LAFORIA Research Report 95/33, December 1995. Older and shorter version under the title "Compromising algorithmicity and plasticity in autonomous agent control architectures : The motivated, social cell", accepted for a poster presentation, European Conference on Artificial Life (ECAL-95), Granada, June 1995.

Compromising algorithmicity and plasticity in autonomous agent control architectures : The autonomous cell [*]

Elpida Tzafestas
LAFORIA-IBP, Université Pierre et Marie Curie,
4, Place Jussieu, 75252 Paris Cedex 05, FRANCE.
brensham@laforia.ibp.fr

Abstract. While algorithmic autonomous agent control architectures demonstrate high efficiency, they suffer from network structure rigidity that shows in the liability to crucial errors. On the other hand, the redundancy inherent in most connectionist architectures allows for continuous self-organization that compensates for limited scale local neuron failures. In this article, we are investigating the relation between algorithmicity and plasticity and attempt a compromise by extending a basic algorithmic cell model. At first, we introduce motivation : the cell becomes a system of multiple independent drives capable of recognizing and consuming different kinds of messages and in competition with one another for the use of the cell resources. Messages may be consumable or catalytic and have identification templates that excite the corresponding cell drives. Those templates are expressed as continuous numeric variables, thus as eigenfrequencies. Unlike usual connectionist models, there are no connections between cells, but message buffers shared by all cells of a level ; this way, cells can be viewed as floating in a common interaction medium. Next to motivation, comes sociality. To ensure the cellular agent's algorithmic operationality, cells should follow the "social standards" of the population. This is achieved via a developmental mechanism of cell eigenfrequency adaptation and the introduction of a parallel immune system, i.e. a population of cells that recognize and eliminate the messages that might be detrimental to the integrity of the cellular agent. The two populations of cells have complementary eigenfrequencies. The model's self-organizational potential is illustrated on the example case of a navigation system. Our simulation results show that the cellular network is recover from various types of failures by "discovering" alternative message flow pathways and that multiple failures slow down the system's responsiveness to external events. Issues such as selectivity and the role of diversity are also discussed.

[*] This work has been the object of three papers (Tzafestas 1995*a*,*b*,*c*), that were merged together, extended and revised to yield this report. Furthermore, the french version of the report (with minor differences) is the seventh chapter of my Ph.D. thesis (Tzafestas 1995*d*).

Résumé. Malgré leur efficacité, les architectures algorithmiques de contrôle comportemental présentent une rigidité structurale qui se manifeste comme susceptibilité aux pannes. À l'opposé, la redondance inhérente dans la plupart des architectures connexionnistes permet une auto-organisation continue et une récupération des pannes de petite échelle. Dans cet article, nous explorons la relation entre algorithmique et plasticité et nous tentons un compromis en élargissant un modèle cellulaire algorithmique de base. En premier lieu, nous introduisons la notion de motivation : la cellule devient un système de plusieurs pulsions indépendantes capables de reconnaître et de consommer différents types de messages. Ces pulsions sont en compétition pour la consommation des ressources cellulaires. Les messages peuvent être consommables ou catalytiques et possèdent des templates d'identification qui excitent les pulsions correspondantes. Les templates sont exprimées par des variables numériques continues, c'est-à-dire par des fréquences propres. À la différence des modèles connexionnistes habituels, il n'y a pas de connexions entre cellules, mais des buffers de messages partagés par les cellules d'un même niveau ; les cellules flottent alors dans un milieu commun d'interactions. Ensuite vient la socialité. Pour assurer l'opérationalité algorithmique de l'agent cellulaire, les cellules doivent suivre les "standards sociaux" de la population. Pour ce faire, nous avons introduit un mécanisme développemental d'adaptation des fréquences propres et un système immunitaire parallèle, c'est-à-dire une deuxième population de cellules chargées de reconnaître et d'éliminer les messages nuisant potentiellement à l'intégrité de l'agent cellulaire. Les deux populations de cellules ont des fréquences propres complémentaires. Le potentiel auto-organisationnel du modèle est illustré dans le cas d'un système de navigation. Nos résultats de simulation montrent que, d'une part, le réseau cellulaire peut "découvrir" des voies alternatives de flot de messages et donc tolérer divers types de pannes imprévues et, d'autre part, que plusieurs pannes ralentissent la réactivité du système aux événements extérieurs. Nous discutons également des questions telles que la sélectivité et le rôle de la diversité.

Contents

1 Introduction

Behavior-based robotics departs from the hypothesis that the source of intelligence is the continuous interaction with an unpredictable world and the need to "survive" in it. On the other hand, the animat approach to cognitive science advocates that primitive low-level capacities of interaction with the world have to be present before higher-level cognitive abilities appear or evolve. Both approaches are bottom-up by definition and rely on the design of artifacts whose main property is that they are able to "survive" in an unpredictable world. Those artifacts are usually called autonomous agents. Their autonomy is viewed as the ability to pursue their goals without external intervention or help, to overcome environmental surprises and to exploit opportunities. Autonomy is therefore relative to a set of predefined goals.

From an architectural viewpoint, autonomous agents (at the robot/animat level) have distributed control systems, either "algorithmic" (as exemplified by Brooks 1986) or "connectionist" (as exemplified by Verschure et al. 1992). We have adopted the algorithmic approach and have been trying to understand the requirements for algorithmic, non-connectionist plasticity, in order to account for recovery from network failures. We have therefore explored the introduction of redundancy in the network. This redundancy might also serve another role : that of cognitive potential to be exploited in further stages of this research.

Our cellular approach purports to develop to a more naturally structured descendant of existing algorithmic behavioral control architectures (subsumption and variants) that suffer from unprincipled design of both individual layers and inter-layer interfaces. The idea is to dispose of a repertoire of primitive components (cells) and to combine them in a proper way to build higher-level structures. Cellularity is defined as the mode of network organization in which the transfer functions are heterogeneous across components but the connections syntax is homogeneous, i.e. the nature of *all* the messages traveling in the connections is the same. Usual connectionist architectures fall into this category as well, since they correspond to the degenerate case of homogeneous transfer functions.

Our motivation behind this line of work has been the intuitive will to adopt a biology-inspired systemic view (Varela 1979) of autonomous agents, i.e. one that presupposes the same organizational principles at all levels : the cell, the organism, the society and so forth. That could be to artificial life what the cell theory is to biology.[1] In (Tzafestas 1995*d*) we explore computational models of *motivation*, *perturbation*, *adaptation* and *sociality/cooperation* (to name just a few of the principles). This article investigates those issues at the autonomous cell level, i.e. one level below that of the autonomous (cellular) robot.

1.1 What is plasticity and why would we want it?

Algorithmic distributed control structures give rise to "preprogrammed" behavior. Such mechanisms are necessary to ensure the operationality of the agents with respect

[1] It doesn't surprise me that by doing artificial life research we are falling on the same principles that we admit as true to biology. Perhaps, this is terrestrial life's way to reproduce itself. I am particularly curious to see what the equivalent molecular level will be and where are we to stop to say that this component is no more an autonomous artificial living agent but just artificial chemistry.

to given tasks. It is however equally important to equip these same agents with self-programming mechanisms, capable of modifying the preexisting structures in order to better respond to the specificities of a given environment. From a designer's viewpoint, this self-programming need reflects the need to minimize human effort during design, as well as the impossibility to determine in advance the details of the agent's activity, either because its environment is a priori unknown or because the agent-environment interactions are of an uncontrollable complexity. Guillot and Meyer (1994) identified three classes of "automatic programming" mechanisms : learning, evolution and development. Learning concerns the self-programming of an otherwise fixed structure according to criteria aiming at more-or-less optimizing the agent's performance ; this kind of self-programming is actually a modification or adaptation of internal parameters of the network components/nodes (in the case of fixed topology connectionnist networks, the patterns of connectivity are reflected in the weights associated with the connections, that may be considered as internal neuron parameters). Evolution concerns the performance optimization through a succession of generations of structures, whereas development corresponds to self-modification of the structure itself thanks to a growth mechanism or equivalent. In what follows, we will be interested solely in the domain of interaction of learning with development within a single cellular agent and we will ignore the option of evolution which calls for a population of cellular agents and for a selection criterion. Those two mechanisms share a functional property : plasticity, that is, *self-organizational possibility of the network*, whether the structure of the network remains fixed (in which case it is called *topology-preserving self-organization*) or not. We say, therefore, that to allow for learning, plasticity is necessary :

> ***Plasticity - Functional viewpoint :*** *Plasticity is the property of a network to self-organize (learn) in response to unforeseen situations, or else its property to adapt to unforeseen input conditions.*

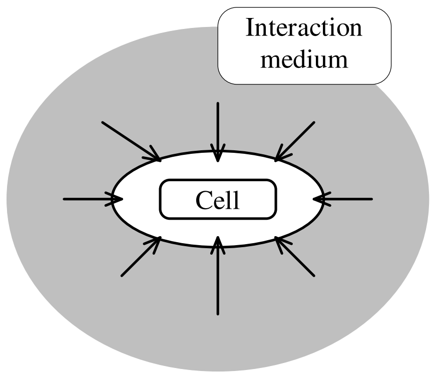

Figure 1 Instructive model (Interaction = Instruction)

From a functional viewpoint, the agent complies with an ***instructive model***, according to which, the environment, or interaction medium, gives learning instructions (fig. 1). There is of course a feedback arrow not visualized in the figure, that of the action of the cell-agent and of the effect that this action has on the interaction medium (otherwise, feedback is itself an instruction of the cell to its world).

An immediate consequence is that the designer's reasoning is that of an observer who attributes a value to what he observes as an instruction or adaptation act according to his *own* intentions and predispositions. A causal theory of the plastic behavior of an observed agent may not have a direct relation with a third party's observations and we are going to try to sketch one below.

1.2 The chicken and the egg : Algorithmicity and plasticity

Plasticity is often seen as the fundamental property used for the generation of high-level categories, such as the "symbols" used in natural languages. For instance, many researchers within connectionism have focused on the emergence of linguistic categories in a neural network having elementary possibilities of self-organization or learning ; however, Clark (1993) insists on the relation between basic structures and emergent functionalities (or "knowledge") and opposes himself to the "tabula rasa" models proposing instead "minimal rationalism" models. Varela et al. (1991) identify the orthodox connectionist paradigm (Anderson & Rosenfeld 1988, Anderson et al. 1991) as the opposite of traditional cognitivism that has been the predominant paradigm in artificial intelligence research until early eighties. They therefore propose an even more radical viewpoint, enaction, where there is an elementary categorization level, that is a "preexistent object taxonomy" in the frontiers of agent-world interaction (ibid., p. 177). Varela et al. (ibid.) cite the work of Rodney Brooks and his group as a robotic demonstration of the importance of a theory based on activity rather than on externally determined functionality. Indeed, a simple solution to a classical spatial learning problem (Mataric 1990, 1991, 1992) reveals the role of the presence of such elementary categories (wall, corridor) ; note the remarkable resemblance with "high-level" qualitative spatial reasoning methods (Kuipers & Byun 1991).

This demonstration brings us anew to the relation between the different behavior-based architectures. Those architectures are generally classified in two major categories : algorithmic and connectionist. Grosso modo, the algorithmic approach (Brooks 1986,1990 ; Arkin 1987,1989) presumes that the problem (robotic task) is given in advance and that the control system may be constructed on this basis, while the connectionist approach (Verschure et al. 1992, Husbands et al. 1993, Yamauchi & Beer 1994) presumes limited "a priori knowledge" that has as a result that the control system (hence the agent!) has to learn by itself the specificities of the problem at hand and -in a way or another- converge to its solution. However, in the case of the algorithmic approach, the agent's operation is not preplanned : on the contrary, the agent's activity depends on continuously received sensor data and environmental feedback, what makes it faurly adaptive and robust to unforeseen courses of action. If all works well, though, the agent wil be seen as following a true, but reactive-adaptive, algorithm. On the other hand, connectionism departs from learning structures and rules specific to the task at hand : it should therefore not be seen as a panacea. Returning to the algorithmic approach, it allows the implementation (programming) of extremely efficient, low-cost distributed control structures (Brooks 1991*a*,*b*), while the designer maintains a sufficient understanding of the system's algorithmic operation. The price to pay is the rigidity/inflexibility of the network structures involved : while the performance of the system degrades rather gracefully when individual components degrade, there is no possibility to recover from such a failure.[2]

[2] Ferrell (1993) implemented on the hexapod robot Hannibal an *ad hoc* algorithm of failure recovery relying on a hierarchical decomposition of the fault detection process and on a clear distinction between hardware (sensors/actuators) and software (program compiled onto the processors). I am

This discussion only makes sense in the case of distributed hardware implementation of individual cells. The usual approach of simulating the distributed control system on an onboard computer does not suffer from this type of failure but necessitates in return greater synchronization efforts. It seems therefore relevant to think on what happens in case of failures if one wants to compare the designed solutions to the natural ones.

