



# TOWARDS SELF-ORGANIZING BUREAUCRACIES

CARLOS GERSHENSON  
New England Complex Systems Institute  
24 Mt. Auburn St. Cambridge MA 02138, USA  
[carlos@necci.org](mailto:carlos@necci.org)  
and  
Centrum Leo Apostel  
Vrije Universiteit Brussel  
Krijgskundestraat 33 B-1160 Brussels, Belgium  
[cgershen@vub.ac.be](mailto:cgershen@vub.ac.be)

## Abstract

The goal of this paper is to contribute to eGovernment efforts, encouraging the use of self-organization as a method to improve the efficiency and adaptability of bureaucracies and similar social systems. Bureaucracies are described as networks of agents, where the main design principle is to reduce local "friction" to increase local and global "satisfaction". Following this principle, solutions are proposed for improving communication within bureaucracies, sensing public satisfaction, dynamic modification of hierarchies, and contextualization of procedures. Each of these reduces friction between agents (internal or external), increasing the efficiency of bureaucracies. Current technologies can be applied for this end. "Random agent networks" (RANs), novel computational models, are introduced to illustrate the benefits of self-organizing bureaucracies. Simulations show that only few changes are required to reach near-optimal performance, potentially adapting quickly and effectively to shifts in demand.

Keywords: eGovernment, self-organization, adaptation, communication, hierarchies

## 1. Introduction

Bureaucracies can be found in governments, corporations, and other social institutions. They have social goals and responsibilities that are achieved by a division of labour that is usually hierarchical. Examples of bureaucracies can be seen in the public and private sectors, e.g. tax collection systems, immigration services, and steering of educational and academic institutions. The efficiency of a bureaucracy is related to the fulfilment of its goals. Thus, it would be desirable to increase functional efficiency in bureaucracies. Ideally, such a system could be designed to reach maximal efficiency. In practice, as most people have experienced, this is far from being the case [Weber, 1968]. Corruption, rigidity, and delays are just few examples of obstacles that hamper efficiency in bureaucracies. It would be naïve to aim for perfect bureaucracies, but certainly the efficiency of actual ones can be improved.

One approach would consist of *optimizing* the bureaucratic functionality, e.g. [Hofacker and Vetschera, 2001]. This approach can provide good solutions if the function or goal of the organization does not change considerably, i.e. when a problem domain is stationary i.e. static. However, the world is changing at

accelerating rates, modifying constantly the requirements of organizations and bureaucracies. Changes cause the shifting of the optimum of a system. And in some cases, the behaviour of the institution itself changes the optimum [Kauffman, 2000]. Thus, a wiser approach would be to design bureaucracies that are able to *adapt* [Holland, 1975] to changing situations. Instead of attempting to predict all the functionality beforehand, an organization could adapt to the changing demands of its environment.

Cybernetics and systems theory proposed some of the first solutions in this direction already a few decades ago, e.g. [Beer, 1966]. The Cybersyn project in Chile was even partially implemented, but was cut short by the 1973 military coup [Miller Medine, 2005].<sup>1</sup> However, this approach is still not widely used in practice, probably because it requires alternative ways of thinking. It is always easier to solve problems for static domains than for dynamic environments.

Organization science has developed several concepts that are useful for improving the self-organization and adaptation of bureaucracies. Noting the cognitive limits of decision makers [March, 1978; Simon, 1982; Cyert and March, 1992] tells us that individuals will not be able to make perfect decisions. This limitation can be overcome in different ways. First, the cognition necessary to solve complex tasks can be distributed [Hutchins, 1995; Weick and Roberts, 1993]. Second, organizations can be designed to adapt to unpredictable events [Carley, 1997]. Organizational learning is one type of adaptation that has been widely studied [Levitt and March, 1988; March, 1991]. Also, computational organization theory [Carley and Prietula, 1994] and agent based modelling, e.g. [Epstein and Axtell, 1996; Axelrod, 2005], have aided in the understanding of the complexity inherent to organizations [Anderson et al., 1999; Anderson, 1999; Levinthal and Warglien, 1999; Lissack, 1999; Axelrod and Cohen, 2000].

