is *self-organizing*. For example, bugs in open-source software are spotted by users and resolved by other contributors.

6. Self-organization through negative feedback

This last point deserves an elaboration. Our assumption is that agents are individually goal-directed. Cybernetics has shown how goal-directedness emerges from negative feedback: perceived deviations from the goal are compensated by counteractions (Heylighen & Joslyn, 2003; Rosenblueth, Wiener, & Bigelow, 1943). This most basic mode of steering is also called *error-controlled regulation*: whatever the origin of the deviation or "error", once it is sensed, its effect is suppressed by a compensatory action. This control mode does not require any anticipation (feedforward), memory, or understanding of what caused the deviation. To efficiently control the effect, it is sufficient that the agent is able to exert an influence in the "opposite direction" of any deviation, independently of its underlying cause (Gershenson & Heylighen, 2005).

This mechanism is well understood for individual agents (Powers, 1973). Stigmergy illustrates how it can be extended to several interacting agents. Imagine a group of non-communicating agents (e.g. ants, or people who do not speak a common language) pushing a large obstacle out of the way across an irregular terrain. Individually, each agent will correct its course based on the perceived movement of the load: e.g. if it shifts too much to the left, the agent will push more toward the right. It does not matter whether the deviation was caused by a hole in the surface, a sudden gust of wind, or the misdirected action of another agent. The overall movement will be determined by the sum of the actions performed by all the individuals. As long as the agents push in generally the same direction, it is irrelevant who did what. The agents can work perfectly independently—perhaps even without knowing that someone else is pushing too—while still producing a coordinated movement.

A similar mechanism is probably involved in bodily coordination, where different muscles, joints, tendons and bones need to work together effectively. A complex movement, such as balancing on one foot, requires the simultaneous contraction or relaxation of many muscles. If the actual position deviates from the desired one, individual muscles will adjust their level of contraction, upward or downward, depending on the perceived direction of the deviation. Since all muscles have simultaneous access to the shared "trace" of their activity (the present position), they also act on it simultaneously. The sum of their actions will normally be sufficient to achieve the desired effect. If not—that is, if one of the muscles contributes too little or too much—either that muscle itself or one or more complementary muscles will react in order to correct the error. The muscles do not need to communicate with each other about their intention to shift position: their actions are instantly coordinated by their simultaneous effect on the trace, just like the actions of the agents pushing the load are coordinated by the perceived position of the load.

To ensure effective collective action, the only assumption we need to add to individual error-controlled regulation is that the goals of the agents are not contradictory—i.e. that error decrease for one agent does not equal error increase for another agent—because then the agents will be involved in a tug-of-war of opposing counteractions. However, note that the goals do not need to be *identical* for coordination to occur: imagine that one group of agents pushes the obstacle to the east, while another group pushes to the north. The net effect is that the load will move northeast, satisfying both groups. It is only when one group pushes eastward and another group westward that a conflict arises, without possibility for a compromise. In this two-dimensional example, the probability for conflict still seems large. However, the larger the number of aspects, components or degrees of freedom of the problem situation, the more freedom there is for agents to focus on different goals without getting in each others' way.

This independence of goal setting is what underlies the automatic division of labor: each agent spontaneously focuses on the task that it deems most important (and for which it is in general most competent). Thus, a variety of agents together can potentially tackle very complex problems that require the achievement, in sequence or in parallel, of many different partial objectives. We may assume that agents have acquired their condition-action rules (and thus their implicit goals) through natural selection of instinctual behavior or differential reinforcement of learned behavior. This means that their condition-action rules are generally appropriate to the local environment, including the other agents with which they regularly interact. Rules that are frequently in conflict with the rules of other agents or the constraints of the environment are likely to be eliminated eventually (Heylighen, 2008).

Therefore, it is plausible to assume that even very different agents, e.g. belonging to different species in an ecosystem, follow rules that are potentially synergetic (Corning, 1995; Heylighen, 2013). Stigmergy appears to be a prime mechanism through which this synergy is realized, by coordinating initially independent actions into a harmonious whole. Thus, the group of agents can achieve much more substantial results collectively than they would if they would work alone. It is this emergence of global order out of local actions that constitutes the hallmark of self-organization (Heylighen, 2001).