Unlike the algorithmic approach, the connectionist paradigm relies on redundant network structures with uniform component functionalities, so that recovery from small scale failures is possible. However, and besides inefficiency, redundancy and uniformity lead to lack of undestanding of the internal functioning of the system and as a consequence make scalability largely impossible.

Leaning closer over this difference between the algorithmic approach and connectionism, we realize that it is in fact just a scale difference : the algorithmic approach seeks minimality for a given problem, while connectionism seeks minimality for a whole class of problems. This difference remains large as long as we cannot yet exactly caracterize what needs to be added in the case of the algorithmic to solve a neighboring problem of the original and what needs to be subtracted in the case of connectionism to solve a particular instance of the class of problems. The algorithmic approach is therefore efficient but structurally rigid, whereas connectionism shows structural plasticity at the expense of efficiency.

Having adopted the algorithmic approach, we tried to understand the nature of plasticity in the opposite sense as the one given so far, that is a form of plasticity necessary for reasons *intrinsic* to the network and that do not necessarily coincide with an observer's point of view. This sort of plasticity oughts to be finally capable of enlarging the application "domain" of the algorithm implemented by the network, that is, of the class of problems that it can solve, and therefore be the true *cause* of functional plasticity.

1.3 "De l'apparence au mécanisme" : Algorithmicity, causality and selectivity

We define plasticity as the property of a network to recover from local failures and to maintain itself as a coherent algorithmic activity unit.

> ***Plasticity - Causal viewpoint :*** *Plasticity is the property of a network to self-organize in order to maintain itself as a coherent unit of innate algorithmic activity, that is the property of a network to maintain its social relations.*

Plasticity is hence viewed as a sort of redundancy and is analyzed here from a "negative" point of view : what kind of redundancy do we need so that the control system succeeds in solving the given problem despite hardware failures and what cascade of failures should there occur to destroy the control system's operationality ? Let us not forget that in healthy living organisms, such as a human or animal body, there are cells dying everyday : in spite of this natural degradation mechanism, organisms manage to survive and even to improve their operation.[3] For an external high-level observer, such as an ethologist, the overall functioning of the organism is not affected by low-level failures. The implementation of a control system on

---

exploring here the idea of failures in a network where such a distinction is absent : everything is of the same nature, everything is cell, and especially the "software" is a network of "material cells", too!

[3] Of course, some of these cells are replaced by new ones, but this is not the case of nervous cells.

distributed hardware should demonstrate the same property. In the future we also plan to analyze this same redundancy from a “positive” viewpoint : if there are no failures, what might this redundancy serve for and what would it take for the control system to “discover” new problems to solve, that is to enlarge the class of problems that it can solve ? Redundancy appears therefore as the cognitive potential that is little by little “exploited” by the agent itself (I am attracted by the idea that learning might just be the reaction of otherwise idle cells that want to participate to the party, that is to the control process[4]).

The task that we are going to study under the prism of plasticity is the navigation of an agent in an environment containing obstacles to avoid and a line to follow eventually in order to arrive to a given goal-position. Our goal in the remaining of the paper will be to *develop* a mechanism of plasticity as defined and also to explore its consequences in terms of network structure. The simulations performed and the results obstained have a qualitative value : they show why and how such a mechanism of hardware plasticity is useful and elucidate its particularities. In an ulterior implementation, we plan to actually *use* this mechanism and the presented cellular model to account for spatial learning of our simulated agent. For the time being, the agent does not learn anything new on its environment, but only how to use best its cellular potential to “implement” its navigation algorithm.

Summary of the plastic model. The extension of the cellular model amounts in equipping each of the cells with multiple “motivations” to process/consume the messages of different types and in removing all the fixed connections in a way so as to make the cells compete for those messages’ consumption. Individual motivations are expressed as continuous variables excited by identification values of the incoming messages, that is as eigenfrequencies. Messages are of two types : the usual messages that are consumed by the cells, and the catalytic messages that trigger the consumption of the other messages by the stimulated cells. As the simulation results will show, genetic variations allow the cellular network to self-organize by “discovering” alternative message flow pathways and multiple failures delay the network’s reactivity to external events, since the cells of intermediate levels are no more specialized to the processing of one type of message, but share their time between many internal motivations.

We also introduce an additional developmental/ontogenetical factor to allow for self-organization even in the case of major failures, when one of the intrinsic eigenfrequencies disappears completely from the network. The most spectacular consequence of the model is, however, that if the messages flowing in the network are false (for instance, they might be mutated or very noisy), the need for an immune system arises. The motivated cell model is therefore further modified to include the functionality necessary for an elementary immune system : two types of cells are defined, the “producer” cells and the defensor cells. The important property of this “two social forces” network configuration is that the two cell populations have complementary eigenfrequencies ! Before proceeding to describe in detail the mechanism as well as the simulation results obtained, I wish to discuss one important property of the new cellular model : ***selectivity***.

---

[4] In a different context, that of a study of self-organization in ant colonies, Drogoul (1993) discussed the need for the presence of idle ants to ensure the robustness of the colony to various perturbations.

Having motivated cells, that is cells that *decide* whether they will process a message or not and what they are going to do with it, implies that cells are no more *directly instructed* by their environment in which they are situated, but rather that they are *selective* toward it ; it is no more *information* transmitted from the world to the agent, but personalized *perturbation* of the agent by a subset of environmental properties. As figure 2 shows, interactions are no longer instructions sent from the world to the cells, but they are instead true actions of selection by the cells. This way, plasticity becomes an essentially *nodal*, rather than a synaptic property. But, let us be honest : selectionist instead of instructionist modeling does by no means change the form and the nature of emergent phenomena. From a design viewpoint, it may simplify the task of programming/construction of control systems that are simultaneously algorithmic and plastic, and it also appears promising for a more causal exploration of these same phenomena (Varela 1979). The importance of selection systems for the future of cognitive science has been stressed by Edelman (1992) who analyzes the *recognition sciences* (evolution, immunology and neurobiology). Recognition sciences study recognition systems, that is the systems having a diversity generator, a heredity mechanism and a mechanism of amplification of selective events. We will come back to the notion of selection in section 5.3.

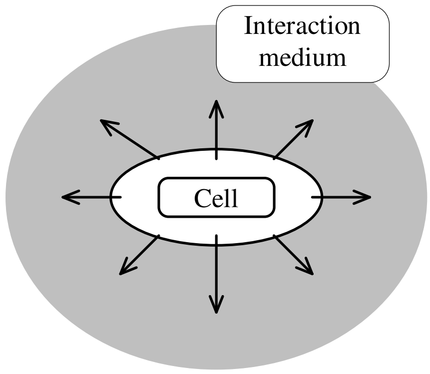

Figure 2 Selective model (Interaction = Selection)

A final question, up to now neglected, concerns the nature and the origin of failures :

> ***Question*** *: If plasticity is in fact the property of a network to self-organize in front of unforeseen failures/lesions, where do these lesions come from ?*

We postulate that those failures are in fact the result of natural *degradation* process of the agent (even of this degradation is not in the same sense as the wear of man-made machines), that is of an aging or senescence process : this hypothesis is analyzed in (Tzafestas 1995*f*) where a senescence mechanism is proposed and described.

We will present first the autonomous cell model and we will discuss its properties before proceeding to a demonstration on the navigation problem. Next, we will show why an immune system is necessary as a defense mechanism against persistent perturbations and we will apply those ideas on the same navigation problem.

2 The autonomous cell : Motivation and sociality

We are summarizing below the relevant aspects of the basic algorithmic cellular model (see the third chapter of (Tzafestas 1995*d*) for a more complete description).

- *The cell is the elementary unit of storage and information processing.* A cell represents and corresponds to a piece of hardware capable of some minimal storage and processing. It has a number of input connection lines, a number of output connection lines, a static or dynamic memory vector and a processing unit, that corresponds to a transfer function plus some side-effect updating of the memory vector (adaptation). Cells are therefore structural primitives differentiated by design or by adaptation in a specific context and represent a uniform substrate for building complex aggregates. The cell's complexity lies between that of a simulated neuron and an augmented finite state machine (AFSM, Brooks 1986), since the transfer function is heterogeneous across different cells, but involves exclusively numeric variables and no symbolic states. Cells also possess an internal execution time which may or may not be controllable by the designer. The transfer function and the size of the memory vector may vary from cell to cell.
- *Cells are categorized into functionally different types.* There are three types of cells, *sensor cells*, *actuator cells* and *processing cells*. Sensor and actuator cells are responsible for primitive sensing or effector operations, respectively. Processing cells categories, named after their transfer functions, include : min, max and find-a-stimulus cells (that output respectively the lowest, highest or just the first non-zero input), time-out cells and priority cells (whose inputs are statically or dynamically ordered and that output the highest priority input). There are also several special processing cell types used in specific substructures in the cell network. Such a special cell type is the decoder cell, which is used in peripheral actuator command systems (such as the navigation system of section 3) and whose role is to combine activity-oriented command input with input from reflex sensors and output the executable command to the corresponding actuator. In general, this command may take one of several values, so that the decoder need only make sure it isn't in conflict with any of the reflex inputs, which represent forbidden command values.

2.1 The motivated cell

The cellular architecture is organized hierarchically and in a generally feedforward manner, by successive layers of cells, corresponding to different algorithmic stages. Let us assume the intermediate cellular level shown in fig. 3. Assume that cell A breaks. For the algorithm to keep functioning, another cell, either B or C, should undertake A's role. For this to be possible, B or C has to fulfill the following three requirements :

- It should be able to switch to A's transfer function,
- It should be able to recognize and process the messages previously processed by A, and
- It should be able to continue doing it previous work despite undertaking also A's role.

The answer to the first two requirements is that the cells should include more than one internal ***drives*** (term borrowed from the animal behavior literature, see for example (Toates 1986)), which are independent "autonomous" components corresponding to those various cell roles. Each of these drives is endowed with an identification template, that is compared against the templates of available messages (template matching corresponds to message recognition). Messages have therefore two parts : an identification template part and an actual data part. The third proposition is true if cells are allowed to include multiple drives executing in parallel.

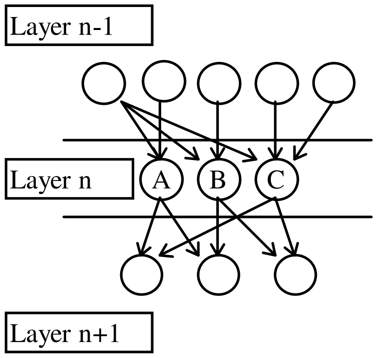

Figure 3 A 3-layer cellular structure

The answer to the first two requirements is that the cells should include more than one internal ***drives*** (term borrowed from the animal behavior literature, see for example (Toates 1986)), which are independent "autonomous" components corresponding to those various cell roles. Each of these drives is endowed with an identification template, that is compared against the templates of available messages (template matching corresponds to message recognition). Messages have therefore two parts : an identification template part and an actual data part. The third proposition is true if cells are allowed to include multiple drives executing in parallel.

To bypass the rigidity of usual symbolic messages, we defined all the data and template parts to be real-valued numbers, and we introduced an extra noise-tolerance in the case of template identification. We may adopt an analogy with physics to say that the identifiers recognized by the various drives are ***eigenfrequencies*** situated in a given scale (in the example that follows, it is the interval from -1 to 1). A cell choosing the interesting messages among all those available is a motivated cell.

To manage the drives' parallel activation and execution, we introduced an additional parameter of ***productivity/speed*** within each drive, which reflects the pace with which the drive scans the input messages. Each drive has a minimum and a maximum speed and a rate of (gradual) change of speed (increase/decrease) according to presence or absence of excitation :

$$speed(t+1)=speed(t)+s\_rate*(final\_speed-speed(t))$$

where *final_speed* is the maximum speed if the drive is stimulated (i.e. if there are messages it can consume) and the minimum speed otherwise. According to this simple adaptation law, the speed of each drive converges to the maximum or minimum speed, if it constantly stimulated or not stimulated respectively, or fluctuates within these two limits in case of uneven and irregular stimulation. We also assume that each cell has a maximum total productivity/speed, so that highly active drives tending to increase their local speed will compete with each other for cell resources (i.e., for fractions of the cell's total productivity). Drives consume messages according to those internally expressed speeds, so that given enough genetic variation in initial speeds, the stronger or more rapid ones prevail by succeeding in consuming the relevant messages first, while the weaker ones gradually fade away.

