ECCO Working paper 2011-12

Stigmergy as a generic mechanism for coordination: definition, varieties and aspects

Francis Heylighen
Evolution, Complexity and Cognition group
Vrije Universiteit Brussel

Introduction

The concept of *stigmergy* was introduced by the French entomologist Pierre-Paul 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" [Bonabeau & Théraulaz, 19], 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 are the pheromone trails left by ants that come back from a food source []. 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 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* [Deneubourg - Insectes Sociaux, 1977; Camazine, Deneubourg, et al., 2002], 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, Théraulaz & Beckers, 1992]. The general ability to tackle complex problems exhibited by such self-organizing multi-agent collectives became known as *swarm intelligence* [Bonabeau, Dorigo, & Théraulaz, 1999]. 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, Théraulaz, 2000] and the optimization of packet routing along communication networks [Kassabalidis, MA El-Sharkawi, RJM]. A similar stigmergic mechanism was recently recognized in molecular biology [Tabony, 2006], in the self-organization of the microtubules that support many functions in the cell. These microscopic tubes change shape and move by absorbing tubulin proteins at one end, and releasing them at the other end. The "trail" of tubulin left at the shrinking end attracts the growing ends of other microtubules, resulting in the formation of a coherent "wave" of microtubules moving in the same direction.

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 [Deneubourg, Goss, Franks, Sendova-Franks, 1991; Beckers, Holland, Deneubourg, 1994]. These robotic implementations inspired the application of stigmergic models to problems of coordination and control in manufacturing [Valckenaers, Van Brussel, Kollingbaum..., 2001]. 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 [Heylighen, 1999; Gregorio, 2001; Dron, Boyne & Mitchell, 2001]. 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]. But it quickly became clear that human collaboration does not need computer support to profit from stigmergy [Parunak, 200; Elliot, 200]. Probably the most famous example of stigmergic self-organization is the "invisible hand" of the market: buying and selling actions leave a trace in the price of the transacted commodities, which 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 [March & Onof, 2006; Susi & Ziemke, 2001; Ricci, ...].

It is clear that since 1990, the concept of stigmergy has undergone a rapid diffusion across an ever-growing number of application domains. While the number of publications that mention the term “stigmergy” appears to have remained roughly constant at about 1 per year in between 1960 and 1990, the following years witnessed an impressive exponential growth in that number: from 3 in 1991, via 70 in 1999, to about 500 in 2006 (as measured via a search on scholar.google.com). It seems likely that this is just the beginning of an explosive development, and that the application of stigmergy in particular to human affairs opens the way to a virtually limitless expansion across the various scientific, technological and human disciplines that study systems, cognition, and behavior.

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 [] noted about human institutions, that the more difficult issue is to find examples where stigmergy does *not* apply, extends to complex systems in general, and in particular to systems that exhibit some form of cognition, cooperation, or organization that is the result of evolution. 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 are often vague, confusing and mutually incoherent [Shell & Mataric]. 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 may be called quantitative, marker-based stigmergy [Parunak, 200]. 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 next section of this paper will clarify the meaning of stigmergy, propose an unambiguous definition, and summarize its benefits in explaining spontaneous forms of coordination. We will then go into greater depth concerning the different components and aspects of the mechanism. This will allow us to situate and classify the apparently very different forms of stigmergy, while remaining focused on their common core. The final sections will apply these concepts to some outstanding issues in the theories of cognition, cooperation and evolution. I will argue in particular that several fundamental paradigms in cognitive science—basically those, like situated and distributed cognition, that focus on the role of environment, and those that see the mind as a "society" of collaborating processes—are reducible to stigmergy.

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], appearing in French in a specialized journal, obviously 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.

Finally, we tend to assume that intelligent organization must be produced by an intelligent agent, as illustrated by Paley's well-known watchmaker argument [Dawkins, blind watchmaker]. Darwin's theory of *natural selection* is a relatively recent conception of a mechanism that can generate organization without presupposing intelligence. However, its counter-intuitive nature may be illustrated by the fact that, in spite of massive empirical evidence, it is still being contended by the "intelligent design" school [Behe]. Another proposed mechanism to generate coordinated activity, *self-organization*, is much more recent, and is still far from being generally accepted and, even less, generally understood. I wish to suggest here that stigmergy is another type of mechanism for generating complex organization and intelligent behavior, which is related to both natural selection and self-organization, but which has some distinct features of its own that may illuminate some outstanding problems with these previously proposed explanations.

