

Getting Things Done: The Science behind Stress-Free Productivity

Francis Heylighen and Clément Vidal

ECCO - Evolution, Complexity and Cognition research group
Vrije Universiteit Brussel (Free University of Brussels)
Krijgskundestraat 33, 1160 Brussels, Belgium
Phone +32-2-640 67 37
Fax +32-2-6440744

<http://pcp.vub.ac.be/HEYL.html>, <http://clement.vidal.philosophons.com>
fheyligh@vub.ac.be, clement.vidal@philosophons.com

Abstract: Allen (2001) proposed the “Getting Things Done” (GTD) method for personal productivity enhancement, and reduction of the stress caused by information overload. This paper argues that recent insights in psychology and cognitive science support and extend GTD’s recommendations. We first summarize GTD with the help of a flowchart. We then review the theories of situated, embodied and distributed cognition that purport to explain how the brain processes information and plans actions in the real world. The conclusion is that the brain heavily relies on the environment, to function as an external memory, a trigger for actions, and a source of affordances, disturbances and feedback. We then show how these principles are practically implemented in GTD, with its focus on organizing tasks into “actionable” external memories, and on opportunistic, situation-dependent execution. Finally, we propose an extension of GTD to support collaborative work, inspired by the concept of stigmergy.

Keywords: personal productivity, personal information management, time management, task management, praxeology, situated and embodied cognition, stigmergy, information overload.

Introduction

Our present society is characterized by quickly growing complexity and change: opportunities, constraints, and objectives are in a constant flux. Managing the situation requires gathering and processing an incessant stream of potentially relevant information. As such, most of our day-to-day activities fall under the heading of *knowledge work* (Drucker, 1973). But how can we efficiently organize such heavily information-dependent work? While there is a large and established literature on how to organize traditional physical activities, such as industrial processes, the literature on knowledge management is as yet much less well developed, and mostly concerns the static storage and reuse of existing knowledge rather than the processing of incoming information. The extensive literature on information processing, on the other hand, mostly concerns computer systems rather than human information processing.

Some aspects of human information processing, such as decision-making, project planning and problem solving, have been well investigated. However, the corresponding theories are not really helpful to cope with the information explosion, since they generally assume a given range of possibilities from which the best possible one is to be chosen. However, in a situation where new information may arrive by the minute, both the relevant options and the criteria for deciding between them are constantly changing. This makes formal optimization methods basically useless in day-to-day knowledge work. As Simon (1997) pointed out long ago, rationality is *bounded*: we never have the full information needed to make optimal choices.

But Simon's alternative strategy of "satisficing" is flawed too: a choice that is satisfactory now, may not appear so good anymore when new information comes in. In practice, people follow a strategy of "bettering": choosing what *seems* best from the available options now, but being ready to switch to something better when new information arrives. This opportunistic mode of decision-making is pervasive in today's fast-paced and uncertain world. However, its lack of a clear focus makes it likely that people would not really know what to aim for, what to do, and what not to do. Moreover, the constant bombardment with new information means that previous plans, decisions and relevant data are often forgotten or neglected.

The last two decades have seen an explosion of methods for "time management", "task management", or "personal productivity enhancement" that try to teach knowledge workers efficient routines for dealing with this overload of ever changing demands (e.g. Covey, Merrill & Merrill, 1994). Most of the recommendations concern concrete tools and techniques, such as installing spam filters, using personal organizers, sharing calendars, etc. Insofar that they look at the wider picture, however, they tend to remain within the optimization paradigm: they suggest first to formulate clear objectives or priorities (optimization criteria), and then to order the different tasks according to (a) how much they contribute to the priorities, (b) how much time, effort or other resources you need to invest in them. The recommendation is then to focus on the tasks that contribute maximally to the chosen objectives while requiring minimal resources.

Although this strategy may seem self-evident, it does not take into account the fact that for knowledge work both priorities and resources are in general ill defined and constantly changing. The reason is that information, unlike material resources, is not a conserved quantity: it can appear (be discovered or communicated) or disappear (become outdated) at any moment. For example, an engineer planning the construction of a bridge can be pretty confident that the amount of concrete and steel necessary for the construction will not suddenly change. On the other hand, an author planning to write a book about how to use this great new communication software may suddenly find out that the software has a fatal security flaw, or that another writer has just finished a comprehensive treatment of the same subject. If that author had planned her complete work schedule around the book project, she would have to start her planning from scratch. More generally, applying an optimization strategy to knowledge work may *produce* rather than *reduce* stress, as people worry about what priorities to accord to different alternatives, and then feel guilty or disoriented when they have not managed to follow their own prescriptions

because of unforeseen changes.

The personal productivity consultant David Allen (2001) has proposed a fundamentally different approach. Based on years of experience in teaching knowledge workers how to deal with their backlog of unprocessed issues, the method is known as “Getting Things Done”, or GTD for short. GTD is intended to minimize stress and anxiety while maximizing productivity—in the sense of maximizing the number of useful tasks performed, not in the sense of maximally achieving a given objective. The method has become remarkably popular in a very short time. According to the Amazon web bookshop, in October 2007 Allen’s book ranked number one in the bestseller lists for both the domains “time management” and “business: health and stress”. At the same time, a search via Google found more than a million web pages referring to this methodology. Numerous software applications and adaptations of existing software have been created to help people implement GTD in their daily life (Wikipedia contributors, 2007). In spite of the many testimonials that GTD works in practice, however, as yet no academic papers have investigated this method. A search (October 2007) for articles referring to (Allen, 2001) turned up a mere 14 documents, none of which discusses the method in any detail.

The present paper intends to fill this gap in the scientific literature. While it would be interesting to test GTD empirically, e.g. by comparing the productivity of people using GTD with the one of people using different methods, this is intrinsically difficult. The reason is that because GTD does not embrace explicit priorities or optimization criteria, there is no obvious standard by which to measure expected productivity enhancements. A simpler approach may be subjective evaluation: how satisfied with their work are GTD users compared to users of other methods? However, this will still teach us little about precisely how and why GTD is supposed to work. The present approach has chosen to address this last issue from a theoretical angle, starting from recent insights in cognitive science and cybernetics.

We wish to view GTD and its proposed theoretical foundation as a first step towards a concrete *praxeology*, i.e. a theory of practical action, with specific application to knowledge work. A *praxeology* has been recognized by the philosopher Leo Apostel as one of the fundamentals components of a worldview (Aerts et al., 1994). Such a theory is independent of any specific goals or values: these are chosen by the individuals performing the actions. In GTD, however, the implicit value is to maximize productivity, i.e. to do more (tasks) with less (time, effort, resources).

In our present information society, mental resources in particular tend to be strained. Indeed, in Simon’s (1971, p. 40) memorable phrase, “a wealth of information creates a poverty of attention”. Given that the amount of attention that we can devote to our work is finite, a growing amount of information clamoring for our attention must at a certain moment produce an overload, where a number of (potentially important) items simply can no longer be processed. GTD is intended to facilitate this unavoidable process of selectively ignoring demands while remaining maximally in control of the situation. Although the method is rooted in practical experience, we will try to show that its success can be justified on the basis of solid theoretical foundations. To do that, we will need to review what the most recent theories about cognition and the brain tell us about information processing in the real world—as opposed to an abstract realm of logic and rationality. But first we need to outline the basic principles of GTD.