Trying to make all the messages available to all cells, we may suppose at first glance an *all-to-all* connectivity between successive layers : this would lead to major bottleneck and inconsistency effects, since the same message would be processed independently by multiple cells and the drive competition mechanism would not function properly. To overtake this problem, we cancelled altogether the connections between levels and defined instead *shared message buffers* where unique instances of messages output by cells of previous levels are put into. Furthermore, those messages that would really need to be replicated for different cells (those that in the minimal algorithmic case stimulate more than one cells) have been defined as catalysts, i.e. messages that are not consumable by drives, but that are just "consulted" by them : actually, individual drives make local copies of the data part of catalyst messages. Moreover, and this corresponds to the chemical notion of catalysis, the consumable messages are ignored if the catalysts are not present.

2.2 The social cell : Adaptation and self-organization

Finally, and to account for recovery from broader scale failures, we introduced an additional developmental/ontogenetical factor[5], that slowly pushes unused drives' eigenfrequencies toward unused message frequencies (the adaptation point of reference is the closest frequency) :

$$\begin{gathered} \textit{eigenfrequency}\ (t+1) = \textit{eigenfrequency}\ (t) \\ + \textit{e\_rate} * \textit{closest\_of\_unused\_input\_frequencies}\ (t) \end{gathered}$$

The *e_rate* parameter has to be very low compared with *s_rate*, to avoid "alienation" of individual drives in case of non stimulation, that is to ensure a slow developmental shift compared with the drives' speed stabilization. This parameter will allow cells to self-organize by adapting "socially", that is to stabilize to common frequencies.

2.3 Summary and properties of the model

The above may be summarized as follows (see also figures 4 and 5) :

- ***Messages have two parts*** : a ***"data"*** part - used by the transfer function of the cell processing the message- and an ***"identification"*** part -that is used for recognition by the various drives competing for its consumption. An important property of this model is that ***identifiers come for free***, which implies that the connections' semantics may be dynamic. There are no connections between cells, but messages *float* in the cell interaction medium before being captured and processed. This corresponds to a shared buffer for each connection level. There are two types of messages : the ***consumable messages*** and the ***catalytic messages***.
- ***Each cell is itself a system of independent drives*** competing with each other for the use of the cell resources ; the cell has bounded and fixed maximum productivity[6]. Moreover, each drive possesses a physiological regulatory mechanism, which is uniform across different cell types, that is, independent of the transfer function, and whose role is to manage the competition relations between drives : the "demands" for speed increase are processed in a manner proportional to the absolute current speed of the sender drives and are normalized according to the

---

[5] I use the term "development" in the same sense as Verschure & Pfeifer (1992), Almássy & Verschure (1992), to show that, if the environment is stable, that is, if it does not change qualitatively, the network will gradually stabilize to a permanent operational configuration without growing.

[6] In a future version, I will introduce a decay factor for this maximal productivity (or an aging function, such as the one in (Tzafestas 1995*f*)).

maximal cell speed. Note that the agent's "needs" are related to those messages' consumption : an agent wants to consume immediately the messages that perturb it, so need satisfaction is equivalent to *volume_of_input_messages=0*. However, and since there is no need to recognize the need satisfaction and the mission completion, those needs are not explicitly described in the agent's control system.

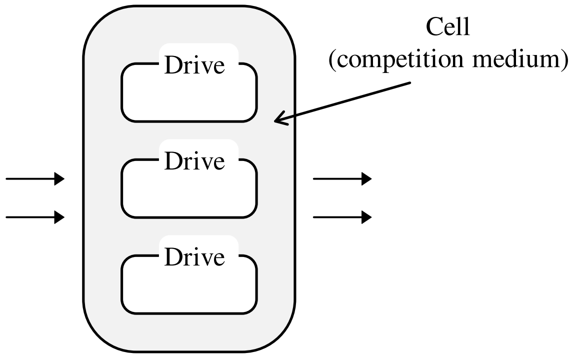

Figure 4 The cell-system

- Every drive is stimulated by consumable or catalytic messages whose identification parts match its eigenfrequencies, that is each drive *recognizes* those messages. Once a drive is stimulated, it executes its transfer function that may be seen as a ***metabolic function*** transforming the input messages into output messages. It is important to stress once again that ***the "motivated" behavior of the cell emerges out of the competition of a set of independent "reactive"components, the drives***.

Drive
speed (between a minimum and a maximum)
transfer function
input eigenfrequencies
catalytic eigenfrequencies
output frequencies
s_rate (for speed modification)
e_rate (for eigenfrequency adaptation)

Figure 5 The drive model

- ***Drives are self-catalyzed by their own activity***, gradually increasing their productivity up to their maximum, as long as stimulation (excitation) persists. Their productivity drops by the same rate, whenever they are not stimulated.[7] This speed regulation mechanism follows an *"eigen-dynamics"* for each drive which is *fast* compared with the task's dynamics (that is, compared with the dynamics of the whole network's operation). In what follows, we will adopt a chemical analogy and we will refer to the drives' local speeds as ***reaction speeds***. The reaction speed self-regulation mechanism is therefore responsible for the *temporary differentiation* of the cells with respect to the input messages.[8]

[7] For the definition of this cellular model, I have been inspired by Monod's (1970) description of the catalytic regulation mechanism of real biological cells, which is accomplished with the aid of allosteric proteins.

[8] Morin (1980) traces the analogy between cellular back-differentiation and despecialisation/ respecialisation in insect societies. The ethological behavioral model of Drogoul (1993) is precisely based on such a dynamic specialisation/despecialisation mechanism by modification of the various

- An additional developmental/ontogenetical factor slowly pushes internal unused eigenfrequencies toward the closest perceived unused eigenfrequencies. The drift takes place at an extremely low rate, to avoid alienation of individual drives. This developmental factor induces a form of ***social behavior*** of the cells, since the adaptation of eigenfrequencies is accomplished with respect to the frequencies present within the interaction medium (within the buffers) and thus with respect to the frequencies coming from other parts in the network.
- ***The cell's physiology*** is the set of mechanisms that link the cell's "program" (i.e., what it is trying to *do*) with its metabolism (that "implements" this program). The drive competition system, the self-catalysis and the drives' adaptation are therefore the physiological functions of the cell.
- As will be shown in the next section, the robustness of the system to failures and to noise is mainly due to the presence of potentially large ***genetic variations*** in the variables that express eigenfrequencies, speeds and adaptation rates *("The real cause of stability in a distributed system is sufficient diversity"*, Hogg & Huberman (1993)).

3 Illustration : Learning to navigate

3.1 The navigation algorithm

We have shown in simulation the model's self-organizational potential for the case of the navigation system of an autonomous robot equipped with 4 obstacle sensors (front, back, right and left) that heads toward a goal point by simultaneously following a line on the floor, if one exists.[9] The algorithmic control system is given in figure 6 and merges the information on the goal with the information on the obstacle and trace (line) sensors. We assume that the heading direction is fed by other parts of the network. In figure 6, the four leftmost "diff" cells compare the input direction with the output of the corresponding obstacle sensor and compute the deviation between the two or they give out a saturation value if an obstacle is perceived. The four rightmost "diff" cells do the same thing for the trace sensors. The cells of the next level sum the differences fed by the upper level ; this sum expresses the total deviation of the corresponding direction from the desired heading and line direction. In this level, it will be necessary to eliminate first the directions blocked by obstacles and then favor those where a line is present. Finally, the decoder cell which is the motor driver chooses between the previous level messages the one proposing the minimum deviation. To this end, the different transfer functions are defined as follows :

obstacle_sensor (dir) : if an obstacle is detected, saturation_value, else dir
trace_sensor (dir) : if a trace is detected, dir, else dir + punishment
*diff : if the difference is less than or equal to 0.8 * saturation_value, difference, else saturation_value*
sum : take the numeric sum
*decoder : eliminate all inputs greater than 0.8 * saturation_value, then choose the smallest*

"motivations" of insect-agents. *"The ability to despecialize, wherever it appears, is an individual regenerative capacity beneficial to the community"* (Morin, ibid., p. 307, translated by us).

[9] This problem has been studied from the robotics viewpoint in collaboration with Philippe Darche, whose robot Pedro (Darche 1994) has served as a thought basis.

The obstacle sensor saturation value and the punishment value are 1000 and 5 respectively.

A first version of this navigation algorithm used a single trace sensor that was giving the best line direction (in order of preference, front, right/left or back) which was subsequently compared with the obstacle sensors output. This version has not been operational, because the agent, on arrival at the end of the line was doing a 180° turn and was going back on its steps (Resnick & Martin (1990) obtained the same result with a symmetrical line-following program on a Lego robot). To solve this problem, the trace sensor should be blind to the presence of a trace in the back, in which case the agent's trajectory would be in return suboptimal for a goal-point in the back. This problem revealed the need to compare the input (goal) direction with *each one* of the directions where a line could exist.

The messages travelling in the connections are real-valued numbers : directions take one out of four values 0.5, 0.25, 0.75, 1 corresponding to front, left, right and back, respectively. The saturation value is defined as 1000, so that the messages arriving to the decoder and whose values surpass 1 may be eliminated.

The corresponding self-organizing network is shown in figure 7. The sets of connections between the first three levels have been replaced by shared buffers, while the connections between the sum cells and the decoder have remained intact (because we did not introduce redundant decoder cells).

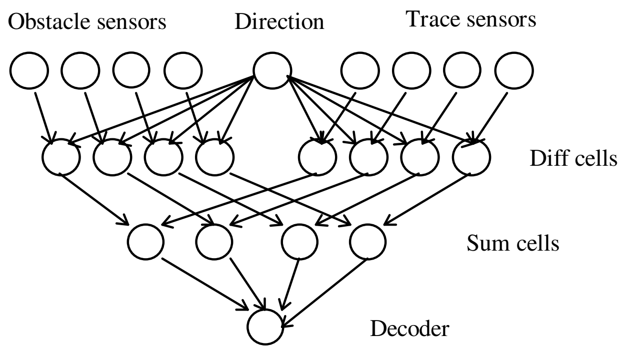

Figure 6 The algorithmic navigation system

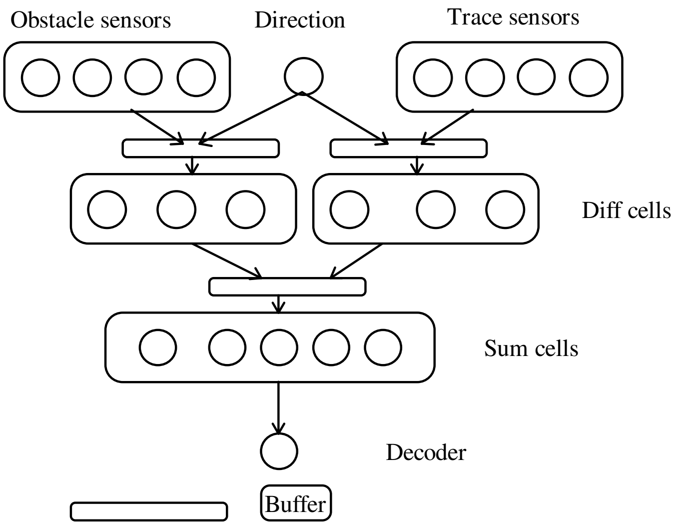

Figure 7 The self-organizing navigation system

We have adopted as eigenfrequencies the same values used for direction discrimination : 0.5, 0.25, 0.75 and 1. This scheme presents an advantage in our case : the heading sent to the motor by the decoder is the eigenfrequency of the chosen message, i.e. of the message corresponding to the minimum deviation (recall, however, that the frequency and the data part of a message are generally *not* correlated). The "diff" and the sum cells' drives need two input eigenfrequencies, so as to merge two incoming messages by taking the difference or the sum, respectively. To prevent obsolete messages from arriving to and being processed by the decoder, all the messages in the shared or private buffers are flushed every time the agent moves (this is equivalent to an additional *"reset"* input for the cell).