The meaning of stigmergy

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 [] defined stigmergy as "the stimulation of workers by the very performances they have achieved" (from the original English abstract).

However, in a recent review paper, Parunak [] 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*) produces 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.). In other words, actions stimulate their own continued execution via the intermediary of the marks they make—where a mark is a perceivable effect or product of an action.

This brings me to my own definition: *stigmergy is an indirect, mediated mechanism of coordination between actions in which a perceived effect of an action stimulates the performance of a subsequent action*.

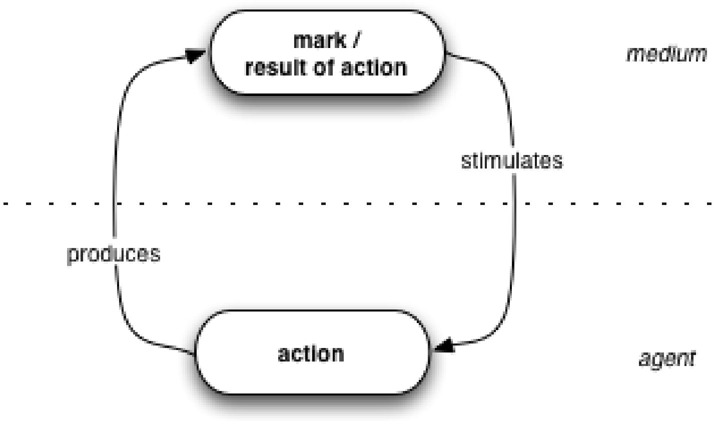

Basic components of stigmergy

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, it appears to me that the concept of agent is not strictly necessary for a definition of stigmergy: as we will see, 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 may 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. (This views fits in with Turchin's [] cybernetic ontology where action is 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 simply as a *condition* → *action* pair, called "production rule" or *production*. It is to be read as "IF *condition* holds, THEN perform *action*". For example, a thermostat will obey the production rule: IF the temperature is below the goal temperature, THEN switch on the heating. 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:

$$NaOH + HCl \rightarrow NaCl + H_2O$$

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 process: the formation of an NaCl molecule together with an H_2O molecule.

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]. 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 complete or deterministic: in general, the condition is neither necessary nor sufficient for the action to occur. According to Grassé's definition of stigmergy, the condition merely *stimulates* the performance of the action. This means that the presence of the condition makes the performance of the action more likely. Formally:

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

where $P(A)$ is the 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]) 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. As we will see, 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.

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. As we will see, 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, or maintains a continuously updated "memory" of its achievements.

Goal-directed action

Let us assume that actions are performed by agents with a minimal form of intentionality or intelligence, i.e. agents whose actions are appropriate to the conditions that trigger them, in the sense that they help the agent to move toward its (implicit or explicit) goals. This is an application of what Dennett [] called the "intentional stance".

This assumption is less strong than it may seem, since natural agents (such as living organisms) have the implicit goal of fitness (i.e. survival and reproduction) built into them by natural selection, while artificial agents (such as thermostats or robots) have their goals specified by their designers. Even natural, non-living objects, such as stones or molecules, can be seen as goal-directed, in the sense that their dynamics can always be modeled as trying to optimize some function of their state [Mesarovic & Takahara] (e.g. potential energy).

This assumption allows us to add a "virtual" component to the stigmergic mechanism, i.e. a component that in a sense only exists for the observer: the *tasks* that the stigmergic system is to perform. Since a stigmergic system does not plan, it generally does not have any awareness or representation of the tasks, jobs or duties that it still has to carry out. But for the outside observer, it may be helpful to use the term "tasks" as shorthand for the actions that are to be performed. A *task* can be defined as an action that:

1) is not yet performed;
2) would contribute to achieving the agents' goals;
3) can be performed on the present medium once the right conditions arise.

This minimal intelligence means that actions are not random or blind, like the mutations that underlie biological evolution, but generally produce some improvement in the agent's situation, i.e. movement closer to the goal. Reaching a far-away goal, however, requires more than a minimal intelligence: this will typically necessitate a complex, coordinated 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 simple or obvious way to go from elementary actions to complex activity schemes. This brings us to the problem of *coordination* [Crowston, 2003], which stigmergy appears to solve.

Coordination

According to the American 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 obviously 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 [Heylighen]. 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, 2003], 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 *allocation of resources* (who receives what?) and 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 components or 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.