The GTD method: summary

GTD is a simple and flexible method for managing your day-to-day tasks or activities, so as to maximize personal productivity. The intended result of applying GTD is being able to keep up with a high workflow in a relaxed manner. The main principle is to get everything that is nagging you out of your mind and into a trusted external memory (file system), so that you can stay focused on what you actually have to do now, rather than on various ideas, plans and commitments for later. To achieve this, GTD provides a compilation of tips and tools, organized around a central flowchart, as depicted in Fig. 1. Organized people will certainly already use calendars, to-do lists, note-taking devices, and other tools. What GTD adds, however, is a method for using those tools systematically together. Allen distinguishes five basic stages in our work:

We (1) *collect* things that command our attention; (2) *process* what they mean and what to do with them; and (3) *organize* the results, which we (4) *review* as options for what we choose to (5) *do*. (Allen, 2001, 24)

Collect (1)

The first phase is to collect *everything* that catches your attention as potentially relevant to your activities, whatever its subject, importance or degree of urgency. This includes incoming letters, emails, phone calls, reports, articles from magazines, agenda items, suggestions and requests from other people, and personal ideas and memories. For the collecting process, you need one or more *collection tools*, which can be physical (trays, folders, notebook, etc.), or electronic (email application, outliner, or word processor, on a computer or a PDA). These together define your “in-basket”.

Collection is just the first step: to gain control over the collected materials, you need to empty the in-basket regularly. Emptying means deciding what to do with—not actually doing—the items in the collection. This happens by processing and organizing the items one by one.

Process (2) & Organize (3)

The processing and organizing phases are summarized in the flowchart (Fig.1).

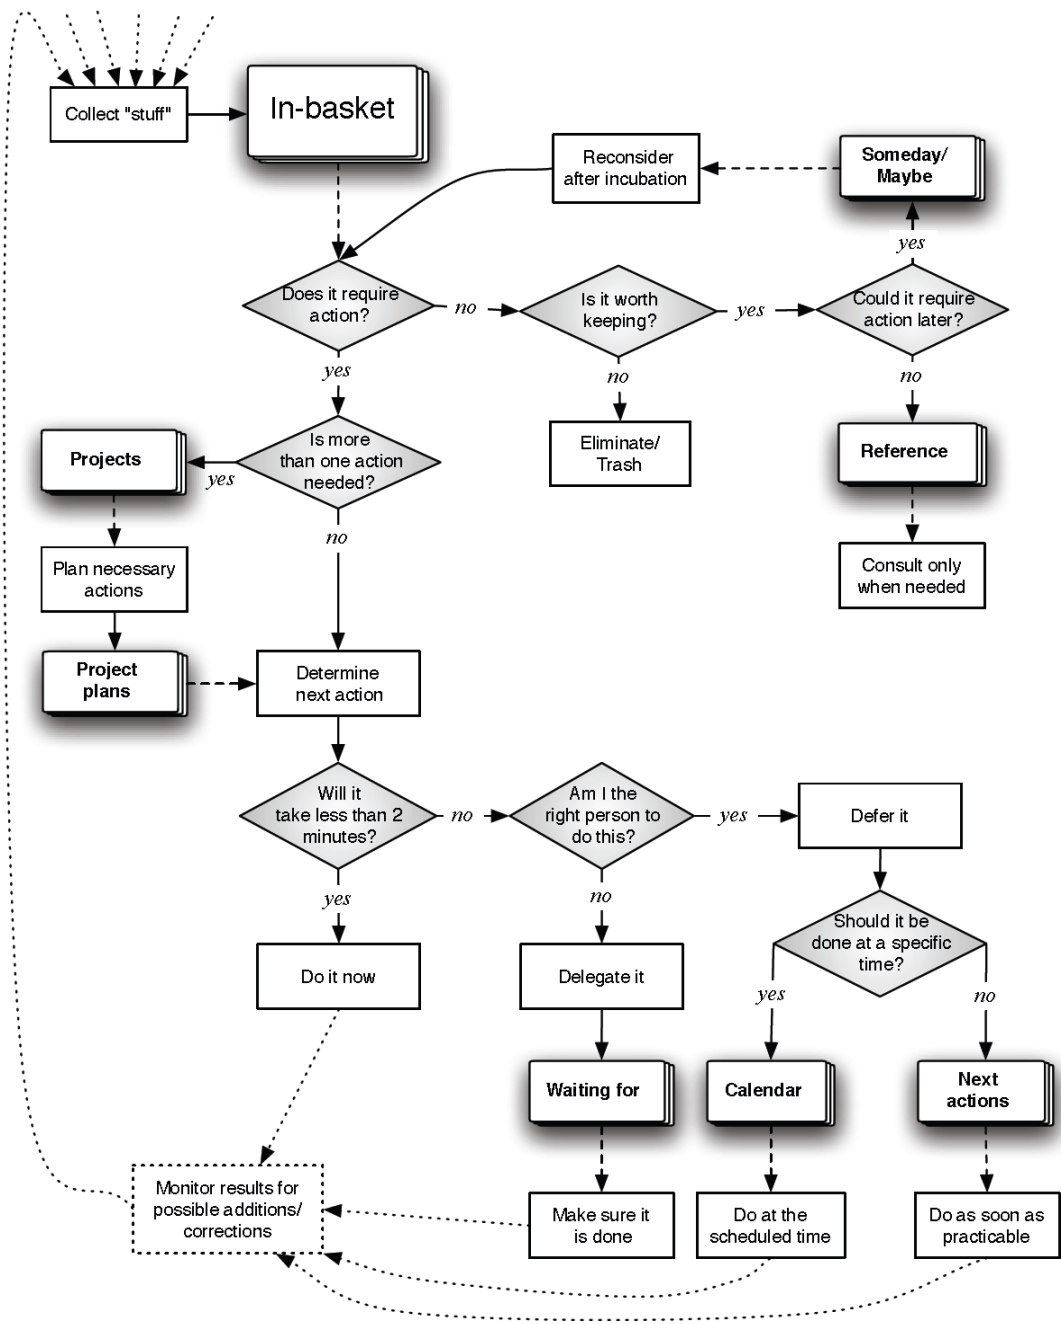


Figure 1: a flowchart depicting the GTD process for organizing and processing incoming “stuff” into action categories (elaborated from (Allen, 2001, p. 32)). Rectangles represent actions, diamonds represent decision points, stacks represent external memories (lists, folders, files, ...) in which reminders are stored. Continuous arrows represent the immediate sequence of processing; broken arrows represent delayed processing, during a review of an external memory; dotted lines represent follow-up processes left implicit in GTD, but whose importance will become clear in the explanation of Fig. 2. The process starts by taking one item out of the In-basket (top-left), and then follows the arrows depending on the answers to the questions. It ends when the item is classified in an external memory or the corresponding action is performed. The most important actions are likely to end up on the bottom-right.

The sequence of decisions to be made starts from the top left of the flowchart and proceeds downwards. The first question to ask is: “What is this stuff?” Note that “stuff” is a catchall word, which can refer to an email, something at the back of your mind, a note, a voice-mail, a scrap from a newspaper, etc., i.e. any item that has been collected. More precisely, the crucial question is: “Is it actionable?”, i.e. does it require you to perform an action?

- if it is *not actionable*, there are three possibilities:
 - *eliminate* the item if you really will not use it (i.e. throw it in the trash bin);
 - *incubate* the item for possible implementation later (i.e. store it in a **Someday/maybe** file that you will review at a later time);
 - *reference* the item if it does not require action but may need to be consulted later (i.e. store it in a **Reference** file, which is organized so that items are easy to classify and retrieve).

- if it is *actionable*, then you have to decide, “What is the next physical action?”
 - if there is more than one action required, store it in your **Projects** list.
 - if the action requires less than two minutes, it is not worth the effort of entering it into the system: better perform it immediately.
 - if you are not the best person to do it, delegate the action to a more qualified person/organization, and keep track of whether you get back the desired result by entering a note in the **Waiting for** file.
 - if the action is to be done on a particular day and time, defer it to this moment, and note it on your **Calendar**.
 - if the only time constraint is that you should do it as soon as you can, put it in your **Next actions** file.