The goal direction message needs be shared among all diff cells, so it is defined as a catalyst. In the present navigation case, this means that the output of the obstacle and trace sensors is neglected if there is no goal direction given. The basic heading-by-line-following algorithm is designed for the case of static obstacles, that is, in an unknown but not hostile world. By construction, the algorithm is such that the directions (eigenfrequencies) about which no messages exist are ignored ; if no messages about the left side make it to the decoder quickly enough, then no leftward movements will be undertaken, even if it is viable.[10] In such a case, the agent will take bypasses or it will zigzag (for instance, an agent with the goal to go forward but irresponsive to the front direction will make a zigzag to do so).

A final point concerns the representation space and its metrics. For a transfer function transforming signal A to signal B, A and B are points in a space and the transfer function represents an operation in this space. We are talking therefore about the action representation space, i.e. about the space in which the designer defines and designs the primitive cellular "actions" (or operations). This issue is discussed in the third chapter of (Tzafestas 1995*d*).

The problem with this navigation algorithm is that the space of the directions has to be a ring, so that the subtraction operation give, for instance, *front-right=front-left*, that is not possible with the real numbers whose space is linear and unbounded (open). Moreover, this cyclic property has to be conserved during origin transposition, for instance 5+0.75 has to be the same with 5+0.25. Of course, I solved this problem by replacing the 0.75 differences by 0.25, since only the values of 0, 0.25 and 0.5 are allowed and 0.75 corresponds to 0.25 if ones traverses the ring in the inverse direction. However, this is more a conceptual than a simple implementation problem : ***the form of the representation space should be coupled with that of the agent-world interaction*** (if not the same). If the problem necessitates a ring configuration for the sensors, the primitive operations should be the operations of a ring. It looks to me that organic chemistry is a representation space extremely rich in properties such as ***symmetry***, ***circularity*** and ***recursivity***. In any case, the metrics and the topology of the representation spaces deserve the attention of artificial life researchers. What is the proper representation (mathematical) space best suited for a given problem and what

[10] By default, the decoder has no criterion of knowing whether all the available messages have already arrived (it might happen that there is no message corresponding to a particular direction). It is then the only cell where activation should be possible with a variable number of inputs (1 to 4). To this end, we have introduced an additional variable that expresses the maximal temporal interval allowed for input messages arrival : if all messages (4) have arrived or the maximal duration has elapsed since the arrival of the first message, the reaction takes place with the present inputs. In the simulations, this variable is assigned a random initial value between 0.3 and 0.5.

are its minimal properties ? Up to now there has only been isolated and partial efforts : Thomas Ray (1991) uses a mechanism for template addressing inspired by molecular biology and stresses the non-euclidean nature of the computer geometry (Ray 1994), while Kephart (1994) insists on the importance of topological structures for population dynamics.

3.2 Sane networks

We simulated the operation of the navigation system for 3 to 5 cells in each of the two “diff” cell sets and in the sum cell set. Each cell was endowed with 1 to 3 drives with eigenfrequencies chosen among the four basic ones. Minimum and maximum speeds have been given random values in the ranges (0, 0.2) and (0.8, 1) respectively, whereas speed and eigenfrequency adaptation rates were given random values in the ranges (0, 1) and (0, 0.05) respectively. A typical trajectory for this agent is shown in figure 8. In all the experiments and keeping a constant sensor sampling rate, the network was found to stabilize quickly (in 10-20 cycles) to a configuration where 2 or 3 of the cells in every set were active, i.e. had drives working at speeds close to maximum, with the rest of the drives and cells being dormant, i.e. at minimum speed. The speed self-catalysis of the dominant cells leads to a self-catalysis of the total response speed, as figure 9 shows.[11] The evolution of the speed of a “diff” cell is given in figure 10 : one can see that, as long as the network is stimulated, that is, as long as a goal heading is fed to the network, the stimulated cells are positively catalyzed and try to stabilize to their maximum speed, but as soon as the stimulation stops -here this happens because the goal is met and the heading cell outputs no message- the cells self-catalyze negatively and their reaction speeds fall back to minimum. Finally, figure 11 gives the reaction speeds of the two drives belonging to the same diff cell and shows the competition between drives that leads to a *partial* stabilization, that is to a speed configuration in which none of the two drives attains maximum speed, but the two drives find some sort of “compromise”.

[11] This total speed is computed as the sum of the reaction speeds of different levels, where the reaction speed of a given level is actually the lowest speed of all the reactions that take place (recall that multi-party reactions represent convergence or fusion points in the network).

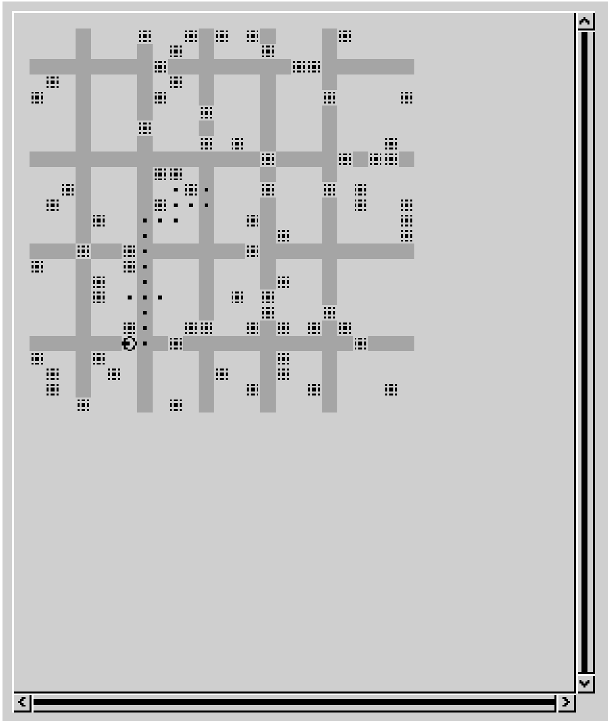

Figure 8 The agent in its navigation universe. Line following with simultaneous obstacle avoidance. The dark parts of the "surface" stand for the traces and the small grids stand for the obstacles. The dot line is the agent's trajectory. The deviations of the agent from the "optimal" trajectory are due to the presence of a degree of noise (here 10%) in the generation of the line direction, what has as a consequence that the chosen direction is not always locally optimal ; this noise allows to the agent to avoid getting trapped in dead-ends.

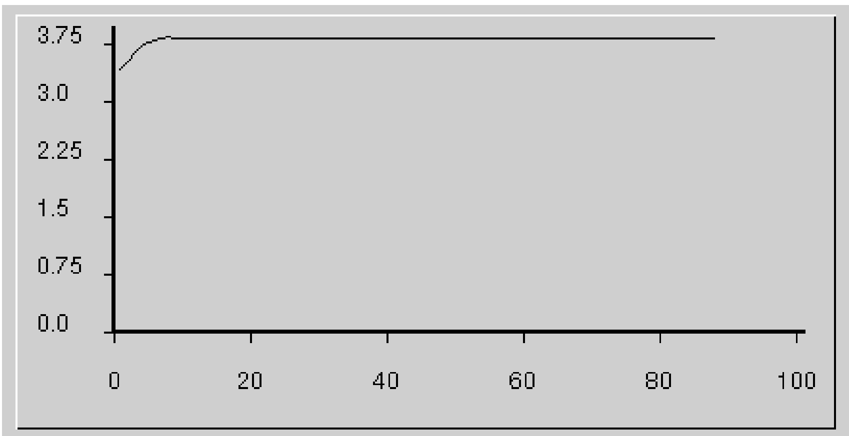

Figure 9 Total response speed stabilization.

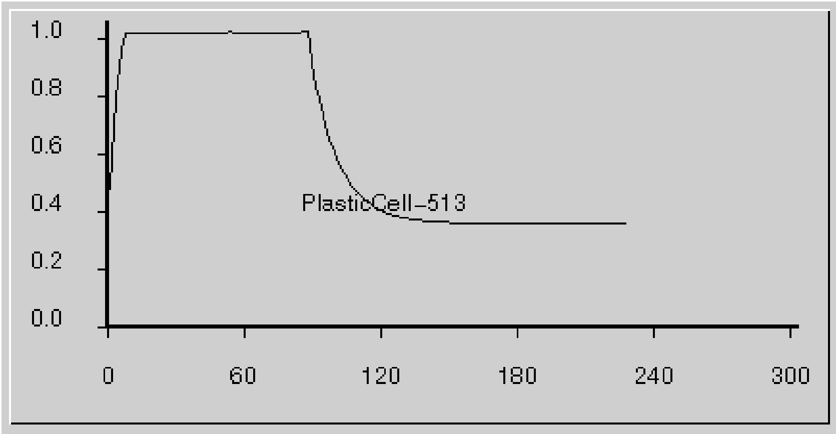

Figure 10 Evolution of the speed of a "diff" cell. In the first part of the curve the cell is stimulated, so its speed rises to its maximum ; in the second part the cell is no longer stimulated and its speed falls to its minimum.

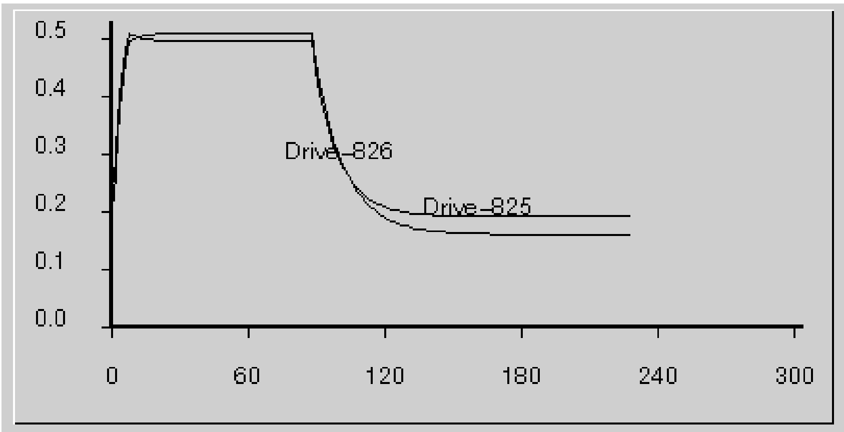

Figure 11 Drive speeds competition for the cell of figure 10. The speeds of the two drives are not their maximal ones, they are away from 1.

3.3 Mild lesions

Next, the following experiment was carried out : once the network had stabilized to a drive configuration (at t=100), we removed one of the active diff cells (i.e., one of the diff cells containing active drives). Once more, it took the agent only a few cycles to re-stabilize to a new configuration ; actually, *the self-organization process is by no mean distinct from the normal network activity*, so the network doesn't stop in order to learn the new configuration before going on. Instead, the navigation system continues operating as before, even if the direction affected by the lesion does not arrive to the decoder on time. Variations in minimum and maximum speeds of drives, as well as in their initial configuration, has as a result that the network does not always lose the direction affected by such a lesion, but often enough it remains sensitive to the direction at the price of sub-optimal response speed (figure 12a). Another relevant remark is that the cell physiology is such that, if suddenly an idle drive is stimulated, the active drives will find a speed compromise that will allow the re-stabilization of the global speed to a more operational value. So, for instance, in figure 12a, once a drive gets stimulated, the second active drive of the same cell yields a fraction of its own speed to compromise with the first.

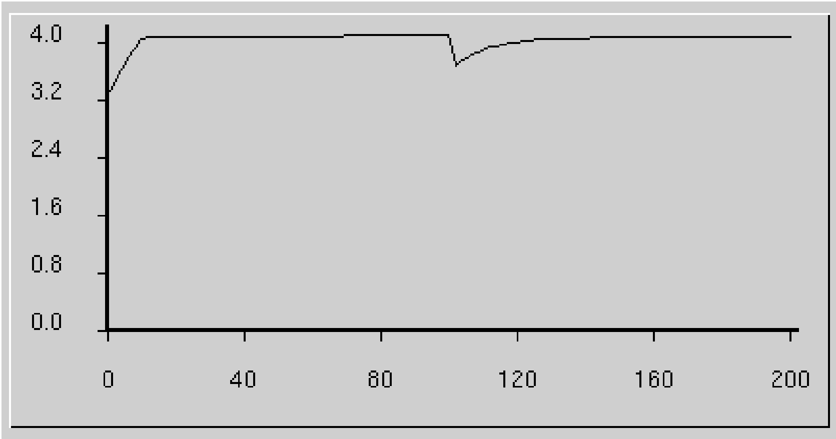

Figure 12a Re-stabilization of total speed. At t=100 removal of a diff cell that used to process the right direction (0.25). The total speed decreases after the lesion and re-stabilizes later (t=200).