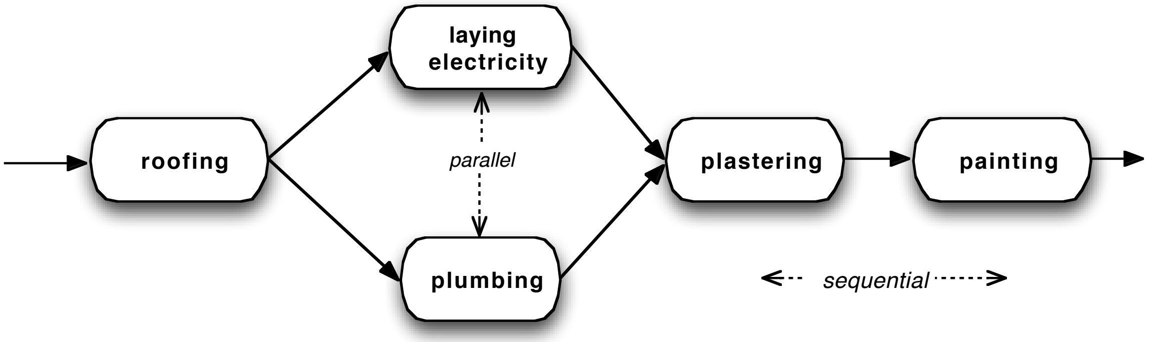

This diagram represents only a small part of the complex of activities that is necessary to construct a building. Construction work and other complex activities are normally planned in detail beforehand, using tools such as project schedules and GANTT charts, to clearly specify the dependencies between 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:

> A pensioner watches two city workers busy in the municipal park. The one digs a series of deep holes at regular distances. The second one then shovels the mounds of earth carefully back into each hole, and flattens the soil. When the pensioner asks him what they are doing, he replies: "Normally we work with a third guy who plants the trees, but today he did not show up..."

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 control 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.

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 not fill any hole unless 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.

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. This is the approach underlying the job ticketing systems used in call centers [Heylighen & Vidal], 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].

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. 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
- **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
- **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
- **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
- **commitment**: agents do no need to commit to a particular task (in contradiction to what Jennings [] claims about multi-agent coordination); what work an agent does is decided on the spot, 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 compensated, as they merely create new conditions stimulating new corrective actions; the activity is *self-organizing*: global organization emerges from local interactions, without any centralized control directing the activity.

Stigmergy as self-organization

This last point deserves a further 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 [Wiener, Bigelow & Rosenblueth, ; Heylighen & Joslyn, 200]. 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 an appropriate compensatory action. This control mode does not require any planning, 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 [Heylighen & Gershenson; de Latil].

This mechanism is well understood for individual agents [e.g. Powers, 19]. 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 towards 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 each individual. 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 may well be involved in bodily coordination, where different limbs and muscles contribute to overall movement of the body. This may be illustrated by Brooks' [] subsumption architecture for the control of the different body parts of many-legged robots, where each limb functions more or less autonomously in helping the robot to move forward, while the only higher-order control imposed is the general direction of movement.

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 that can only end if the stronger subdues the weaker. 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 movement 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 continually in conflict with the rules of other agents or the constraints of the environment are likely to be eliminated eventually [cf. Heylighen, 2007].

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; Wright, Heylighen,]. Stigmergy seems 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, Bonabeau]. It implies in particular that organization arises spontaneously from local activity, without planning, centralization or external control. Problems, contingencies or disturbances will be tackled by the same local action: since there is no plan, there can be no deviation from the plan and therefore no true "error"; everything is contingent and subject to the same, incessant activity of adaptation or improvement.

Affordances, disturbances and feedback

Stigmergy exhibits another fundamental "signature" of self-organization [Bonabeau, Heylighen]: positive feedback. Error-controlled regulation typically assumes negative feedback: the *reduction* of deviations away from the goal. However, goal-directed action can as well make use of positive feedback: the *amplification* of movements towards the goal [Maruyama]. In the traditional cybernetic perspective, changes in the situation not controlled by the agent tend to be interpreted as perturbations, since they move the system away from a previously achieved goal state [Maturana & Varela; Heylighen & Joslyn]. 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 can call them *affordances* [cf. Nelson]. 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. Assuming that agents are implicitly goal-directed, we may infer that they will counteract the disturbances and reinforce or build upon the affordances, i.e. exert a negative, respectively positive, feedback to negative, respectively positive, diversions. Similarly, their reaction to a neutral diversion will be neutral: neither amplifying it nor suppressing it.