7. Self-organization through positive feedback

Stigmergy exhibits another fundamental "signature" of self-organization (Heylighen, 2001; Theraulaz & Bonabeau, 1999): *positive feedback*. Error-controlled regulation assumes negative feedback: the *reduction* of deviations away from the goal. However, goal-directed action can also make use of positive feedback: the *amplification* of movements toward the goal. In the cybernetic perspective, changes in the situation not controlled by the agent are interpreted as perturbations, since they move the system away from a previously achieved goal state (Heylighen & Joslyn, 2003; Maturana & Varela, 1980). However, as long as no final goal is reached—which is the default situation in long-term, on-going projects, such as building, extending and maintaining a termite hill—such contingent events may as well facilitate as hinder the further movement toward the goal. When they hinder, we will call them *disturbances*; when they facilitate, we may call them *affordances* (Gibson, 1977). In the most general case, we may call them *diversions*, since they divert action from its on-going course, whether in a positive, negative or neutral way (Heylighen, 2012). Assuming that agents are goal-directed, we may infer that they will counteract the disturbances and reinforce or build upon the affordances.

Let us illustrate these notions with the paradigmatic case of termites erecting a pillar as part of their nest construction. A bit of mud that is accidentally dropped in a particular place, either by a termite, the wind or a passing bird, constitutes a diversion. In this case, the diversion constitutes an affordance, since it provides a foundation on which a taller mud structure can be erected. Stigmergic stimulation will lead termites to add mud to the emerging heap, rather than to the flat surfaces surrounding it. The taller the heap grows, the stronger the stimulus it will exert on termites passing by and therefore the faster its further growth. This positive feedback loop results in an accelerated exploitation of the opportunity, diverting effort away from less promising alternatives and thus efficiently allocating agents and resources to the most productive activities.

Suppose now that a fragment of the thus erected column breaks off. This constitutes a negative diversion, i.e. a disturbance. In this case, the perception of the missing mud will stimulate the termites to fill the hole with new mud, thus counteracting the deviation from the ideal column shape. Finally, a neutral diversion may arise, such as a breeze blowing some termites off-course, so that they end up near a different column than the one they were heading to. While making the activity deviate from the original course, this event neither facilitates nor hinders the work. Therefore, it will be neither counteracted nor reinforced.

A similar dynamics occurs in Internet communities centered on a particular forum, website or page. More activity on a particular site tends to produce more interesting traces, such as discussions, wiki pages, or comments. These attract more people and therefore more contributions, which in turn incite more activity. Vice versa, less activity reduces the level of interest for the products of that activity and thus the number of potential further contributions. Thus, web communities and their activities are subject to a positive feedback, where the initially most promising "projects" grow very quickly, while the less promising ones dwindle, losing the competition with the others and potentially disappearing altogether. In this way, work is divided efficiently across a wide variety of projects, ensuring that the most promising ones quickly "take off" without dissipating too much energy in less promising ones.

8. Conclusion

Our theoretical analysis of stigmergy has illustrated how wide-ranging and fundamental this mechanism is. Virtually all evolved processes that require coordination between actions seem to rely at some level on stigmergy, in the sense that subsequent actions are stimulated by the trace left by previous actions in some observable and manipulable medium. The trace functions like a registry and map, indicating which actions have been performed and which still need to be performed. It is shared by all agents that have access to the medium, thus allowing them to coordinate their actions without need for agent-to-agent communication. It even allows the coordination of "agentless" actions, as investigated e.g. by Chemical Organization Theory (Dittrich & Fenizio, 2007).

Thus, stigmergy can be seen as a fundamental mechanism of *self-organization*: it allows global, coordinated activity to emerge out of local, independent actions. Like self-organization in general, stigmergy relies on *feedback*: action elicits action, via the intermediary of the trace. This feedback is typically positive, in that actions intensify and elaborate the trace, thus eliciting more intense and diverse further actions. The resulting virtuous cycle explains in part why stigmergic organization is so surprisingly effective, enabling the construction of complex structures—such as a termite hill, a network of trails, or a world encyclopedia—in a very short time, even when starting from scratch. When necessary, feedback can also be negative: errors, disturbances or "overshoots" that make the trace deviate from its ideal shape will elicit actions that correct the deviation.