When you review your **Projects** list, for each project you should start developing a **Project plan**. This in general does not mean a formal scheme with milestones, deadlines and objectives, but a formulation of the overall goal or desired outcome, with a focus on the list of next actions required to move towards this goal. Once these actions are defined, they need to be reviewed, which means that they should follow the part of the flowchart that describes the decision tree for actionable items. There is in general no need to plan actions far ahead: once the first “next action” is done, the next “next action” will probably become obvious.

To make this summary more concrete, table I shows an example of a very simple GTD list of reminders arranged in their appropriate categories. Note that items are susceptible to move from one category to another. For example, the item “plane tickets for Brussels” was initially in the **Project Plan** “Travel to Belgium”, reminding you to order the tickets; now you are **Waiting for** them to arrive by post; if they don’t arrive, it will become a **Next Action** to call the company about the tickets; after you have used them, you may store the tickets as a **Reference**, so that later you could potentially use them as proof of expenses made.

<p>Next Actions</p> <ul style="list-style-type: none"> - Buy a present for Ellen - Call Peter about the new contract <p>Project Plan: Travel to Belgium</p> <ul style="list-style-type: none"> - Book hotel - Phone tourist office - Change money <p>Calendar</p> <ul style="list-style-type: none"> - Oct. 29: Ellen's birthday - Nov. 12: departure for Brussels <p>Waiting for</p> <ul style="list-style-type: none"> - The plane tickets for Brussels <p>Someday/Maybe</p> <ul style="list-style-type: none"> - Read that novel set in Belgium <p>Reference</p> <ul style="list-style-type: none"> - Visa pin code: 4576
--

Table I: an example of some reminders listed under their corresponding GTD categories

Review (4)

The reviewing phase is crucial to remind you of what you still have to do. The *daily review* includes reviewing first your **Calendar** (which are the things that you have to do imperatively on this day?), and then your **Next actions** list (which are the things that you should do as soon as practicable?). The *weekly review* is a more in-depth review of all your (potentially) actionable files (**In-basket, Calendar, Next actions, Projects, Project plans, Waiting for, Calendar and Someday/maybe**). It is essential to get an overview of what has to be done in the coming period, and thus get the feeling of being in control. Concretely, it means that you make sure that the different files in your external memory are kept up-to-date. This will include a complete cleaning of your desk, email, and other collection places, and thus some further processing and organizing according to the flowchart.

A regular review is important in order to develop and maintain genuine trust in your system. Most people who are not applying GTD do this kind of review a few times a year, for example at the beginning of a new year. This gives them a great feeling of clarity, control, and purpose. These good intentions, however, quickly dissipate when new, unprocessed things start to accumulate, and previous plans become outdated because of changing circumstances or lack of follow-up. If they would do such a review systematically every week, this feeling of control and goal-directedness could become permanent.

Do (5)

Having all your lists of to-dos up-to-date, what should you do right now? Allen proposes three models for deciding which action to perform. The first is the “four-criteria model for choosing

actions in the moment”, which advises that you consider the following factors in the listed order:

1. **Context:** What can you do here and now? You cannot do the same actions when you are at your desk as when you are walking in the street. The context limits your choices to the tasks you can (practically) perform. If you have a large number of “next actions”, it is recommended to classify them by context (“office”, “home”, “errand”, etc.), so that actions requiring the same context can easily be performed together.

2. **Time available:** How much time do you have now? Fit the duration of the next action you choose to the amount of time available: if the time is limited, do only short actions.

3. **Energy available:** How much energy do you have at this moment? Adapt your choice of action to your level of physical and mental energy: when you are tired do only routine actions, and keep difficult actions for when you feel more energetic.

4. **Priority:** What are your priorities? Given the context, the time and energy available, what action should be done first? The two following models help you to answer that question.

The “threefold model for evaluating daily work” proposes the following possible strategies:

- (1) Do work as it shows up
- (2) Do predefined work
- (3) Define your work

Is the *work that shows up* (1) the most urgent thing you have to do? When you accomplish tasks as they appear (answering a phone call, chatting with a colleague passing by, replying to an email that just arrived, etc.) by default it means that you are deciding that these tasks are the most important ones at this moment. Alternatively, you can decide, if possible, to postpone the work that shows up, in order to focus on your *predefined work* (2). This means that you systematically go through your **Next actions** list. If you do not have any next actions listed, or if you do not feel confident that the listed “next action” is the best thing to do, you have to *define your work* (3). This is similar to the reviewing phase, where you clear your mind by updating your system of to-dos.

Still, to be confident that what you are doing is truly important, you need a deeper insight in your general goals and values. The “six-level model for reviewing your own work” can support such clarification. Allen uses an airplane analogy to define the levels (Allen, 2001, p. 51):

- 50, 000+ feet: Life
- 40, 000 feet: Three- to five-year vision
- 30, 000 feet: One- or two-year goals
- 20, 000 feet: Areas of responsibility
- 10, 000 feet: Current projects
- Runway: Current actions

You can define goals for different terms or time-spans, from tasks to undertake immediately (Runway) to missions that extend over the rest of your life span (Life). The latter require you to answer almost philosophical questions, like “What is my purpose in life?” It is important to engage from time to time in this “vertical thinking” (Allen, 2001, p. 20-21), and write down and review those lists of goals, so as not to be constantly chasing priorities at the runway level.

Cognitive foundations of knowledge work

Knowledge work consists of various forms of human information processing, which includes such activities as data gathering, interpretation, classification, problem solving, and decision-making. These mental processes have been studied since the 1950's by cognitive science (e.g. Luger, 1994; Thagard, 2005).

Limitations of symbolic cognition

Initially, the guiding metaphor for analyzing cognitive processes was the manipulation of symbols according to a complex program or algorithm. This led to the symbolic paradigm for cognition. Its basic assumption is that knowledge is an abstract, internal representation of the external environment. The main task of cognition is to solve problems, i.e. answer queries about that environment and design plans to achieve goals in that environment. This is done by manipulating the elements (symbols) of the representation according to given inference rules in order to find the combination that best solves the problem. The symbolic paradigm thus sees reasoning, planning and (bounded) rationality as the essence of cognition.

The symbolic paradigm was implemented in artificial intelligence (AI), a general approach to simulate cognitive processes by means of computer programs. However, symbolic AI has been much less successful than expected—in particular in terms of reproducing actual human performance. In contrast to the logical reasoning of AI programs, people's reactions are based on intuition, which is rooted in their subjective experience of the situation. This makes them much more flexible in dealing with complex and unforeseen circumstances. In part as a result of these failures, the symbolic view of cognition has come under harsh criticism over the past two decades (e.g. Bickhard & Terveen, 1996; Clancey, 1997; Suchman, 1990). It has now been largely overtaken by a “new” cognitive science, which is inspired more by the concrete functioning of the human mind (biologically, neurologically, psychologically, socially) than by abstract theories of logic and computation.

One fundamental criticism of symbolic theories is that if you try to represent all the relevant aspects of the real world with symbols, your representation becomes much too complex to be systematically explored by a computer, and *a fortiori* by the human brain. Indeed, the brain is limited by the famous “magical number seven” (Miller, 1956): not more than about seven items can be held simultaneously in working memory. A sufficiently detailed description of a real-world situation will typically include hundreds of symbols (words, concepts, features) that can be combined in millions of different ways, making it essentially impossible to manipulate these symbols in order to systematically explore all their potentially relevant combinations.