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.

The combination of positive and negative feedbacks is typical for complex, adaptive or self-organizing systems [Holland, Heylighen]. It makes the system very flexible, allowing it to energetically act and grow when given the opportunity, while maintaining a stable and robust configuration in the face of disturbances. It also produces differentiation, by amplifying minute differences or chance fluctuations into robust macroscopic structures [Prigogine; Helbing]. Finally, it makes the system intrinsically non-linear, which implies that for the outside observer its evolution is both unpredictable and uncontrollable.

The benefits of parallelism

As noted, (stigmergic) coordination has two aspects: *parallel* and *sequential*. Agents or rules working in parallel simply add their effects together. Because their actions are simultaneous, there is no time to *interact*, i.e. for the one to causally affect the other. Therefore, their total effect is simply the aggregate, superposition or sum of their individual effects. Rules working in sequence, however, by definition interact, since the result of the former affects the performance of the latter. This allows *non-linearity*, i.e. a total effect different from the sum of the individual effects. This total may be larger—in which case there is amplification or positive feedback—, or smaller—which means suppression or negative feedback. As we saw, amplification is useful to exploit affordances, suppression to control disturbances. The combination of parallel and sequential—or linear and non-linear—aspects provides maximal opportunities for an efficient exploration and exploitation of the situation.

Since the power of non-linearity in self-organization is well known [e.g. Bonabeau; Heylighen], it is worth paying special attention here to the less studied benefits of linear or parallel activity. We can distinguish two cases of parallel action: 1) two or more actions are performed on the same object or task, i.e. the same part or aspect of the medium; 2) actions are performed on separate, independent parts of the medium. The first case may be exemplified by termites adding mud to the same pillar, or agents pushing the same load. The second case can be found when two termites add mud to two different pillars, or when the one is busy repairing the nest while the other is collecting food.

This latter case is perhaps most intuitive, since it underlies the mechanism of the *division of labor*. The advantages of the division of labor are well known: it enables specialization, so that each agent can focus on the task it has most expertise with, and thus the task it can perform most efficiently. We already argued that stigmergy tends to automatically allocate tasks to the most competent agents. To maximally benefit from the division of labor, we moreover need to ensure a sufficient diversity in the competencies of the agents [Martens]: the more diverse their expertise, the more likely it is that at least some agent(s) will be particularly competent for a certain task. As a result, a more diverse group of agents will normally be more productive than an equally large, but more homogeneous, group.

This general principle can be illustrated by a classic ecological experiment: if two identical patches of land are seeded with plants that belong either to one or a few species, or to several different species, the more diverse patch will produce more biomass than the more homogeneous one [Naeem et al.; Hector et al. Plant Diversity]. (The overall yield increases with the logarithm of the number of species.) The reason appears to be that plants of different species use the available nutrients in a somewhat different way, thus together being able to exploit the resources more completely ("niche complementarity"), while moreover helping each other through synergies [Hector et al. Plant Diversity]. This is an example of parallel stigmergy where synergetic interaction is mediated by the shared environment (land).

The benefit of diversity is not limited to situations where agents work in different places or perform different tasks. When diverse agents tackle the same problem in parallel, their aggregate solution will in general be better than the one of any single agent or agent type. This phenomenon has been referred to as the "wisdom of crowds" [Surowiecki]. It can be exemplified by the situation where a crowd of people are asked to guess how many beans are contained in a particular jar, or how heavy a particular ox weighs. In such cases, the average of all the guesses is typically much more accurate than any particular guess. The reason is the law of large numbers: if we assume that guesses exhibit a random deviation from the correct answer, then these random deviations tend to cancel each other out when a large number of them are aggregated. Each individual deviation is caused by the limited experience or inaccurate perception of that individual. But when experiences are diverse, the shortcomings of the ones tend to compensate for the shortcomings of the others, providing a more balanced, and therefore accurate, global perception [Heylighen,]. This general observation is confirmed by the multi-agent simulation of a maze-learning task by Johnson []: the average of the choices made by agents with different experiences turned out to provide a better solution than even the best of the individual solutions.