In conclusion, the concept of stigmergy allows us to explain a broad variety of ill-understood phenomena of spontaneous coordination, in which some kind of apparently intelligent activity emerges out of simple causal processes, without requiring direct communication or central supervision. Because of its non-intuitive nature, its power of explanation is as yet poorly recognized. I hope that the present paper will inspire further researchers to pick up the thread and examine stigmergic mechanisms in an ever-wider range of human and non-human activities. This should not only clarify deep theoretical issues, but suggest a variety of as yet unimagined practical applications.

References

Beckers, R., Holland, O. E., & Deneubourg, J. L. (1994). From local actions to global tasks: Stigmergy and collective robotics. *Artificial Life IV, 181*, 189.

Bolici, F., Howison, J., & Crowston, K. (2009). Coordination without discussion? Socio-technical congruence and stigmergy in free and open source software projects. In *Conjunction with the international conference on software engineering, Vancouver, BC, Canada*.

Bonabeau, E., Dorigo, M., & Theraulaz, G. (1999). *Swarm intelligence: From natural to artificial systems*. Oxford: Oxford University Press.

Camazine, S., Deneubourg, J. L., Franks, N. R., Sneyd, J., Theraula, G., & Bonabeau, E. (2003). *Self-organization in biological systems*. Princeton: Princeton University Press.

Corning, P. A. (1995). Synergy and self-organization in the evolution of complex systems. *Systems Research, 12*(2), 89–122.

Crowston, K. (1997). A coordination theory approach to organizational process design. *Organization Science, 8*(2), 157–175.

Deneubourg, J. L. (1977). Application de l'ordre par fluctuations a la description de certaines etapes de la construction du nid chez les termites. *Insectes Sociaux, 24*(2), 117–130.

Deneubourg, J. L., Goss, S., Franks, N., Sendova-Franks, A., Detrain, C., & Chrétien, L. (1991). The dynamics of collective sorting robot-like ants and ant-like robots. In *Proceedings of the first international conference on simulation of adaptive behavior on from animals to animats* (pp. 356–363).

Deneubourg, J. L., Theraulaz, G., & Beckers, R. (1992). Swarm-made architectures. In *Toward a practice of autonomous systems: Proceedings of the first European conference on artificial life*.

Dipple, A., Raymond, K., & Docherty, M. (2014). General theory of stigmergy: Modelling stigma semantics. *Cognitive Systems Research, 31–32*, 61–92.

Dittrich, P., & Winter, L. (2005). Reaction networks as a formal mechanism to explain social phenomena. In *Proc. fourth int. workshop on agent-based approaches in economics and social complex systems (AESCS 2005)* (pp. 9–13). <http://users.minet.uni-jena.de/~dittrich/p/DW2005.pdf>.

Dittrich, P., & Fenizio, P. S. di (2007). Chemical organisation theory. *Bulletin of Mathematical Biology, 69*(4), 1199–1231.

Dittrich, P., & Winter, L. (2008). Chemical organizations in a toy model of the political system. *Advances in Complex Systems, 11*(04), 609.

Dorigo, M., Bonabeau, E., & Theraulaz, G. (2000). Ant algorithms and stigmergy. *Future Generation Computing Systems, 16*(8), 851–871.

Dron, J., Boyne, C., & Mitchell, R. (2001). *Footpaths in the stuff swamp*. <http://files.eric.ed.gov/fulltext/ED466580.pdf>.

Elliott, M. A. (2007). *Stigmergic collaboration: A theoretical framework for mass collaboration*. PhD thesis, Victorian College of the Arts, University of Melbourne. <http://eprints.infodiv.unimelb.edu.au/00003977/01/elliott_phd_pub_08.10.07.pdf>.

Gershenson, C., & Heylighen, F. (2005). How can we think complex? In K. Richardson (Ed.), *Managing organizational complexity: Philosophy, theory and application* (pp. 47–61). Charlotte: Information Age Publishing.

Gibson, J. J. (1977). The theory of affordances. In *Perceiving, acting and knowing: Toward an ecological psychology* (pp. 67–82). Boston: Houghton-Mifflin.

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.

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. (2001). The science of self-organization and adaptivity. *The Encyclopedia of Life Support Systems, 5*(3), 253–280.

Heylighen, F. (2006). Mediator evolution: A general scenario for the origin of dynamical hierarchies. In *Worldviews, science and us*. Singapore: World Scientific.

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

Heylighen, F. (2011). Self-organization of complex, intelligent systems: An action ontology for transdisciplinary integration. *Integral Review*.