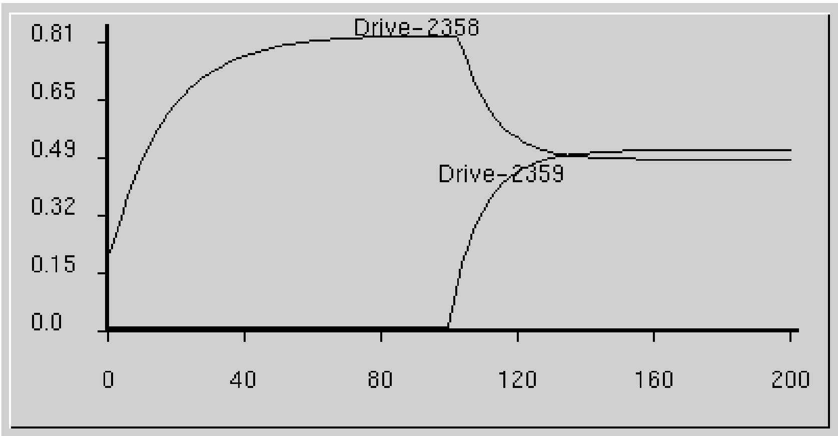

Figure 12b For the experiment of figure 12a, speed re-stabilization of two drives after the lesion (t=100) and the recovery.

3.4 A serious wound

Next, we removed all the "diff" level drives that were excited by a given frequency (actually, all the drives of one of the two diff cell sets). As before, the navigation system became irresponsive to the removed eigenfrequency, but this time it took around 50 cycles before a new message pathway was discovered (figure 13), that is, before one of the unused drives (actually, the most unstable/changeable one!) stabilized to that frequency (figure 14). This is of course an expected consequence of *e_rate* being low, what necessitates a persistent stimulation of the network to ensure reconvergence to a stable configuration. Note also that the total speed of the agent decreased after the lesion, because the remaining cells had to redistribute their productivity among different drives.[12].

[12] I tried to elaborate a stabilization model that would remind what happens in the living organisms, such as animals, where, after recovery from nervous system lesions, the performance of the system is lower as far as the response speed is concerned (my lesion study resembles the one carried out by Beer (1990)).

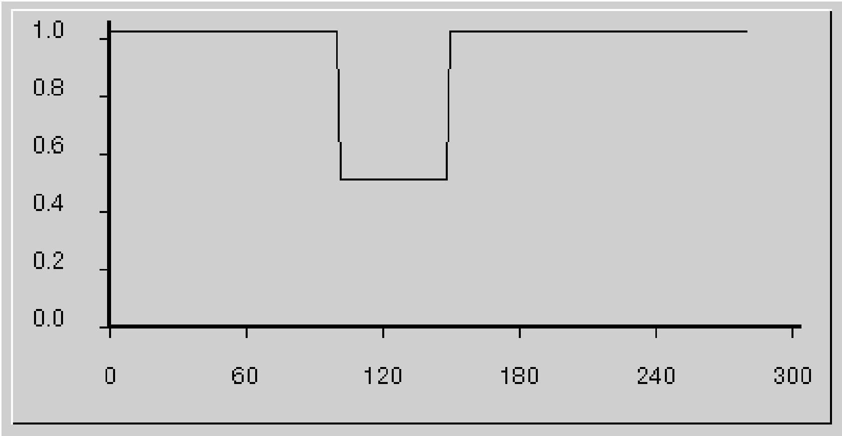

Figure 13 The network's response to a direction before and after lesion (1 if the network responds to the direction, 0.5 otherwise). At t=100 the only cell that responded to the front direction (frequency=0.5) was killed. It took some 50 cycles before the network became again sensitive to this direction.

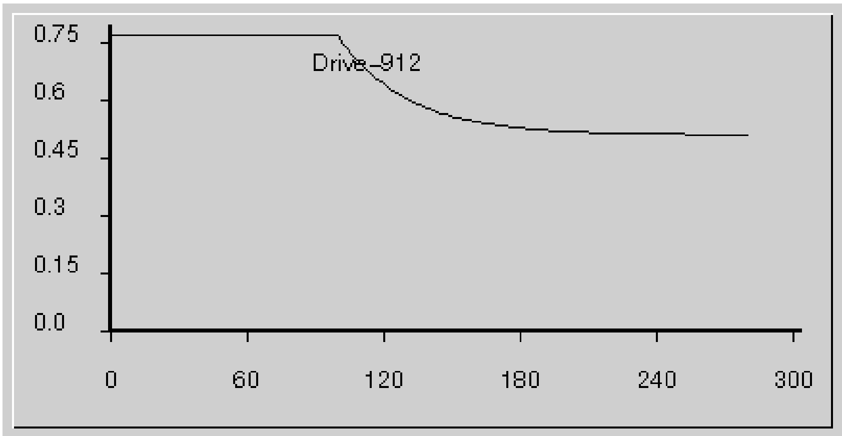

Figure 14 In the experiment of figure 13, the response direction (or response identifier) of the drive that undertook the role of "front" (with continuous developmental adaptation, cf. section 3.5)

3.5 The need for continuous developmental adaptation

"Developmental" adaptation, that is adaptation of drive eigenfrequencies, has to be continuous, that is the drive has to adapt its eigenfrequencies when it is stimulated as well as when it is not. The difference lies in the adaptation point of reference : when the drive is stimulated, it should adapt its eigenfrequencies according to the frequencies that have stimulated it, otherwise it should adapt its eigenfrequencies according to the closest input frequencies. Drives can then stabilize to persistent frequencies in a way that allows a co-stabilization of different level drives, that is in a way that allows *the construction of a social standard or a common language in the network*. To demonstrate the role of continuous developmental adaptation, we carried out the following experiment : we let the network work for 300 cycles without continuous adaptation (that is, with developmental adaptation only in absence of stimulation) and at t=80 we removed one of the stimulated diff drives. At about t=100 the network stabilized to a new configuration.

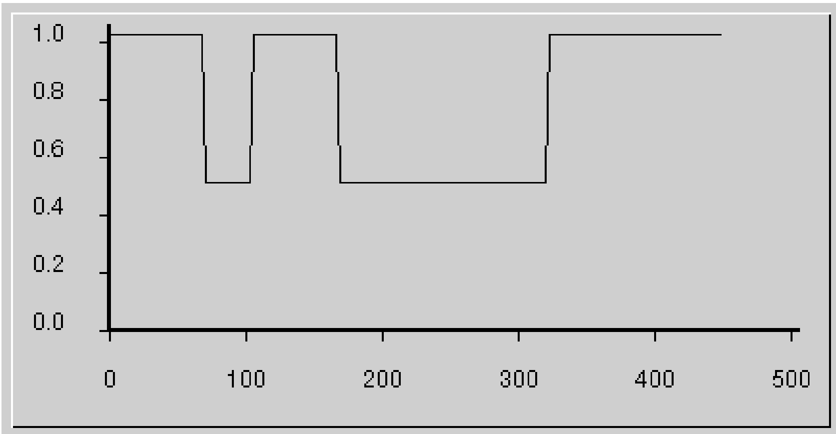

Figure 15 Need for continuous adaptation : response to the front direction (0.5). First part (up to t=300) without continuous adaptation. In that case, with a single level lesion (diff), the system re-stabilizes, but with a second lesion in the sum level it is blocked (see text). In the second part, we introduced continuous adaptation. About 30 cycles later, the network became sensitive again to the front direction (frequency of diff cell = 0.499, frequency of sum cell = 0.505).

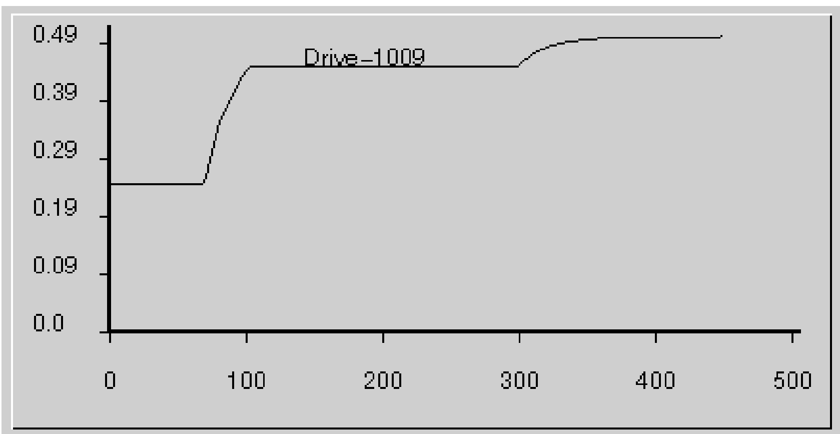

Figure 16 The evolution of the eigenfrequency of the "diff" drive that re-stabilized in the experiment of figure 15.

Then at t=160 we removed a second stimulated drive : we chose the sum level drive that corresponded to the frequency affected by the first lesion (we had taken care not to have other drives of the same eigenfrequency, so as to force developmental adaptation). We let the network re-stabilize and we realized that it was no more sensitive to the direction corresponding to the frequency affected by the two lesions. Why so ? The drives that undertook the roles of those manually removed stabilized separately to a frequency close to the input one (the diff drive stabilized to 0.45 and the sum drive to 0.55[13]) and did not change since. This way, in the sum level, the "fusion" of the two frequencies could no more take place, the difference between the two being too large compared to the identifiers matching threshold, which is around 0.05 (recall that the input and output frequencies of the diff drives are identical, hence the diff drive produces a message with a 0.45 frequency, while the sum drive is expecting messages with a frequency around 0.55). The solution has been to allow drives to adapt *continuously* to the input frequencies : at this point (t=300), we modified the adaptive behavior of the drives and defined developmental adaptation as continuous ; the result was that, this time, the drives whose frequencies were previously dissociated, stabilized anew to frequencies much closer to the input ones (the first stabilized to 0.495 and the second to 0.505) in a way that they could be merged (fusioned) at the sum level. The results of this experiment are given in figures 15 et 16.

3.6 Discussion of the model and the results

The introduced cell model is an extension of the basic model of section 2 and differs from traditional neuron models in two fundamental ways :

- There are no connections betwen cells, but cells "float" in a common medium and share message buffers with other cells. This allows and simultaneously necessitates a competition between cells for the consumption of available input messages,
- Cells respond to messages at their own pace and according to their local preferences/motivations. They are not simply " input machines" and therefore manipulable machines, but they possess a degree of individual autonomy that shows as selectivity toward the input messages. ***Selectivity is coupled with a reaction speed regulation mechanism, that is with a mechanism that modifies the dynamics of interaction with the world.***

Just like connectionist models, though, cells adapt continuously and in parallel with activity, that is the learning and action phases are not separate (not only does the agent adapt without a supervisor, but it has no idea that it adapts). The memory of the system is medium term, and so it does not present the usual forgetting and short memory problems. If it needs to forget, it will ! This adaptation process is based on purely endogenous parameters and does not depend on a benevolent teacher that would have to prepare the learning data, choose the language and the abstraction level of representations (van de Welde (1991) insisted on the need to carefully rethink those questions). According to Kaelbling's (1991) criteria, such a cellular network is ensured to *converge* to a *correct* configuration using an adaptation mode of *constant*

---

[13] The diff drive adapted with respect to the frequency of the message sent by the corresponding sensor, whereas the sum drive adapted with respect to the frequency of the message sent by the second diff level that remained intact throughout the experiment.

spatio-temporal complexity. Let us not forget, however, that the agent does not learn anything new on its environnement, but only how to use best its innate potential ; a scaling study to topological or other learning problems remains to be done.

This model gives room to the introduction of *genetic variations* and to the emergence of division of labor between cells that arises. Furthermore, such a network of motivated cells remains largely understandable and in the case of major failures an engineer may intervene directly to the drive level and upgrade the system manually, since the algorithmic character of activity/message processing is conserved. In this sense, *variation allows self-organization that ensures qualitative stability, or algorithmic operationality*. Besides, the continuous nature of data and frequencies will allow a dynamical analysis of these networks. The adaptation rates as well as the maximum/minimum speeds may also be defined as variable throughout the agent's life ; I can imagine decreasing adaptation rates that will make the agent decreasingly robust to failures and decreasing speed limits that will make the agent decreasingly productive (see also (Tzafestas 1995*f*)).