In summary, the greater the variety of agents (or productions) that work in parallel on a particular task or set of tasks, the better we can expect their overall performance to be. Note that parallelism or independence is important for this conclusion to hold [Surowiecki]: if the agents work in sequence, later ones may compensate for the limited experience of earlier ones, but because of non-linear interactions early choices are likely to be amplified. This can lead to a much quicker exploitation of a good solution, but also to the system settling in a far from optimal solution. The creation of pheromone trails by ants, which happens partly in parallel, partly in sequence, illustrates the precarious trade-off between the benefits of parallel/linear approaches (more wide-ranging exploration) and those of sequential/non-linear ones (more efficient exploitation) [Heylighen].

Note also that it is stigmergy that enables such massive parallelism [cf. Manderick & Moyson]. In computing, traditional approaches to parallel processing that rely on central supervision have been notoriously difficult. One reason is that it is very complicated to synchronize the different processes. In a stigmergic setting, no synchronization or supervision is needed, since the result of one process is kept in the medium until another process needs it. For simultaneous processes, it is the medium itself that aggregates their individual results into a global result.

Varieties and aspects of stigmergy

Within the broad category of stigmergic mechanisms, we can distinguish many examples and special cases. To bring some order in these phenomena, it is useful to develop a basic classification of 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 binary (“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 own previous results.

Individual vs. collaborative 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 its environment.

An elegant example discussed by Bonabeau & Théraulaz [] 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, …, S_5. These subsequently perceived conditions or stimuli each trigger a fitting action or response: $S_1 \rightarrow R_1$, $S_2 \rightarrow R_2$, …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, since the present stage of the activity is directly visible in the work already realized.

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 Bonabeau & Théraulaz [] 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 competencies, adding more agents merely increases manpower and therefore the size of the problem that can be tackled, the speed of advance, and 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.

The only complication added 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 [], 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]. Each wolf is attracted to move towards 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.

Quantitative vs. qualitative stigmergy

Quantitative stigmergy [Bonabeau & Théraulaz] 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 my definition of stimulation in terms of conditional probability: the stronger the trace, the larger 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 positive developments.

But quantitative stigmergy can also be 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 good (action) reduces the supply and thus increases the price (quantitative trace resulting from the buying and selling activity). A higher price will normally reduce the probability that someone would buy additional stock of that good. Thus, a higher price reduces demand, which in turn will reduce the price. This form of distributed control [Heylighen] corresponds to the "invisible hand" of the market, which stabilizes prices and efficiently allocates production capacity to the goods that are most in demand [Smith].

Qualitative stigmergy [Bonabeau & Théraulaz] 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 stigmergy) will elicit a corrective action (e.g. writing a new definition).

Different types of inaccuracy, 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 number of different 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. As the one condition becomes more similar to the other, the probabilities become more similar too. In the middle, the two probabilities may become equal, as in the situation of Buridan's ass, which had to choose between two equally attractive options. Most generally, we may assume that the probability is equal to $P(a_i \mid c)$, where $\{a_i \mid i = 1, \ldots, n\}$ is a discrete set of possible actions, while the condition $c \in C$ varies continuously over the space C of all states that the world can have. In this model, the probability (and therefore frequency or intensity) of an action varies approximately continuously (quantitative variation), while the action itself is chosen from a discrete range of options (qualitative variation).

Sematectonic vs. marker-based stigmergy

Grassé's original definition of stigmergy concerned stimulation by the performed work itself: in his observation, termites are stimulated by the mud heaps they have already built. E. O. Wilson [], in his "Sociobiology", called this stimulation *sematectonic*. However, in many cases social insects appear to be stimulated by pheromone traces, which are left expressly as a means of communication, not as a contribution to the work itself. In fact, it 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 makes sense knowing that termites are blind). The situation is even clearer with 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 grass 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, pheromones, 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", is usually called *marker-based* stigmergy [Parunak].

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 additional cost and complication 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 [Heylighen]. 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 material that still requires 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 [], 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 possibly vital information, which is useful, both for the receiver, who is warned of a potentially dangerous rival, and for the emitter, who can use it to frighten away newcomers from his 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 along 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" that economists consider essential for reliable transactions.

In this case, we see how something (smell) that was merely a side effect of a primary action (getting rid of waste products) turned into an intentional, communicative signal, even though the 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. A human example is an artist making a sketch. The sketch functions both as a first step towards 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 remind the artist of different approaches to the subject, or to convince a sponsor who may be interested to order the finished work. The first function is sematectonic, the second one marker-based.