Heylighen, F. (2012). *A tale of challenge, adventure and mystery: Towards an agent-based unification of narrative and scientific models of behavior. ECCO working papers no. 2012-06*. Brussels, Belgium.

Heylighen, F., & Joslyn, C. (2003). Cybernetics and second-order cybernetics. In R. A. Meyers (Ed.). *Encyclopedia of physical science and technology* (3rd ed., Vol. 4, pp. 155–169). New York: Academic Press.

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. (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). Lehmanns Media.

Heylighen, F. (2013). Self-organization in communicating groups: The emergence of coordination, shared references and collective intelligence. In À. Massip-Bonet & A. Bastardas-Boada (Eds.), *Complexity perspectives on language, communication and society* (pp. 117–149). Berlin: Springer.

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

Holland, O., & Melhuish, C. (1999). Stigmergy, self-organization and sorting in collective robotics. *Artificial Life, 5*(2), 173–202.

Jennings, N. R. (1993). Commitments and conventions: The foundation of coordination in multi-agent systems. *The Knowledge Engineering Review, 8*(03), 223–250.

Kassabalidis, I., El-Sharkawi, M., Marks, R., Arabshahi, P., & Gray, A. (2001). Swarm intelligence for routing in communication networks. In *Global telecommunications conference, 2001. GLOBECOM'01. IEEE* (Vol. 6, pp. 3613–3617).

Kennedy, J. (2006). *Swarm intelligence. Handbook of nature-inspired and innovative computing* (pp. 187–219). Heidelberg: Springer.

Manderick, B., & Moyson, F. (1988). The collective behavior of ants: An example of self-organization in massive parallelism. In *Proceedings of the AAAI spring symposium on parallel models of intelligence, Stanford University*.

Marsh, L., & Onof, C. (2008). Stigmergic epistemology, stigmergic cognition. *Cognitive Systems Research, 9*(1), 136–149.

Maturana, H. R., & Varela, F. J. (1980). *Autopoiesis and cognition: The realization of the living*. Dordrecht: Reidel.

Orrick, D., Bauer, J., & McDuffie, E. (2000). Remedial help desk 101 at Florida State University. In *Proceedings of the 28th annual ACM SIGUCCS conference on user services: Building the future* (pp. 222–226). New York: ACM.

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.

Powers, W. T. (1973). *Behavior: The control of perception*. Chicago: Aldine Transaction.

Raymond, E. (1999). The cathedral and the bazaar. *Knowledge, Technology and Policy, 12*(3), 23–49.

Ricci, A., Omicini, A., Viroli, M., Gardelli, L., & Oliva, E. (2006). Cognitive stigmergy: A framework based on agents and artifacts. In *Third international workshop on environments for multiagent systems (E4MAS 06)*.

Robles, G., Merelo, J. J., & Gonzalez-Barahona, J. M. (2005). Self-organized development in libre software: A model based on the stigmergy concept. In *Proceedings of the 6th international workshop on software process simulation and modeling (ProSim), 16*.

Rosenblueth, A., Wiener, N., & Bigelow, J. (1943). Behavior, purpose and teleology. *Philosophy of Science, 10*(1), 18–24.

Shell, D., & Mataric, M. (2003). On the use of the term "'stigmergy'." In *Proceedings of the second international workshop on the mathematics and algorithms of social insects*.

Sumpter, D. J., & Beekman, M. (2003). From nonlinearity to optimality: Pheromone trail foraging by ants. *Animal Behaviour, 66*(2), 273–280.

Susi, T., & Ziemke, T. (2001). Social cognition, artefacts and stigmergy: A comparative analysis of theoretical frameworks for the understanding of artefact-mediated collaborative activity. *Cognitive Systems Research, 2*(4), 273–290.

Tabony, J. (2006). Microtubules viewed as molecular ant colonies. *Biology of the Cell, 98*, 603–617.

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

Turchin, V. F. (1993). The cybernetic ontology of action. *Kybernetes, 22*(2), 10–30. http://dx.doi.org/10.1108/eb005960.

Valckenaers, P., Van Brussel, H., Kollingbaum, M., & Bochmann, O. (2006). Multi-agent coordination and control using stigmergy applied to manufacturing control. *Multi-Agent Systems and Applications*, 317–334.

Weber, S. (2004). *The success of open source*. Cambridge: Cambridge University Press.