The temporal behavior of the network is such that it is ensured to converge to a stable configuration if continuously stimulated and in any case it will "follow" any "capricious" stimulation regimes, according to its internal motivations. The self-catalytic cellular behavior makes the agent appear lazy (work at slow speed) in the absence of important perturbations, but increase productivity drastically when stimulated. This self-regulatory behavior of the cell demonstrates the significance of time as a design parameter (issue raised and analyzed in the dynamical neural networks literature, for instance Husbands et al. (1993), Yamauchi & Beer (1994)). As long as the network operates near random initial conditions or for a while after the occurrence of lesions, the behavior of the agent appears as non-coordinated and suboptimal, for instance the agent might wander a little where one would expect it to go straight ahead.

4 Toward immune control systems

We have seen that the developmental adaptation mechanism allows the network to find a stimulation pathway even in the case of major failures in the network. Actually, cells adapt to the frequencies of the input messages *without* deciding on the legality of those frequencies. This may lead to a network alienation if those messages are more or less false, for instance in case that those messages' frequencies are subject to "mutations". If mutations persist, drives will stabilize to frequencies detrimental for the network's integrity. What do we do ? The solution is to equip the network with a possibility to distinguish between valid and invalid messages, that is to introduce an immune system[14] :

- The alien frequencies should be detected and the carrying messages eliminated.

It is easy to imagine that this situation corresponds to a *cooperation game* : the "intruder" frequencies ought to be recognized as different from the legal ones -

[14] For the development of the rudimentary immune model that follows, I was inspired by Golub and Green's (1991) overview of the immune mechanisms and theories and more specifically by the clonal selection theory. The term "identification template" should not mislead the reader into thinking that the adopted approach follows an instructive template theory (ibid.) ; to prevent confusion, in what follows we will use exclusively the term "eigenfrequency" that prompts to a selective model.

therefore as detrimental- to be eliminated. This observation leads us to a view of the cell's social behavior as a behavior based on similarity measures (*genetic kinship*, in terms of theory of evolution language). To this end, another type of cells has been introduced, the defensor cells, entitled to recognize and eliminate intruders. Note, however, that the initial idea of defensors recognizing the good eigenfrequencies and eliminating those with low affinity has been found faulty : in that case, and given that more than one eigenfrequencies are supported by the network, a frequency considered as an intruder by a certain defensor might be perfectly valid from the viewpoint of another one and the network would then be instable. The alternative is to have defensors that recognize directly the bad frequencies to eliminate : the defensor cells would therefore have eigenfrequencies values known to be bad, i.e. eigenfrequencies complementary to the valid ones. This formulation of the recognition problem has two important features :

- The cellular model remains selective. Defensor cells follow the same cellular model as the producer cells, with the only functional difference that they "swallow" the messages that excite them without metabolizing them.
- The network comprises therefore ***two complementary social forces***, "production" force, that is responsible for the completion of the agent's task, and the "supervision" or "police" force, that is responsible for the detection and the elimination of the alien messages for the same task.

To model this social cellular behavior, we used the quantitative tit-for-tat model (Tzafestas 1995*e*) according to which a social stimulus is only considered cooperative when its value is superior to a certain threshold ($f(stimulus) \geq T \Rightarrow cooperation$). This model allows the description of a whole class of social relations - *participation* as well as game relations - with variations in the parameters (such as the threshold T) as well as in the structures (such as the perception and cooperation functions).

The social behavior of the control system cells may be described in terms of this model as follows :

- *perception function* (or social stimulus) $f(x_1, ..., x_n)=min|x_i-e|$, where x_i are the frequencies of the input messages and e is the eigenfrequency of the drive (actually, the social adaptation affects all the eigenfrequencies of a drive),
- *cooperation function* : developmental adaptation, that is eigenfrequencies shift, plus execution of the transfer function in case of stimulation, like in the motivated model (a specific function for each drive belonging to a producer cell and a swallowing function for each drive belonging to a defensor cell),
- *defection function* : do nothing for both types of cells.

The final type of studied failures is alien messages that are not forced into the buffers by some external agent but are produced by capricious, malevolent or simply faulty cells. In that case, eliminating the messages is not sufficient ; on top of that, we need a way to prevent those cells from taking control of incoming messages and metabolize them abnormally.

- The network should have *self-correcting potential*, i.e. the malignant cells should converge back to a proper behavior.

To this end, we added the possibility of *internal drive adaptation* (adaptation/ convergence of the output frequency toward the input frequency) according to a social participation criterion :

$$participation < 0 \Rightarrow internal\ adaptation$$

where $participation = f_{output} - f_{input}$, $f_{input} = 1$ if the drive is stimulated and 0 otherwise, f_{output} is the proportion of output frequencies corresponding to messages present in the output buffers (actually, the mean of proportions concerning all the output buffers). The inequality holds if one or more output messages are consumed by the defensor cells.

The social cell behavior and the social network configuration may be summarized as follows :

- There are two parallel cellular systems : the producer cells and the defensor cells whose excitation frequencies are complementary[15]. The frequencies recognized by the defensors are detrimental for the social network's integrity. The defensors' transfer function is a function of swallowing of the stimulating messages (no message output).
- The social behavior of all drives follows a quantitative tit-for-tat model : the cooperative behavior is developmental adaptation, while the non-cooperative one is to do nothing (no social participation). All drives detect the messages' affinity to their own eigenfrequencies according a mono-dimensional distance citerion.
- There is an additional measure of social participation of drives that stimulates an internal adaptation process : if participation is negative (that is, if messages "produced" by the drive disappear quickly thanks to defensors intervention), the adaptation affects the output rather than the input frequencies.
- Once more, the power of the social cell model lies in its potential for large variations in parameters (thresholds, rates, frequencies etc.) as well as in structures (metabolic transfer functions and adaptation variants).

The issue of the population. The participation measure used is "binary" : if the drive participates normally to the process, then this measure will be 0, otherwise it will be inferior to 0, in which case "internal" adaptation will take place. Another possibility considered in perspective is to endow the cells with an additional speed regulation mechanism according to a measure of continuous participation depending on the proportions of input and output frequencies in the corresponding buffers. In this case, we will be able to talk about regulation by social feedback that will show as "division of labor" and that will become necessary in case of large populations with redundancy at all levels : as far as stability is concerned, an algorithmic solution - as well as a minimal solution - will suffer from occasional centralization (as is for example the case of the decoder cell) and could therefore not be able to profit from diversity which is the basis of stability in distributed systems. A diversity generator in combination with the cooperation game between the cells of the two types may give rise to phenomena that might be viewed as learned behaviors. An additional study that I intend to carry out concerns the dynamics of a producer-immune control network

[15] Note however that, as defined, the defensor cells do *not* constitute a system in the true sense of the term (there are no relations between them). In an extension of the model, it will be necessary to define a true defensor cell network parallel and structurally homomorphic to the production network.

especially designed for a problem of navigation with spatial learning : what kind of dynamics does it emerge ? And what are the necessary types of dynamics that may allow for such learning ?[16]

The limits. For the sensor and actuator cells that do not have input and output frequencies (and that are therefore situated in the frontiers of the agent with its world), it does not make sense to speak about developmental adaptation (in the implementation, the developmental adaptation factors of those cells' drives have been fixed to 0). The interaction modes with the world are fixed and correspond to the "physical" perception and actuation abilities : a mammal could decide to fly, but it could never fly. Those limits are therefore the cell's adaptation limits and by extension the learning limits of the cellular agent. I imagine that an artificial nervous system coupled with an artificial immune system would demonstrate the same inferiority in terms of adaptation possibilities.[17] I also imagine that it would be interesting to introduce completely internal boundary "nervous" parts, that is neither perception nor actuation nervous systems, that would determine fixed modes of esoteric interactions.

The world of representations. It is interesting to evoke once more the problem of the representation space : a diversity generator would have to provoke "mutations" whose gradual accumulation could take the form of the messages far away from their initial point. However, those mutations should not be too abrupt, otherwise the semantics of interaction between cells could collapse and the network to disintegrate. Resistance to "noise" is therefore necessary. But mutation and resistance to noise are not everything : it would be preferable to have multiple alternative trajectories between two points in this space, in which case the diversity generator would produce a range of mutated forms that would explore the space of the possibilities in parallel. The gain does not lie in the parallelism, though, but in the existence of multiple alternatives, hence of redundancy : in the space of real numbers there is unfortunately only one possibility to go from 0.3 to 0.5, the metrics of this space is linear. Once more, organic chemistry appears to me the perfect space for such a diversity.[18]

5 Illustration : The sickly robot

5.1 Buffer intruders and viruses

We have carried out a first experiment to show that even with a ***random intrusion*** there may be alienation of the network in the sense that certain drives stabilize to useless frequencies and one or more navigation directions (that are essential for the navigation task) disappear from the network. Random intrusion means that with a certain probability (the intrusion rate) the frequency of a message of the concerned buffer is mutated randomly ; the new frequency may therefore be valid or not for the network, but its generation remains stochastic. The results of a simulation with an

[16] Stewart et Varela (1991) stress and analyze the importance of two coupled dynamics, the "reaction" dynamics and the meta-dynamics that generate and select new reactions. In our model, it would be necessary to include a possibility of developmental "revision" of eigenfrequencies.

[17] I cannot stop thinking that there has to be a good reason why nervous cells don't proliferate.

[18] I would like to know which is the minimal "molecular" level that we need. I tend to think that messages would be more resistant to noise, if they were systems, too : a string of symbols, such as *abbcabxa* or *10001101101*, is not a system because there is no relation between the different symbols and each one may mutate independently of its neighbors - but this is not the case of proteins.

intrusion rate of 0.4 in the input buffer of one of the two diff levels are given in figures 17a and 17b. We can see that, 100 cycles later, the network is no more responsive to the right direction, since the drive that was responsible for this direction has been alienated. Generally, as a number of experiments have shown, a high intrusion rate results in developmental adaptation, such that there is a diversity of uncorrelated frequencies, i.e. there is no *meaning* for these frequencies.

We repeated the same experiment with the additional possibility of intrusion of virus, that are messages having an alien data part as well as an alien frequency part (an alien data part is a data part randomly mutated). As simulations have shown, this time, not only is the network alienated and certain frequencies disappear from it, but, since there may be messages with a good frequency and a bad data part, a second type of alienation exists, a *more semantic alienation* that causes collisions of the agent with the walls and the obstacles (see figure 18).

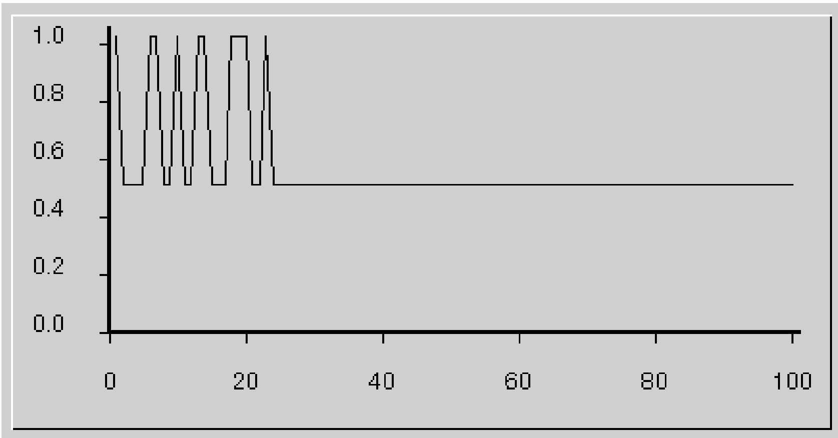

Figure 17a Alienation of the network by intruders : 100 cycles with intrusion. Loss of the reponse to the front direction with an intrusion rate of 0.4 in the buffer of one of the two diff buffers.

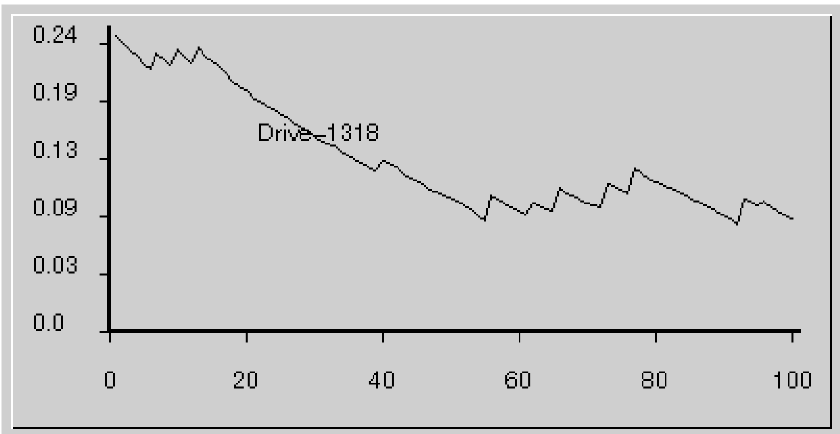

Figure 17b Alienation of the eigenfrequency of the drive that responded to the right direction (0.25) in the experiment of figure 17a.