Instead, the brain relies on its long-term memory, which can store millions of facts, to quickly *recognize patterns* in the incoming information. Recognized patterns function as stimuli that trigger appropriate responses or actions. Unlike a computer program, the neural network structure of the brain is very good at identifying patterns, at associating perceived patterns with the appropriate actions, and at storing patterns and associations in long-term memory. However, it is very poor at simultaneously keeping several such patterns actively in mind while reasoning, because the corresponding patterns of neural activation tend to interfere with each other. Moreover, activation quickly decays because of diffusion and neuronal fatigue.

Finally, while long-term memory is very effective at *recognition*, it is rather poor at *recall*, i.e. reviving memory patterns without perceptual stimulation. This is illustrated by the “tip of the tongue” phenomenon, where a fact, such as a colleague's name, cannot be recalled—even though you know the memory is there. In that sense, human memory is much less reliable than a computer memory for retrieving a fact outside of the concrete context that reminds you of that fact.

Situated and Embodied Cognition

One of the key insights of the new cognitive science is that cognition is necessarily *situated* and *embodied* (Clark, 1997, 1999; Clancey, 1997; Anderson & Michael, 2003). This means that a cognitive system, such as the human mind, is always interacting with its environmental situation via its bodily sensors (eyes, ears, touch...) that perceive, and effectors (hands, vocal chords...) that produce actions. The complexity of the real world is dealt with not by manipulating an abstract internal representation, but by manipulating the world itself, i.e. by performing actions and monitoring their results via perceptions. This interaction is controlled via *sensory-motor feedback*:

- perceptions trigger actions;
- actions produce changes in the environmental situation;
- these changes are again perceived,
- these perceptions trigger new actions to—if necessary—correct or extend the effects of the previous actions.

Different situations will produce different perceptions, and therefore trigger different actions. Both cognition and action therefore are *situated*: they are determined much more by the concrete external situation than by internal reasoning or planning (Suchman, 1990; Susi & Ziemke, 2001; Clancey, 1997). This shifts most of the burden of memory and reasoning from the brain to the environment: instead of having to conceive, predict and remember the potential results of an action, the action is simply executed so that its actual results can be read off from the environmental situation.

Effective actions leave their mark on the environment. Insofar that this mark is made in a stable medium, such as stone, paper or silicon, it functions like an objective registration of what has happened, storing the information for later inspection by the brain. In that way, the brain can “offload” information and store it in an external memory that is more reliable and less energy consuming than its own working memory. In this case, we may say that the mind *extends* into the physical environment (Clark & Chalmers, 1998), or that cognition is *distributed* across the brain and various material supports (Hollan, Hutchins & Kirsh, 2000; Hutchins, 1995). A simple example is taking notes. The markings on the paper change as the results of our actions (writing). On the other hand, they remain safely stored while we do not interact with the paper. When perceived (read), they trigger thoughts and corresponding new actions, such as adding a related item to the list of already registered items.

A useful paradigm to conceptualize the dynamics of such environmentally mediated activity is the concept of *stigmergy* (Parunak, 2006; Susi & Ziemke, 2001; Heylighen, 1999, 2007). An activity is stigmergic if the action by an agent leaves a mark (*stigma* in Greek) in the environment that stimulates an agent (the same or another one) to perform further work (*ergon* in Greek). This subsequent action will leave another mark which in turn will stimulate yet another action. Thus, different actions indirectly trigger each other, via the traces they leave in the environment. For example, upon noticing that someone has used up all the paper, you leave a note to your secretary to buy paper; the subsequent appearance of paper reminds you to print that long report; the printout in turn stimulates you to study its recommendations; etc. Stigmergy was initially conceived by Grassé (1959) to explain the activity of social insects, such as termites and ants. This collaborative activity, such as nest building, is apparently complex, intelligent and goal-directed. Yet, the individual insects are intrinsically very dumb, lacking anything like a working memory or ability to plan. Thanks to the mechanism of stigmergy their work is efficiently coordinated.

The environment not only provides a passive medium that registers the effect of actions: it

actively intervenes in the agent's activity, producing opportunities to perform new actions or disturbances that make the actions' result deviate from what was intended. In situated cognition, opportunities for action created by the presence of specific objects or situations are called *affordances* (Norman, 1999). For example, the presence of a phone affords you the opportunity to make a call. Because our brain has evolved to quickly adapt to its environmental situation, our perception is especially tuned to the recognition of both disturbances, that create problems that need to be addressed, and affordances, that may help us to solve problems and achieve our goals (Gibson, 1986).

Being in control

Another simple paradigm to understand this agent-environment interaction is the cybernetic notion of *feedback control* (Powers, 1973; Heylighen & Joslyn, 2001), which is also known as error-controlled regulation. A goal-directed agent, such as an ant or a human, tries to achieve its goals by eliminating any difference between its present situation (perception) and its desired situation (goal). A goal here should not be understood as a completely specified objective or end-state, but merely as an (explicit or implicit) preference for certain situations over others. For every perceived difference between the present situation and the goal, an action is performed to reduce that deviation, i.e. bring the situation closer to the preferred one. If the result as perceived is not sufficient, a next action is performed to again bring the situation closer to the goal, and so on, until the agent is satisfied. Although some actions may be counterproductive (in that they increase the distance to the goal), the overall process tends to zoom in efficiently on the goal because of negative feedback: every new action tends to correct the errors created or left over by the previous action. External disturbances are dealt with in the same way: whatever caused the deviation or error, the system's reaction is to try to maximally reduce it, until there is no deviation left. In that way, the system remains in control of the situation, by efficiently counteracting any movement away from its desired course of action. In feedback control, there is no need for planning or for complex reasoning. This makes the mechanism very robust, and able to deal with the most complex circumstances (Gershenson & Heylighen, 2004).

This cybernetic notion of control is at the basis of the psychological state of *flow* (Csikszentmihalyi & Nakamura, 2002). Flow is the pleasurable state that people experience when they are absorbed in an activity that demands their full attention, but such that they feel in control, i.e. able to effectively move towards their goal, however far away this goal still may be. The psychologist Csikszentmihalyi (1990) discovered the flow state by finding common patterns in those activities during which people reported the highest level of pleasant feelings, as measured by the method of experience sampling. Flow is characterized by a clear sense of goals, and by continuous feedback indicating in how far the last action brought the situation closer to the goal. To experience flow, challenges should match skills, i.e. the task should be neither too difficult, which would produce stress and anxiety, nor too easy, which would produce boredom. During flow, people tend to forget their worries and even their notion of time, focusing completely on the task at hand. Typical flow producing activities (for those who are good at them) are playing a video game, performing music, painting, playing tennis, or climbing rocks. But flow can also be achieved during everyday work—even during something as prosaic as assembly work on a factory conveyor belt—provided the above conditions are met (Csikszentmihalyi, 1990).

Situated cognition: conclusion

We may conclude that the feelings of stress, anxiety, and information overload (Shenk, 1997) that are often experienced during knowledge work may be avoided by restoring a sense of control.

Given the limitations of the brain, this is best achieved when the intrinsically difficult functions of information processing, memory, and the triggering of actions are as much as possible delegated to the environment (cf. Kirsh, 1996, 2000). This means that we should choose or arrange the external situation in such a way that it can reliably store information, stimulate new actions, and provide feedback about the effectiveness of previous actions. In that way, it will allow a complex train of activity to be efficiently sustained, coordinated, and steered towards its intended goals.

The different components of this mind-environment interaction are summarized in Fig. 2. We can distinguish two nested levels of mind: 1) the traditional idea of mind as inherent in the brain; 2) the “extended mind” (Clark & Chalmers, 1998) which encompasses the brain together with any external memories that are used to support information processing. In the traditional perspective, external memory is part of the environment. In the cybernetic or distributed cognition perspective, however, it is part of the agent, since it is completely controlled by the agent. The part of the environment that is not under control—i.e. which does not perform merely as the agent expects—intervenes in the agent’s activity via what we have called affordances and disturbances. These, together with the feedback received via the environment about previous actions and the reminders stored in the external memory, determine the situation as perceived by the agent, and therefore the agent’s further actions.