Transient vs. persistent traces

After discussing basic aspects of stigmergy that are well recognized in the literature [e.g. Parunak], 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 conceptualized as part of the “productions” or (agentless) actions that contribute to the overall process of self-organization. We have conceptualized the medium as the passive part of the stigmergic organization, undergoing shaping and molding by the active productions, but not contributing to the activity itself.

But even a passive medium is subjected to dissipation, diffusion, or the increase of entropy entailed by the second law of thermodynamics. This means that structures and differences tend to spontaneously decay into formless homogeneity—unless they are actively maintained and reconstructed [Prigogine]. Examples are the evaporation of pheromones and the erosion of termite hills by rain, wind and gravity. This spontaneous decay should not be seen as 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 receive continuing 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 so-called *learning parameter* in neural networks [], which determines the size of new changes in connection strength relative to the accumulated effect of previous ones. 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 towards 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.

Given that what counts is the relative attractiveness of different options for action, the "learning" parameter, which measures the importance of new contributions to the trace, is in practice equivalent to a "forgetting" parameter, which measures the speed of decay of the existing trace. The optimal value of this parameter will depend on the speed with which information becomes obsolete. This will itself depend on the variability in the environmental diversions and the specific 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 towards a deer—which is sematectonic). The reason is that sound and light, because of their wave nature, spread and decay almost immediately. An intermediate decay speed is typical for chemical diversions in a liquid environment, where concentrations of molecules may change within minutes. An example of this kind of stigmergic coordination are the chemical signals broadcasted by bacteria that encounter either an affordance, such as food, or a disturbance, such as a concentration of toxins [Ben-Jacob et al.]. The first type of diffusing signal will create a chemical gradient that incites bacteria of the same colony to swim towards the food source, so that they too can profit from it. In the second case, the gradient will incite them to move away from the danger threatening their congener.

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 or productions 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 fax, 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 long 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 disturbances and coordination errors can be corrected without delay.

Synchronous communication is rarely conceived as stigmergic, since it seems to fit in better with the conventional paradigm of direct interaction, as exemplified by human conversation. Yet, the examples we listed, such as the warning cry, are still indirect, since they are targeted at no one in particular but merely "released" in the medium. The stigmergic nature of synchronous interaction is even clearer when the signal is sematectonic. For example, a bird suddenly 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 [Okubo], where the agents are continually adjusting their trajectory on the basis of real-time perceptions of the movements of other agents, as well as diversions such as obstacles, predators or prey.

Broadcast vs. Narrowcast

Another basic component of the stigmergic taxonomy proposed by Parunak [] is the topology of the medium (or "environment"). Here the same difficulties arise as with the dynamics: the potential topologies are unlimited in number and complication, making classification intrinsically hard. Again, I suggest replacing this multidimensional, qualitative notion by a one-dimensional, quantitative aspect: the range or *scope* of the stigmergic process. The scope represents the size of the "neighborhood" across which a stigmergic signal is perceivable. The two ends of the scope continuum may be called *broadcast* and *narrowcast*. Broadcasted traces can be perceived by all agents involved. Narrowcasted traces are perceivable by only one or a few agents. This will obviously depend on the topology of the medium: a large trace in an uninterrupted, flat plain will be visible from afar; the same trace in a landscape cut through by rocks, valleys and trees will only be visible in a small region. It will also depend on the degree of diffusion of the trace: traces such as sounds or smells that propagate easily will have a wider scope than traces that remain localized, such as shapes and inscriptions.

As yet, there does not seem to be much research that can clarify the differences between broadcast and narrowcast. Implicitly, most studies of stigmergy assume broadcast within a given group of agents, such as an ant colony. But obviously, these groups themselves are limited in scope, and therefore there is always a degree of narrowcast.

The situation becomes more complex—but also more interesting—when different actions or agents have a different scope, so that A's traces e.g. may reach B, C and D, while D's traces reach B and E. In this case, the topology of the stigmergic system becomes equivalent to a *network* where different nodes (A, B, C…) each are connected to (i.e. can deposit traces perceivable by) different other nodes. The implication is that the network paradigm—which is increasingly popular for modeling various complex and self-organizing systems such as neural networks, social networks, citation networks, etc. [Heylighen, Newman]—could be viewed as a special case of the stigmergic paradigm, albeit a rather complicated one. The stigmergic paradigm remains more general than the network paradigm in the sense that the scope of a stigmergic interaction can vary, while the "scope" of a network connection is fixed. For example, a more intense trace (e.g. more concentrated pheromone) will typically spread over a somewhat larger scope, and thus influence more agents.