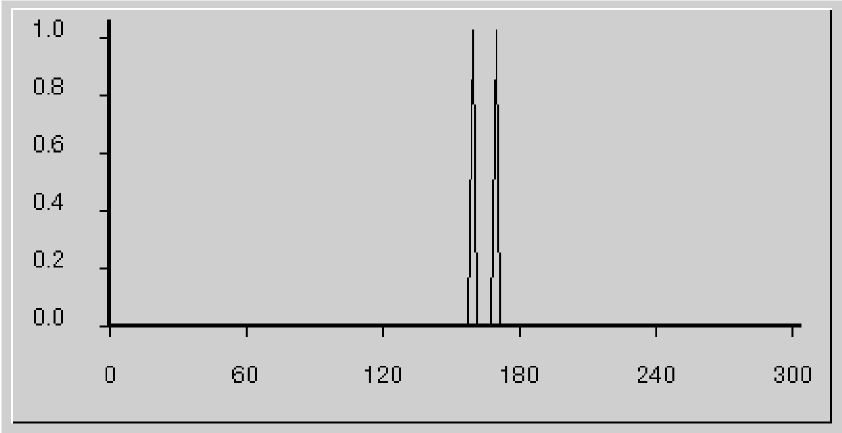

Figure 18 Alienation of the network by viruses : 300 cycles with an intrusion rate of 0.3 with virus possibility (buffer of one of the two diff levels). There has been a total of 9 collisions, 2 of which are visualized (the value of 1 corresponds to a collision).

We have therefore added around each of the buffers two to four defensor cells having one to four drives with eigenfrequencies taking one of the following values : 0.125, 0.375, 0.625 or 0.875. Those frequencies are situated in the middle of the controlled intervals, that is the middle of the intervals between successive operational frequencies. The cooperation detection thresholds of the defensors' drives have been defined in the same way, that is as approximately half of the controlled intervals, around 0.06. The intrusion of mutated messages or virus as before continued to cause partial destabilization of the network in the sense of occasional loss of one or more navigation directions, but did not lead to network alienation. In any case, to ensure the network's operationality, the defensors' drives ought to have high minimum speeds and their number should be sufficiently high.[19]

5.2 *Malignant cells and auto-immune response*

Finally, we experimented with faulty producer cells whose output messages were distorted : their frequencies as well as their data parts were taking random values (always between 0 and 1). In any case and given enough defensor cells, the messages carrying alien frequencies have been correctly eliminated. However, for the messages encapsulating a good frequency with a bad data part, results have been subtler : even if alienation of the network has not been possible, there are zigzags and bypasses and this up until the faulty cell discovers the proper pathway (often, the agent may no more be able to attain its goal-point). Figures 19a and 19b give the results of a simulation with such malignant cells.

This experiment shows that the sole existence of a passive "social" restabilization mechanism of malignant cells is not sufficient, and the reason is twofold :

- It is possible that the malignant cells have only a limited restabilization possibility (or even none!). This is the case of very low developmental rates (close to 0).
- The consequences of the loss of sensitivity to a given message type for such a long period may be fatal.

[19] Since defensor cells should not proliferate, their number should be high enough from the beginning. A good estimate to take is the size of the controlled buffer.

What is necessary is the possibility to recognize not only the bad message, but also the *source* of this message. This way, an isolated message, be it an intruder or a virus, will simply be eliminated. On the contrary, for a message coming from another cell in the network, this cell has to be recognized and "punished" in a way or another. Note that the "social" stabilization mechanism relies on such a "punishment" measure (the absence of participation to the activity of the social network). However, such "distributed" punishment is not enough : the social network should include an authority that will be charged with punishment. In the case of faulty cells, the authority, that is the immune network, should either intervene directly in the cells' program or simply kill them, in any case show direct aggressivity.[20] The first option would necessitate a rather complex intervention program, while the second appears to be a simple solution in the case of cells that cannot proliferate.[21]

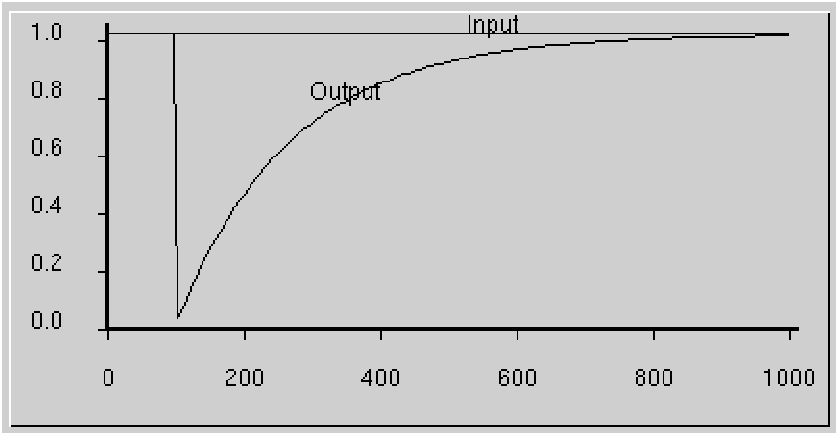

Figure 19a Re-stabilization of malignant cells. 100 cycles with a normal agent. At t=100, transformation of a "diff" drive that corresponded to the back direction into a deceiver. Very low developmental rate (approximately 0.0059). The evolution curves of input and output frequencies are given. The network re-stabilizes around t=1000, but the goal is always not attained.

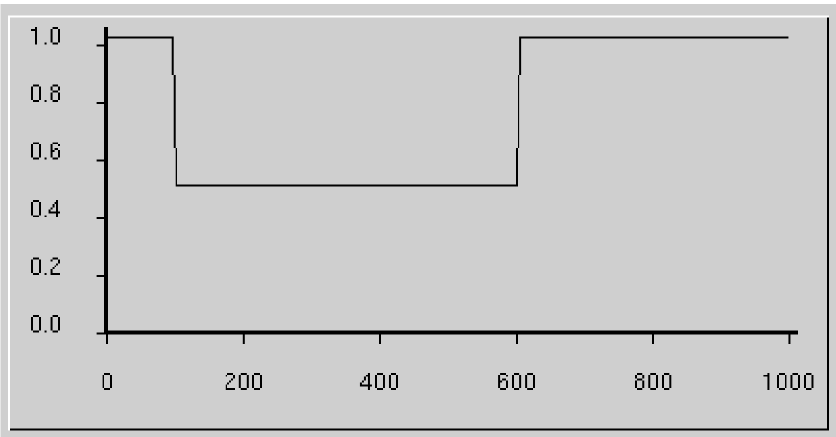

Figure 19b Loss of back direction for about 500 cycles in the experiment of figure 19a.

---

[20] I found the same result, that is the need for diect aggressivity in social relations, during my first experiments with the quantitative tit-for-tat model (Tzafestas 1995*e*).

[21] Stahl & Goheen (1963) discuss the form and nature of molecular algorithms and postulate that *there is no particular reason for which certains algorithms are not allowed, there is no violated principle! It is just a complexity constraint.* This argument is given for lamarckian adaptation but it may apply equally well to many other cases.

The adaptive and self-organizational abilities of immune systems have stimulated research in optimization and adaptive control (Bersini & Varela 1991, Bersini 1992). On the other hand, the idea to apply principles borrowed from immunology to robotics is not new. The biological adaptation systems in general (including the immune system) have been vaguely proposed as an inspiration source for the development of process control systems (Renders & Hanus 1992). The immune system paradigm is used in the action selection mechanism of (Ishiguro et al. 1995), which is a direct descendant of Maes' mechanism (1991). Our mechanism has been developped with the aim to reveal principles of cellular organization rather than to provide solutions to engineering problems ; the navigation problem has only illustrated these principles.

5.3 General discussion

We have exposed the philosophy behind plasticity and self-organization in an algorithmic cellular network. We have shown that, in order to allow this kind of plasticity, it has been necessary to do two major extensions to the original cell model : (a) introduction of local motivations/preferences in the cells, that are expressed in an uniform manner as excitation eigenfrequencies and that are compared to the identifiers of the input messages, and (b) replacement of the connections between levels by shared message buffers. Another extension consists in defining two message types, the usual consumable messages and the catalytic messages whose sole presence is enough to trigger the consumption of other messages by the excited cells ; all messages have two parts, an identification and a data part. Some simulation results reveal that the cellular network self-organizes in case of failures by discovering alternative message pathways and that multiple failures slow down the system's reactivity to external events, since the cells of intermediate levels are no more specialized to one message type but share there time between multiple motivations. Furthermore, the presence of an additional developmental/ontogenetical factor allows for self-organization even in the case of major failures in the network, that is when an eigenfrequency disappears completely from the network.

Note that the social behavior of the motivated cell is by definition "cooperative", that is the cell will try to stabilize to a new frequency, even if this frequency is detrimental for the network's integrity. What is necessary in this case, is *an immune system, that is a cellular system parallel to the first, supervising and regulating the "social" adaptation of the other cells*. These cells are motivated, too (they also ought to be social, but in the studied example their sociality has not been used) and their excitation eigenfrequencies are complementary to those of the producer cells, what has as a consequence that recognition and elimination of enemy-messages is straightforward. By default, the social behavior of all cells is cooperative, regardless of their type : ***the integrity of the network is therefore preserved thanks to the coupling of two social forces***, the production and the police. A study of the integrity of the network is therefore translated into a study in the meta level, the level of the *matching adaptive motivations network*, i.e. the network of the cells' eigenfrequencies. Note also that there is no need for a meta-immune system, since, in principle, the immune system is allowed to act upon itself [22] provided that defensor cells are sufficiently numerous and diverse.

For the set of the two systems (the production and the immune sytem) to be coherent and non-manipulable, the population of frequencies has to be developped by the agent itself during its life as a result of interactions in the network. The most important perspective of the presented work is the study of the co-development of the two networks *in the abstract level of the eigenfrequencies representation space*, that is in the level of information (be it unidirectional, like the real numbers space, or not) : I am convinced that a "learnt" behavior is in reality a "developped" behavior.

From selectivity to selectionism. If the cells/drives are selective, a learnt behavior will in fact be a behavior *generated* by the network itself and *reinforced* in a particular environment. Selectivity has to be combined with a diversity generator and an environmental amplification (or reinforcement) mechanism to allow for apparent learning. In the case of the autonomous cell model, it would be necessary to introduce a mechanism for eigenfrequency generation (or revision). There would then be a tug-of-war game between the two cell populations, in which the ones would try to generate new frequencies and the others to eliminate them ; in this case, the recognition and elimination of alien eigenfrequencies would be just a side-effect of the cells' social game. And the question would become : in a stable environment, does the network stabilize to a configuration ? Does it "learn" something ? Manderick (1992) analyzes and compares three selective systems (natural evolution, the immune system and the brain) stressing the importance of selectionism for the design of autonomous cognitive system in the context of the nouvelle (or behavior-based) AI. (Steels 1994) remains so far the only example of variation and selection inside the control system of a robot and is inspired by the selectionist brain theories developped by Edelman.

6 Conclusion : The cell, the animat and autonomy

This article presented the problem of a cellular organization of dynamic fault-tolerant topology. A network organization is cellular when the components have heterogeneous transfer functions, but the connections syntax is homogeneous. The basic input-output cellular model is extended to a *selective and motivated model* of autonomous cell. The extended model relies on the definition of the cell as a system of components, the drives, executing in parallel and being in constant competition with one another for the consumption of the messages received by the cell. While drives are responsible for the messages consumption, the cell's physiology is responsible for the coordination and the conflict resolution between drives. In order to work, the mechanism necessitates compound messages, that comprise an identification and a data part. Some of these messages ought to be catalysts that will trigger the metabolic activity (the reactions) of other cells without being literally consumed. The operationality of the model and its robustness to various types of failures are ensured by an additional mechanism that modifies the dynamics of the interaction of the drives with the world of the messages ; this modification concerns the parameters that

[22] Pitrat (1991) insists that all knowlege is situated at a meta level and that one single level of meta-knowledge with the possibility to act upon itself, that is with reflexivity, suffices for the development of cognitive abilities comparable to those of humans and even greater. Inversely, the absence of a regulatory meta level may have disastrous consequences : AIDS attacks the meta part of the immune system.

determine the dynamics of agent-world interaction (here, message consumption, or reaction, speeds) and is homeostatic. Sociality shows as the adaptation of the cell to the form of the encountered messages, thus as a sort of "goal revision" of the cell. In the case of alien messages and viruses that may alienate the network, and make it give false answers to its perturbations, a cellular defense system parallel to the first one, i.e. an immune system, appears necessary. The defense system relies on the recognition and elimination of alien messages. The coupling between the immune and the production system shows in the complementarity of the messages recognized by the two systems, it is hence a coupling between two social forces within the same cellular agent. The problem of the representation space is also evoked and briefly discussed in a future learning perspective.