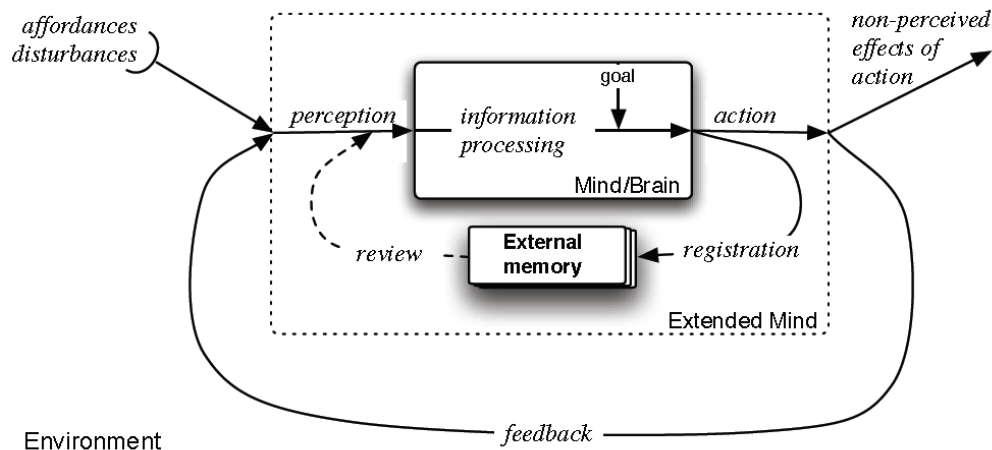


Fig. 2: the major components of mind-environment interaction.

The environmental situation with its affordances and disturbances is perceived by the mind/brain. The information in this perception is processed and compared with the goal or preferred situation. This determines an action to correct any deviation between perception and goal. The action affects the situation, and some aspects of this new situation, influenced by further disturbances and affordances, are again perceived (feedback via the environment). Some actions merely function to register information for later review in an external memory, which is not affected by disturbances. The external memory together with the mind/brain constitutes the “extended mind”, i.e. everything that is under the direct control of the agent.

Cognitive paradigms applied to GTD

Given the situated and embodied perspective on cognition and action, we are ready to provide a scientific motivation for the different recommendations of GTD. We can first note that Fig. 2 can be seen as a simplified version of the GTD flowchart in Fig. 1, with the different external

memories collapsed into a single one. The affordances and disturbances of Fig. 2 are simply the “stuff” collected in the **In-basket** of Fig. 1. The feedback in Fig. 2 makes explicit the fact that the monitoring of performed actions generally suggest further actions to be added to the **In-basket**. Let us now summarize the most important innovations proposed by GTD and interpret them from within this cybernetic/cognitive framework.

Externalizing memory

The first basic message of GTD is that you should as much as possible get everything out of your mind and into a trusted external memory, e.g. by writing it down on paper or in a computer file. In that way, not only won't you forget important or simply interesting tasks, plans, references or ideas, but you will feel much less stressed by the need to remember all that “stuff”. Indeed, the limitations of both working and long-term memory are such that you cannot rely on them to recall all the important facts when they are needed. Trying to do that will merely overburden the brain, as it requires several patterns of neural excitation to be kept activated without getting distracted or undergoing interference with new information coming in. The brain is an intrinsically active medium where patterns are always in flux. As such, it is poor at keeping track of unchanging details. The passive media of paper or hard disk are much better at storing information in an invariant way, so that you can be sure that what comes out is exactly what you put in.

Stigmergic action

The next basic message of GTD is that you should register information as much as possible in an “actionable” form, i.e. in a way that stimulates you to act when you review your external memory. This fits in with the perception→action logic that underlies situated cognition or cybernetic control. Reviewing your external memory means re-entering it into your brain so that its underlying patterns can be recognized by your long-term memory. If the meaning of those patterns is not clear, the brain will need to further process them, by combining them with various other related patterns, in the hope that some new pattern will emerge in which everything fits. This pattern may then suggest a specific action. While such interpretation processes are necessary in complex or novel situations, they demand a lot of additional effort, without any guarantee of success. Therefore, to work efficiently, such processes should as much as possible be avoided, or at least be performed independently of the actions that eventually need to be executed.

GTD recommends performing this reflection *before* the pattern is registered in the external memory. In that case, reviewing the external memory will avoid remaining vagueness and ambiguity, and the procrastination that this typically engenders. Instead, the reviewed item should directly suggest the action to be taken, maintaining the flow of activity without interruption for further reflection. The whole activity can then be performed in a quasi-automatic, “stigmergic” mode, where the note read immediately triggers an appropriate action. Moreover, GTD makes items more actionable by classifying them in a number of discrete categories, each demanding a specific type of action (**Next action** → perform, **Project** → plan, **Someday/maybe** → reconsider...), so that there is no doubt in your mind about what the next step is.

Situated action

Another basic principle of the GTD method is that the decision to perform an action should depend first of all on the *situation*, i.e. the local circumstances that determine in how far the

action is easy to perform here and now. This is considered more important than ordering to-dos by priority, project, or planning. For example, it is recommended that you arrange all phone calls you have to make together in an “at phone” context, and all things you have to discuss with your boss in a “meeting with boss” context. When deciding which of several possible actions to do first, you moreover take into account more subjective situational factors, such as “how much time do I have?”, and “how much energy do I have?”. Only after all these factors have been considered should you think about priorities when deciding about what action to do now.

The principle is that an action is performed most efficiently in the presence of the mental and physical resources, triggers, and affordances that facilitate performing it. For example, sitting in a quiet room next to your phone with its preprogrammed numbers makes calling easy. In principle, you could perform the same calls standing next to a public phone on a noisy street corner, but obviously this will seriously reduce your productivity while increasing your stress level. On the other hand, the street corner may constitute the appropriate situation for buying flowers, as the presence of a flower shop not only affords you the opportunity of purchase (which also exists via the phone or Internet), but also stimulates your senses of sight, touch and smell, so that you can intuitively pick the best option. A similar facilitation occurs on the mental level. For example, it is easier to reflect about how to tackle a particular project after you just had a conversation or read a report about that project, because the relevant aspects are still fresh in your mind. Popular culture knows about this principle through the proverb “you should strike while the iron is hot”.

On the other hand, switching (mental or physical) context costs time and energy, so it is better to minimize it. For example, after you just read a report about project A, somebody calls you to discuss project B. After the phone call, to continue with A you will have to put B out of your mind and try to remember the relevant issues in A. This refocusing effort is a pure waste of mental resources: if you had finished your work regarding A before addressing B, the whole operation would have consumed less time and attention, and most likely have had better results. This is why disruptions are to be avoided. Frequent interruptions, e.g. by incoming email messages or phone calls, significantly reduce a worker’s productivity, presumably because the mind finds it difficult to reacquire its focus after having to shift its attention (Czerwinski, Horvitz & Wilhite, 2004).

The principle of staying within the same context also appears in the “two minute rule” of GTD: if it takes less than two minutes to perform an action, do it immediately rather than file it for later processing. Indeed, considering an item in order to decide which type of action it requires already sets up a mental context around that item. A short action will be performed most easily within that context. If instead the item is classified for later processing, this mental set will have to be recreated. This may take more time than the two minutes it takes to perform the action now and thus eliminate the item from the to-do list.