This stigmergic perspective may actually clarify some problems in traditional network models. For example, we know that in the brain connections are not fixed, since neurons can grow axons to connect with remote other neurons. To guide this growth pattern, some neurotransmitter-like signal molecules must be able to diffuse outside of the existing neurons and synapses, implying a more "broadcast" form of stigmergic communication. Similarly, social networks are everything but well defined and fixed in their scope: people's actions will typically have repercussions well beyond their present friends and acquaintances, potentially bringing them in contact with a much wider circle of people. We will leave these issues for future work, and just note that a stigmergic analysis may extend even to typical network models, such as connectionist theories of learning and thinking in the brain.

Bibliography

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

Bolici, F., Howison, J., Crowston, K., & Syracuse, N. Y. (2009). Coordination without discussion? Socio-technical congruence and Stigmergy in Free and Open Source Software projects. In *Socio-Technical Congruence Workshop held at the International Conference on Software Engineering (ICSE'09), Vancouver, Canada*.

Bonabeau, E. 1999. Special issue on stigmergy. Artificial Life 5, no. 2.

Deneubourg, Jl, G. Theraulaz, and R. Beckers. 1992. Swarm-Made Architectures. Toward a Practice of

Autonomous Systems: Proceedings of the First European Conference on Artificial Life.

Dorigo, M., E. Bonabeau, and G. Theraulaz. 2000. Ant algorithms and stigmergy. FUTURE GENER COMPUT SYST 16, no. 8: 851-871.

DOWNING, HA, and RL JEANNE. 1988. Nest construction by the paper wasp, Polistes: a test of stigmergy theory. Animal behaviour 36, no. 6: 1729-1739.

Grassé, P. P. 1959. La reconstruction du nid et les coordinations interindividuelles chezBellicositermes natalensis etCubitermes sp. la théorie de la stigmergie: Essai d'interprétation du comportement des termites constructeurs. Insectes Sociaux 6, no. 1: 41-80.

GRASSE, PP. 1960. The automatic regulations of collective behavior of social insect and" stigmergy". J Psychol Norm Pathol (Paris) 57: 1-10.

Gregorio, J. 2003. Stigmergy and the World-Wide Web. Online document] 2003 13: 12.

Holland, O., and C. Melhuish. 1999. Stigmergy, self-organization, and sorting in collective robotics. Artificial Life 5, no. 2: 173-202.

Parunak, H. V. D. 2006. A survey of environments and mechanisms for human-human stigmergy. Environments for Multi-Agent Systems II: 163-186.

Ricci, A., A. Omicini, M. Viroli, L. Gardelli, and E. Oliva. 2006. Cognitive stigmergy: A framework based on agents and artifacts. Third International Workshop on Environments for Multiagent Systems (E4MAS 06).

Robles, G., J. J. Merelo, and J. M. Gonzalez-Barahona. Selforganized development in libre software: a model based on the stigmergy concept. Proceedings of the 6th International Workshop on Software Process Simulation and Modeling (ProSim 2005).

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*.

Surowiecki, James. 2005. *The Wisdom of Crowds*. Anchor.

Susi, T., and T. Ziemke. 2001. Social cognition, artefacts, and stigmergy: A comparative analysis of theoretical frameworks for the understanding of artefact-mediated collaborative activity. Cognitive Systems Research 2, no. 4: 273-290.

Tabony, J. 2006. Microtubules viewed as molecular ant colonies. Biology of the Cell 98: 603-617.

Theraulaz, G., and E. Bonabeau. 1999. A Brief History of Stigmergy. Artificial Life 5, no. 2: 97-116.

Theraulaz, G., and J. L. Deneubourg. 1992. On formal constraints in swarm dynamics. Intelligent Control, 1992., Proceedings of the 1992 IEEE International Symposium on: 225-233.

Valckenaers, P. M., H. V. Kollingbaum, O. Brussel, C. Bochmann, and C. B. Zamfirescu. 2001. The Design Of Multi-Agent Coordination And Control Systems Using Stigmergy.

Valckenaers, P., H. Van Brussel, M. Kollingbaum, and O. Bochmann. Multi-agent coordination and control using stigmergy applied to manufacturing control. dimension 5: 6.