The essence of the agent's autonomy is the continuous change, the continuous modification or adaptation of the dynamics of interaction with the world (see also van Gelder & Port (1995)). Change or adaptation goes hand in hand with the agent's selectivity, and its purpose is to make the agent increasingly selective to the relevant aspects of the environment and decreasingly selective to the irrelevant ones. This adaptation is in fact a self-regulation of the parameters determining this dynamics, hence the self-regulation of parameters having a temporal character, rates or speeds. Concerning the relation between autonomy and self-regulation of adaptation rates, the same result has been obtained at the cellular agent level (chapters 4 and 5 of (Tzafestas 1995*d*)) and supports our claim for *recursive organizational principles*, i.e. for organizational principles applicable at all agent levels. Of course, adaptation possibilities are by no means limited to parameter regulation and further alternatives are considered in perspective -in particular the developmental approach, especially when coupled with an immune control system, as discussed in section 5.3.

In the case of the cellular self-organization mechanism, to ensure continuous change, it could be useful to have a mechanism inducing failures and thus injecting noise continuously (for instance a senescence function, such as the one developped in Tzafestas (1995*f*)).[23]

Finally, it is important to stress once again the role of diversity and self-organization inside the agent ; besides the pure operationality considerations presented above, we have sketched the relation of diversity with selectionism and the perspectives that open up in the direction of developmental learning.

References

Almássy, N., P. Verschure (1992). Optimizing self-organizing control architectures with genetic algorithms: The interaction between natural selection and ontogenesis, *Technical Report No. 92.10, Department of Computer Science, University Zurich-Irchel*, 10p, also *Proceedings of the Second Conference on Parallel Problem Solving from Nature (PPSN '92)*, R. Männer & B. Manderick (Eds.), North-Holland, September 1992.

Anderson, J.A., E. Rosenfeld (1988). *Neurocomputing, Foundations of research*, MIT Press.

Anderson, J.A., A. Pellionisz, E. Rosenfeld, Eds. (1991). *Neurocomputing 2, Directions for research*, MIT Press.

[23] The organizational role of noise has been underlined by various researchers, for instance Atlan (1979).

Arkin, R. (1987). Motor schema based navigation for a mobile robot : An approach to programming by behaviour, *Proceedings 1987 International Conference on Robotics and Automation*, Raleigh, NC, March.

Arkin, R.C. (1989). Motor-schema based mobile robot navigation, *International Journal of Robotics Research*, 8(4):92-112.

Atlan, H. (1979). *Entre le cristal et la fumée - Essai sur l'organisation du vivant*, Éditions du Seuil, Paris.

Beer, R. (1990). *Intelligence as adaptive behavior - An experiment in computational neuroethology*, Academic Press, San Diego, CA.

Bersini, H., F.J. Varela (1991). The immune recruitment mechanism: A selective evolutionary strategy, *Proceedings ICGA 1991*.

Bersini, H. (1992). Immune network and adaptive control, in F.J.Varela & P.Bourgine (Eds), *Toward a practice of autonomous systems, Proceedings of the First European Conference on Artificial Life*, Paris, December, MIT Press/Bradford Books.

Brooks, R.A. (1986). A robust layered control system for a mobile robot, *IEEE Journal of Robotics and Automation*, 2(1):14-23.

Brooks, R.A. (1990). The Behavior Language - User's Guide, *M.I.T. AI Memo No. 1227*, 35 p.

Brooks, R.A. (1991*a*). Intelligence without Representation, *Artificial Intelligence, Special Issue on the Foundations of Artificial Intelligence*, 47(1-3):139-159.

Brooks, R.A. (1991*b*). Elephants don't play chess, in P. Maes (Ed.), *Designing Autonomous Agents : Theory and Practice from Biology to Engineering and Back*, MIT Bradford Press, reprinted from Special Issue of the Journal on Robotics and Autonomous Systems, 6(1&2), June 1990.

Clark, A. (1993) Minimal rationalism, *Cognitive Science Research Paper 270*, Univ. of Sussex, 19p.

Darche, P. (1994). Le paradigme acteur appliqué aux systèmes embarqués communicants : ActNet, un réseau d'acteurs robotiques, *Thèse de Doctorat de l'Université Paris VI*, Mars.

Drogoul, A. (1993). De la simulation multi-agents à la résolution collective de problèmes - Une étude de l'émergence de structures d'organisation dans les systèmes multi-agents, *Thèse de Doctorat d'Université Paris VI*, November.

Edelman, G.M. (1992). *Bright air, brilliant fire: On the matter of mind*, Basic Books, 1992, traduit en français sous le titre *"Biologie de la conscience"*, Éditions Odile Jacob.

Ferrell, C. (1994). Failure recognition and fault tolerance of an autonomous robot, *Adaptive Behavior*, 2(4):375-398.

Golub, E.S., D.R. Green (1991). *Immunology - A synthesis*, 2nd Edition, Sinauer Associates, Sunderland, MA.

Guillot, A., J.-A. Meyer (1994). Synthetic animals in synthetic worlds, *Groupe de BioInformatique, École Normale Supérieure*, Paris.

Hogg, T., B.A. Huberman (1993). Better than the best: The power of cooperation, in L. Nadel and D. Stein (Eds.), *SFI 1992 Lecture Notes in Complex Systems*, pp. 163-184, Addison-Wesley.

Husbands, P., I. Harvey, D. Cliff (1993). Analysing recurrent dynamical networks evolved for robot control, *Cognitive Science Research Paper 265*, Univ. of Sussex, January.

Ishiguro, A., Y. Watanabe, Y. Uchikawa (1995). An immunological approach to dynamic behavior control for autonomous mobile robots, *Proceedings IROS'95*, Pittsburgh, PA.

Kaelbling, L.P. (1991). Foundations of learning in autonomous agents, *Robotics and Autonomous Systems*, 8(1991):131-144.

Kauffman, S.A. (1993). *The origins of order - Self-organization and selection in evolution*, Oxford University Press.

Kephart, J.O. (1994). How topology affects population dynamics, pp. 447-463, in C.G.Langton (Ed.), *Artificial Life III*, Addison-Wesley, Reading, MA.

Kuipers, B., Y. Byun (1991). A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations, *Robotics and Autonomous Systems*, 8:47-63.

Maes, P. (1991). A Bottom-Up Mechanism for Action Selection in an Artificial Creature, in J.-A. Meyer & S.W.Wilson (Eds.), *From Animals to Animats, Proceedings of the First International Conference on Simulation of Adaptive Behavior*, Paris, 1990, Bradford MIT Press.

Manderick, B. (1992). Selectionist systems as cognitive systems, in F.J.Varela & P.Bourgine (Eds), *Toward a practice of autonomous systems, Proceedings of the First European Conference on Artificial Life*, Paris, December, MIT Press/Bradford Books, pp. 441-447.

Mataric, M.J. (1990). Environment learning using a distributed representation, *Proceedings 1990 IEEE International Conf. on Robotics and Automation*, Cincinatti, Ohio, May.

Mataric, M.J. (1991). Navigating with a rat brain : A neurobiologically-inspired model for robot spatial representation, in J.-A. Meyer & S.W.Wilson (Eds.), *From Animals to Animats, Proceedings of the First International Conference on Simulation of Adaptive Behavior*, Paris, 1990, Bradford MIT Press.

Mataric, M.J. (1992). Integration of representation into goal-driven behavior-based robots, *IEEE Transactions on Robotics and Automation*, 8(3):304-312.

Monod, J. (1970). *Le hasard et la nécessité - Essai sur la philosophie naturelle de la biologie moderne*, Éditions du Seuil, Paris.

Morin, E. (1980). *La méthode 2. La vie de la vie*, Éditions du Seuil, Paris.

Pitrat, J. (1991). *Métaconnaissance - Futur de l'Intelligence Artificielle*, Hermès, Paris.

Ray, T.S. (1991). An approach to the synthesis of life, pp. 371-408, in C.G.Langton, C.Taylor, J.D.Farmer & S.Rasmussen (Eds.), *Artificial Life II, Proceedings of the Workshop on Artificial Life*, February 1990, Santa Fe, New Mexico, Addison-Wesley.

Ray, T.S. (1994). An evolutionary approach to synthetic biology: Zen and the art of creating life, *Artificial Life*, 1(1&2):179-209.

Renders, J.-M., R. Hanus (1992). Biological learning metaphors for adaptive process control : A general strategy, *Proceedings of the 1992 International Symposium on Intelligent Control*, Scotland, August.

Resnick, M., F. Martin (1990). Children and Artificial Life, *M.I.T. E&L Memo No. 10*, November.

Stahl, W.R., H.E. Goheen (1963). Molecular algorithms, *Journal of Theoretical Biology*, 5(1963):266-287.

Steels, L. (1994). A case-study in the behavior-oriented design of autonomous agents, pp. 445-452, in D.Cliff, P.Husbands, J.-A. Meyer & S.W.Wilson (Eds.), *From animals to animats 3, Proceedings of the 3rd International Conference on Simulation of Adaptive Behavior*, Bradford/MIT Press, Cambridge, MA.

Stewart, J., F.J. Varela (1991). Morphogenesis in shape-space, elementary meta-dynamics in a model of the immune metwork, *Journal of Theoretical Biology*, 153(1991):477-498.

Tzafestas, E.S. (1995*a*). Variation et auto-organisation à l'intérieur du robot: Le cas d'un modèle de fréquences propres, pp. 181-192, *Actes des Journées de Rochebrune "Évolution et organisation"*, Rochebrune, France, March.

Tzafestas, E.S. (1995*b*). Compromising algorithmicity and plasticity in autonomous agent control systems: The motivated, social cell*, Accepted for a poster presentation 1995 European Conference on Artificial Life (ECAL'95)*, Granada, Spain, June.

Tzafestas, E.S. (1995*c*). Peripheral cellular control: An eigen-frequency model and a case-study in self-organization, pp. 483-488, Vol. I, *Proceedings 1995 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'95)*, Pittsburgh, PA, August.

Tzafestas, E.S. (1995*d*). Vers une systémique des agents autonomes : Des cellules, des motivations et des perturbations*, Thèse de Doctorat de l'Université Pierre et Marie Curie*, Paris, December.

Tzafestas, E.S. (1995*e*). Beyond cooperation and competition : Explorations with a quantitative tit-for-tat model*, LAFORIA Research Report*, to appear.

Tzafestas, E.S. (1995*f*). Aging agents*, LAFORIA Research Report*, to appear.

van de Welde, W. (1991). Toward learning robots, *Robotics and Autonomous Systems*, 8(1991):1-6.

van Gelder, T., R. Port (1995). It's about time : An overview of th dynamical approach to cognition, *Research Report 116*, Indiana University at Bloomington, Dept. of Cognitive Science, October 1994, 39p., in *"Mind as Motion : Explorations in the dynamics of cognition"*, R. Port & T. van Gelder (Eds.), Bradford/MIT Press.

Varela, F. (1979). *Principles of biological autonomy*, Elsevier/North-Holland, New York, French edition *"Autonomie et connaissance : Essai sur le vivant"*, Editions du Seuil, Paris.

Varela, F.J., E. Thompson, E. Rosch (1991) *The embodied mind - Cognitive science and human experience*, MIT Press.

Verschure, P.M.F.J., B.J.A. Kröse, R. Pfeifer (1992). Distributed adaptive control : The self-organization of structured behavior, *Robotics and Autonomous Systems*, 9(1992):181-196.

Verschure, P.M.F.J., R. Pfeifer (1992). Categorization, representations, and the dynamics of system-environment interaction : A case study in autonomous systems, pp. 210-217, in J.-A. Meyer, H.L.Roitblat & S.W.Wilson (Eds.), *From Animals to Animats 2, Proceedings of the Second International Conference on Simulation of Adaptive Behavior*, Hawaii, 1992, Bradford/MIT Press.

Yamauchi, B.M., R.D. Beer (1994). Sequential behavior and learning in evolved dynamical neural networks, *Adaptive Behavior*, 2(3):219-246.