Adapting more important than planning

Unlike other “project management” or “time management” methods, GTD does not emphasize explicitly defined priorities, milestones, or deadlines, i.e. formalized planning schemes and objectives. These may be necessary for large-scale but well-defined projects, such as building a factory or organizing a customer survey. However, they tend to be counterproductive for everyday tasks and duties, such as answering your mail, arranging a meeting, or simply organizing your thoughts. One reason is that setting up a plan demands quite a lot of mental effort, involving the kind of abstract symbol manipulations for which our brain is not very well suited. For simple, routine activities, starting the job with just a few reminders of what should be done will get you to the desired result more quickly.

Moreover, in our quickly evolving information society we are bombarded with new constraints, challenges and opportunities (what we have called affordances and disturbances), so that priorities and plans constantly need to adapt. What seemed to be a good idea two months ago may well appear outdated today. As a result, you cannot look ahead in any detail for more than a few months. Applying GTD means being ready for any opportunity that arises, but without forgetting earlier commitments. To achieve that, you simply need to register all the interesting opportunities and decide whether you commit to them now or merely file them as **Someday/maybe**. When the situation changes or new information comes in, some of the **Next actions** you had committed to may no longer appear so important, whereas a **Someday/maybe** may now turn out to be worth committing to. In any case, the interesting opportunities will still be available in your external memory, ready to produce actions—unlike a more rigid plan where everything will have to be rescheduled once it turns out that some objectives are no longer worth achieving.

This flexible and pragmatic approach fits in with the cybernetic paradigm, which notes that error-controlled regulation or feedback (reacting after the event) is more basic and dependable than anticipatory regulation or feedforward (acting on the basis of plans or predictions) (Heylighen & Joslyn, 2001; Gershenson & Heylighen, 2004). The reason is that predictions can never be fully reliable: there are always unforeseen events that disturb the most carefully laid out plans. Feedback control, on the other hand, is specifically intended to cope with disturbances. Whatever the nature of the disturbance, once it has been assessed, a counteraction is produced to reduce its effect. If this corrective reaction occurs quickly enough, the disturbance will be dealt with at the early stage when it is still easy to handle, and not have the time to grow into a serious problem.

Planning, of course, is still useful—and necessary in those cases where problems may be foreseen, such as catastrophes, that are too large to be counteracted after the event. But the planning mode advocated by Allen (2001, p. 54) is loose and flexible, emphasizing a clear sense of overall purpose coupled with a spontaneous “brainstorming” approach where different ideas on how to approach the goals are written down in an external memory, and then organized according to their intuitive relationships, rather than an imposed, formal structure. This “natural” planning method fits in much better with the way our brain works, and is more likely to adapt easily to unforeseen circumstances.

Indeed, the situated action approach (Suchman, 1990) reminds us that plans must always remain subordinated to the situation: whenever something unexpected happens, control switches back to the feedback mode, and any plans will have to be adapted or improvised from scratch. This applies in particular to basic research, which is in a sense the epitome of knowledge work. There is an unfortunate tendency in science funding to demand detailed and explicit planning of research projects. Research, by definition, is intended to be creative or innovative. This means that its results cannot be predicted: if you could anticipate a discovery, it would not be a discovery. Moreover, as an almost purely information-based enterprise in a very rapidly changing environment, its objectives constantly have to adapt to new insights and observations. Requiring the achievement of a priori fixed objectives, deadlines, milestones and deliverables is absolutely counterproductive to innovation, as it forces practitioners to restrict their goals to safe and predictable outcomes, while ignoring unexpected opportunities.

Organizing from the bottom-up

Again in contrast to more traditional management methods, GTD starts from the bottom (concrete issues you have to deal with) rather than from the top (high-level goals and values). The rationale is that modern work and life are so complex that if you start from abstract, idealistic goals and try to work your way down to their concrete implementation, you will simply be overwhelmed by the number of possibilities you have to take into account. This is likely to result in a scheme that is

either unworkably ambitious, or rigidly limited. GTD proposes that you first tackle the concrete issues that presently demand your attention, until you feel more or less in control. Only then should you start considering long-term implications of what you are doing, at increasingly more abstract levels. If this long-term extrapolation appears unsatisfactory, it may be time to redefine your higher priorities and change course, safe in the knowledge that at least the short-term problems are under control.

This recommendation can again be motivated from cognitive and cybernetic principles. Long-term planning requires the kind of abstract symbol manipulation that is intrinsically very demanding on the brain. Moreover, given the lack of sensory feedback, the plans that are laid out in this way are likely to remain vague, abstract and unrealistic. Making them more concrete will run into all the contingencies and unforeseen perturbations that make detailed plans intrinsically undependable. On the other hand, any unsolved present issues will remain nagging, creating a sense of anxiety and lack of control, that makes it difficult for the mind to focus on something faraway. When daily activities are running smoothly and on course, it becomes easier to extrapolate this course towards an increasingly distant future, thus getting a sense of where long-term priorities are best laid.

Using feedback to keep on track

Without planning, the danger is that you would just wander from one thing to the next, without clear goal or direction. To counter that, GTD teaches you to couple a sense of overall purpose with a concrete list of **Next actions**, i.e. the very next steps you need to take to move your project(s) forward. Each time you have performed one of these tasks, you can mark it off, thus getting the concrete feedback signal and satisfaction that you are moving forward, and be ready to perform or define a subsequent “next action”. In that way, you are constantly advancing towards your goal at the most efficient speed, without the need for a deadline or otherwise artificially imposed time schedule to make sure that you attain your objectives.

Such feedback-driven, uninterrupted advance towards your goals, at the highest pace you still feel comfortable with, is precisely what Csikszentmihalyi (1990) found to produce the experience of flow. Allen (2001, p. 10) refers to the corresponding mental state as the “mind like water” experienced in martial arts. The idea is that if your GTD task management system is set up well, doing your work becomes stress-free, seemingly effortless, and a source of continual satisfaction. While we personally have not yet reached that Zen-like state while dealing with various administrative hassles, Csikszentmihalyi’s (1990) work makes it very plausible that applying GTD, with its emphasis on clearly defined goals, feedbacks and efforts adapted to the concrete challenges of the situation, would indeed bring one closer to a flow state.

Extending GTD to support collaborative work

GTD is intended as a method to enhance the productivity of individual knowledge work. However, as Allen (2001, p. 255-256) points out, its application in an organizational framework will moreover produce collective benefits. Most obviously, if every individual in the organization becomes more efficient, the group as a whole profits. More specifically, GTD is intended to make individual work more dependable, by reducing the risk that commitments are neglected. If you are less likely to forget or postpone the promises you made to your co-workers, your co-workers will have more trust in your contributions. If all people in an organization become similarly more reliable in performing the tasks they have committed to, the organization as a whole will function much more efficiently, profiting from increased trust, synergy and social capital, while being less

vulnerable to friction, conflict and confusion.

However, in addition to these spontaneous organizational “side-effects” of GTD, we can envisage more direct contributions to organizational efficiency, by extending the underlying cognitive and cybernetic principles to collaborative work. To do that, we can build further upon the paradigm of stigmergy, which was initially proposed to describe the collaborative organization of social insects. The advantage of externalizing information into the environment is not only that it supports individual information processing, but also that it facilitates *sharing* between different individuals. Indeed, an external memory, such as a library or database, can typically be used by many people—unlike the memories in your brain. But the stigmergic/GTD paradigm focuses on more than mere information storage: it demands the externalization of tasks, to-dos or “next actions”, i.e. the registration of concrete stimuli that trigger an activity when encountered in the right context. By sharing these, coordination between different agents becomes much easier.

Let us illustrate this with a classic example of insect stigmergy (Bonabeau et al., 1999; Heylighen, 1999): the creation of a network of pheromone trails by ants. When an ant finds food, it will leave a trail of pheromones (smell molecules) on its way back to the nest. An ant setting out from the nest looking for food will preferentially walk along such a pheromone-marked path. If it too finds food, it will come back along the same route leaving more pheromone. The larger the food source, the more ants will thus come back from it while adding pheromone, and thus the stronger the trace will become. The stronger the trail, the more ants will follow it to find food. Once the food is exhausted, no more pheromone will be added and the trail will quickly evaporate. In this way, ants are efficiently steered towards the presently most promising locations for carrying out their main task: bringing food to the nest. They need neither to individually remember locations, nor to communicate them to other ants: the pheromone network performs the function of both a shared external memory and an indirect communication medium that triggers productive action.

Let us try to imagine how such a mechanism could be implemented in a collective GTD-like system. Most obviously, we can provide shared access to most components of the GTD filing system. Modern computer and network technology makes it easy to create a shared reference system, where all bits of information that are potentially useful for one of the members of an organization are stored for all to be consulted. For example, if you get an announcement of an interesting new publication or the address of a potential customer, you can store it in the organization database, where others can find it by entering relevant keywords. Another already existing tool is a shared calendar, where members of a working group can mark meetings, presentations, or other events that are relevant to more than one individual. Similarly, we have recently seen the appearance of group tools for brainstorming, mindmapping or outlining, which can be used to support the “project planning” stage of GTD.

More complex workflow systems can support the process of delegating tasks to co-workers. However, these tools typically assume a rigid scheme that specifies exactly who does what when. Such explicit plans run counter to the philosophy of adaptability, opportunism and self-organization that characterizes GTD and stigmergy. A more flexible approach is suggested by job ticketing systems, which are used in organizations such as call centers that provide support about technical problems concerning software or hardware. When a customer calls in with a specific question, an expert needs to be located that has the relevant knowledge and that is available as soon as possible. Rather than immediately delegating the task to a specific individual, the system creates a “job ticket” with a short description of the type of problem. These tickets are added to a shared pool of tasks to be performed. When one of the technicians has finished a task, he or she will immediately consult the pool and select the task that best falls within his or her domain of expertise as a “next action”. In that way, tasks are performed in the most efficient way without need for any advanced planning, and thus without a risk of unanticipated problems (such as a job requiring more time than expected so that the designated technician remains unavailable for a newly delegated task).

What will be needed for a collaborative GTD system is an integration and coordination of these different systems so that an organization-wide task monitoring system is created. Incoming items will first have to be processed and classified individually according to the existing GTD scheme, except that now an additional decision has to be made about whether to file the item in the individual or in the organizational memory system. Items for the organizational system will have to be classified as **Reference, Someday/maybe, Calendar, Next action** or **Project**. Note that items that individually may fall in one category (e.g. trash) may collectively fall in another (e.g. **Reference**): what is irrelevant for one person in the organization may be relevant for someone else.

The most important items are the ones that are actionable. Here the additional decision needs to be made *who* will perform the action. In a truly flexible, stigmergic system that decision is ideally made by the individual who commits to the action, not by a boss who delegates the action to a subordinate without knowing precisely whether that subordinate is available, competent or willing to perform it. The philosophy of GTD is that people commit to a certain action on the basis of personal criteria, such as context, time, energy and priority—not because it is imposed on them. Normally, the individuals themselves are the ones best able to judge whether they are ready to perform a task. However, such freedom entails the risks that certain important tasks are never executed, or that certain individuals do not perform their fair share of the workload. To avoid this, items could be entered into the shared work pool with a number of points attached to them, where the points represent an estimate of the importance of the task for the organization. Employees who satisfactorily perform one of the tasks in the work pool receive the corresponding points. At the end of the month, their wages or bonuses may be calculated in function of the total number of points they have earned. This would ensure that everyone is motivated to tackle as many important tasks as possible.

The system would moreover stimulate an efficient and flexible division of labor, since employees would tend to select those available tasks for which they have most skills and the most appropriate situation. Indeed, they would perform these tasks more efficiently than less qualified colleagues, and therefore become available more quickly to collect a subsequent task and its associated points. This ability to work on the task that one feels most competent in is part of the explanation for the surprising success of open-source software development, where the programmers themselves decide what they contribute to (Benkler, 2002; Heylighen, 2007).

If it turns out that certain tasks still have not been performed after an extended period, in spite of the points they offer, this may be a signal for the management that the task is either not that interesting and therefore should better be withdrawn, or—if it is deemed really important—that the task is more difficult than expected and therefore deserves more points. In that way, the pool of tasks with their associated points and the pool of available workers will mutually adapt. Thus, the task pool could start to function like an internal job market whose “invisible hand” efficiently matches supply (of worker’s efforts) and demand (tasks in the pool that require effort). Like the ant trail network, such a job market is an example of *quantitative, marker-based* stigmergy (Parunak, 2006; Heylighen, 2007), i.e. the quantity of markers (points, or pheromones) attached to a task determines the amount of effort that is invested in performing it.

Conclusion

The bombardment with information that knowledge workers presently undergo produces a lot of stress and confusion (Shenk, 1997). Traditional methods for task and time management only provide superficial relief, because they fail to address the central problem: new information typically requires reconsideration of priorities, objectives and resources. When priorities are inconsistent, methods based on optimization or detailed planning become ineffective. In his best-selling book of the same name (Allen, 2001), David Allen has proposed an alternative method: “Getting Things Done”, or GTD. In GTD, the focus has changed from establishing priorities to meticulously keeping track of opportunities and commitments for action. When (or even whether)

these opportunities are followed up depends more on the affordances of the present situation than on any shifting plans for later. Referring to plenty of personal experience, its practitioners claim that this method minimizes stress, while ensuring that work proceeds smoothly towards maximal productivity.

While there are as yet no empirical studies confirming these claims, we have argued that they can be justified on theoretical grounds. For this, we have reviewed a variety of concepts and insights emerging from the new science of situated and embodied cognition, which has largely overtaken the older symbolic paradigm within cognitive science. Proponents of situated cognition assert that the basic functioning mode of the human mind is not reasoning and planning, but interacting via perception and action with the environmental situation. The kind of abstract, internal reasoning envisaged in the symbolic view of cognition is intrinsically hard on the brain, because of its strict “magical number” limitation on working memory and the unreliability of recall from long-term memory. The more natural approach to problem solving is simply trying out actions in the environment and using sensory-motor feedback to correct the situation when errors or disturbances make it deviate from the goal. Further actions are typically triggered by such feedback together with the affordances and disturbances of the environment, i.e. by new information coming in through the senses—not by pre-existing plans. Moreover, the burden on memory can be very much reduced by “offloading” information into a stable external memory, where it is safely stored and ready to trigger action later on.

Although Allen (2001, p. 72) merely hints at the perspective of distributed cognition, GTD appears to implement these same principles. It does this by insisting that all task-related information be stored in a system of external memories, and this so as to be maximally actionable, i.e. ready to stimulate action. To achieve this, GTD proposes a detailed flowchart (Fig. 1) that formalizes the process of collecting and organizing incoming information into a set of action categories. This is followed by reviewing and performing the registered to-dos. The emphasis is on first doing the actions that best match the affordances and constraints of the present situation, rather than the actions with the highest priority. Implicitly, GTD assumes that all tasks in the external memory are worth performing; if in practice not all of them can be done, then it is better to do as many as presently possible. To achieve this you should start with the ones that require the least time and effort *given the constraints and affordances of the situation*. Priorities are subjective and likely to change. Affordances are objectively given, but remain available only as long as the situation lasts. Therefore, maximizing productivity means optimally exploiting the present affordances. This means being ready with a comprehensive list of worthwhile actions to perform whenever the occasion presents itself.

GTD’s claim of making work stress-free can be justified on two grounds. First, GTD minimizes the burden on memory and reasoning by systematically exploiting external memories. As argued by Allen, this will reduce the anxiety caused by not being sure that you remember everything you need to remember. Second, and more fundamentally, the consistent application of GTD should promote all the features that characterize the flow state: a clear sense of purpose; regular feedback as to-dos are “ticked off” one by one; on-going, unrestrained advance towards the goals; and challenges (tasks) adapted to skills (affordances and personal abilities). As extensively documented by Csikszentmihalyi (1990), activities with these features are characterized by a sense of control, focus, and well-being—in sharp contrast to the confusion, anxiety and procrastination that accompany the all-too-common situation of information overload (Shenk, 1997). Of course, we are never fully master of our own destiny, and from time to time challenges will be imposed upon us for which we lack the necessary skills. Therefore, GTD cannot guarantee the absence of work-related stress, but it clearly seems like an important step in the right direction.

The present reinterpretation of GTD on the basis of recent scientific theories does more than justify GTD’s experience-based claims. By grounding the concrete recommendations in a broader theoretical framework, it enables a further generalization, improvement and extension of the methodology. We have discussed one example of such a suggested extension of GTD, from

individual to collaborative work. This extension was inspired by the concept of stigmergy, which explains how shared external memories can support the coordination of collective activity. Further research will be needed to explore this and other implications of our approach.

In the meantime, we hope that both practitioners and theoreticians will be inspired by our paper to apply, test, and further develop GTD as a general method for knowledge work. Even if they do not use GTD proper, we hope that they will take to heart the general principles that we have reviewed, such as the importance of external memory and situation, and the priority of adapting over planning. The philosophy underlying GTD is that true productivity should be measured not by the number of *planned objectives* that are achieved, but by the number of *intrinsically worthwhile results*. Whether these results were foreseen or not is completely irrelevant to their ultimate value. What counts is the total amount of progress made. As we have argued, a flexible and opportunistic approach such as GTD is intrinsically better prepared to maximize productivity in this sense.

References

- Aerts, D. Apostel, L. De Moor, B. Hellemans, S. Maex, E. Van Belle, H. Van der Veken, J. (1994) *World views*. VUB Press. <http://www.vub.ac.be/CLEA/pub/books/worldviews.pdf>
- Allen, D. (2001) *Getting Things Done: The Art of Stress-Free Productivity*, (Penguin).
- Anderson, D. & Michael, L. 2003: Embodied cognition: A field guide, *Artificial Intelligence* 149, p. 91–130.
- Benkler, Y. (2002) Coase's penguin, or, Linux and the nature of the firm. *The Yale Law Journal*, 112, 369–446.
- Bickhard MH, L Terveen (1996) *Foundational issues in artificial intelligence and cognitive science*, Elsevier Science Publishers. <http://www.lehigh.edu/~mhb0/AIFull.pdf>
- Bonabeau, E. Dorigo, M. Theraulaz, G. (1999) *Swarm Intelligence*, Oxford University Press.
- Clancey W.J. (1997), *Situated Cognition*, Cambridge University Press, Cambridge.
- Clark A. (1997) *Being There: Putting the Brain, Body and World Together Again* (MIT Press, Cambridge, MA).
- Clark A. (1999): An embodied cognitive science, *Trends in Cognitive Science* 3:9, p. 45—351.
- Clark A. and Chalmers D. (1998): The Extended Mind, *Analysis* 58, p. 7-19.
- Covey, S., Merrill A.R. & Merrill R. R. (1994) *First Things First* (Simon & Schuster)
- Csikszentmihalyi, M. & Nakamura, J. (2002). The concept of flow. In: Snyder, C. R., & Lopez, S. J. (Eds.). *The handbook of positive psychology*. Oxford University Press, p. 89 – 105
- Csikszentmihalyi, M. (1990). *Flow: The Psychology of Optimal Experience*, Harper Perennial.
- Czerwinski, M. Horvitz, E. Wilhite, S. (2004) A diary study of task switching and interruptions, *Proc. 2004 conference on Human factors in computing systems* (ACM Press, New York), p. 175 – 182.
- Drucker, P. F, (1973). *Management: Tasks, Responsibilities, Practices*. Harper & Row, New York.
- Gershenson C. & F. Heylighen (2004). How can we think the complex? in: Richardson, K. (ed.) *Managing the Complex* Vol. 1 (Institute for the Study of Coherence and Emergence/Information Age Publishing)
- Gibson, J.J. (1986). *The Ecological Approach to Visual Perception*. Hillsdale (NJ): Lawrence Erlbaum.
- Grassé P.-P. (1959). La Reconstruction du nid et les Coordinations Inter-Individuelles chez *Bellicositermes Natalensis* et *Cubitermes* sp. La théorie de la Stigmergie. *Insectes Sociaux*, 6:41-84.
- Heylighen F. & Joslyn C. (2001): Cybernetics and Second Order Cybernetics, in: R.A. Meyers (ed.), *Encyclopedia of Physical Science & Technology* (3rd ed.), Vol. 4 , (Academic Press, New York), p. 155-170.
- Heylighen F. (1999): Collective Intelligence and its Implementation on the Web, *Computational and Mathematical Theory of Organizations* 5(3), p. 253-280.

- Heylighen F. (2007). Why is Open Source Development so Successful? Stigmergic organization and the economics of information, in: B. Lutterbeck, M. Baerwolff & R. A. Gehring (eds.), *Open Source Jahrbuch 2007*, Lehmanns Media, 2007, p. 165-180.
<http://pespmc1.vub.ac.be/Papers/OpenSourceStigmergy.pdf>
- Hollan, J., Hutchins, E. and Kirsh, D. (2000): Distributed cognition, *ACM Transactions on Computer-Human Interaction* 7:2, p. 74—196.
- Hutchins E (1995): *Cognition in the Wild* (MIT Press).
- Kirsh, D. (1996) Adapting the Environment Instead of Oneself. *Adaptive Behavior*, Vol 4, No. 3/4, 415-452.
- Kirsh, D. (2000): A few thoughts on cognitive overload, *Intellectica* 1:30, p. 19—51.
- Luger, G. (1994). *Cognitive science: the science of intelligent systems*. San Diego: Academic Press.
- Miller, G.A. (1956): The magical number seven, plus or minus two, *Psychological Review* 63:2, p. 81-97.
- Norman, D.A. (1999). Affordances, conventions and design, *Interactions* 6 (3) pp. 38 – 43.
- Parunak, H. V. D. (2006) A survey of environments and mechanisms for human-human stigmergy, in: *Environments for Multi-Agent Systems II*, (Lecture Notes in Computer Science, Vol. 3830, Springer Berlin), p. 163-186.
- Powers, W. T. (1973) *Behavior: the Control of Perception*. Aldine, Chicago.
- Shenk D. (1997): *Data Smog: Surviving the Information Glut* (Harper, San Francisco).
- Simon H.A. (1997) *Models of Bounded Rationality*, MIT Press.
- Simon, H.A. (1971): Designing organizations for an information-rich world, in: *Computers, Communications and the Public Interest*, Martin Greenberger, ed., The Johns Hopkins Press.
- Suchman L. A. (1990): *Plans and situated action*, Cambridge University Press
- Susi, T. & Ziemke, T. (2001). Social Cognition, Artefacts, and Stigmergy. *Cognitive Systems Research*, 2(4), 273-290.
- Thagard P. (2005) *Mind, Introduction to Cognitive Science*. (2nd ed.) (Bradford Books)
- Wikipedia contributors. (2007) Getting Things Done. Wikipedia, The Free Encyclopedia. http://en.wikipedia.org/wiki/Getting_Things_Done . (accessed on 6 November, 2007).