Self-organization [Heylighen, 2003] has been used as a principle in many domains such as computer science and robotics [Kohonen, 2000], the Internet [Bollen and Heylighen, 1996], and traffic light control [Gershenson, 2005], just to name a few. In organization science, self-organization has been studied as a phenomenon, e.g. [Comfort, 1994; Morgan, 1996, p.233]. The goal of this work is to use it conceptually as a *tool* to improve the efficiency of bureaucracies.

Even when it has been widely studied and applied in many domains, self-organization is a concept difficult to define [Gershenson and Heylighen, 2003]. Nevertheless, for practical purposes, a notion can suffice: A system *described* as self-organizing is one in which elements *interact* in order to achieve *dynamically* a global function or behaviour [Gershenson, 2007]. A classical example can be seen with flocks of birds, or schools of fish: there is no leader in the group, and all individuals follow local rules, interacting with their neighbours, and only this produces coherent global behaviour. In a similar way, local rules of interaction can be designed, to produce dynamically robust and adaptive behaviours. In this work, some of such rules are proposed, as a way to improve the performance of bureaucracies.

Following these lines of research, this work suggests methods that improve the efficiency of bureaucracies via self-organization, exploiting current technology. In the next section, general ideas for the design of self-organizing systems are presented. In the following sections, different aspects of bureaucracies, and self-organizing

---

<sup>1</sup>One aspect of the Cybersyn system implemented a "nervous system" for the country, where every day information with the productions and demands from factories was telegraphed to a "brain" room where people decided which demands were more urgent and route the productions accordingly. Like this, the government was able to stock the country in spite of widespread strikes and scarcity.

improvements are proposed, namely in the areas of communication, sensors, hierarchies, and context. Afterwards, random agent networks (RANs) are defined to model bureaucracies. Computational experiments with RANs illustrate the benefits of self-organization for improving the performance of abstract bureaucracies. The paper is aimed more at government bureaucracies, but in principle the ideas could be applied to business bureaucracies.

## 2. Designing Self-Organizing Systems

Organizations can be modelled as systems of information processing agents [Radner, 1993; Van Zandt, 1998; DeCanio and Watkins, 1998]. An agent can be defined as a description of an entity that *acts* on its environment [Gershenson, 2007, p.39]. They could also be described as *cognitive* systems [Gershenson, 2004a]. Thus, not only people can be described as agents, but also departments, ministries, and governments. Agents can have goals [Simon, 1964], that are described by an observer. Agents can be said to be cognitive because they need to “know” which actions to take to reach their goals. We can define the agent’s “satisfaction” as related to the achievement of their goals: the better the goals are achieved, the higher the satisfaction. The goals are determined by the expected functionality, e.g. the goal of a school is to provide the best education possible for its students. Also, the efficiency is determined by the expected functionality: a bureaucracy is efficient if it does what it is supposed to do while minimizing costs. Thus, a description of a bureaucracy can be made in terms of agents trying to fulfil goals to increase their satisfaction, leading to higher efficiency. The public can also be described as an agent or several agents, interacting (externally) with the bureaucracy. The concepts discussed in this section have been formalized in [Gershenson, 2007, Ch. 4], but for the purposes of this paper, their description will suffice.

The satisfaction of one agent (e.g. a clerk) can be in conflict with the satisfaction of another agent (e.g. the minister). It can be argued that decreasing the “*friction*” or interference of agents at one level (e.g. personnel level), i.e. how one agent decreases the satisfaction of another agent, will result in an increase of satisfaction at the higher level (e.g. ministry level) [Gershenson, 2007; Helbing and Vicsek, 1999]. Higher satisfaction implies a better fulfilment of goals of agents. If one agent at the lower level is not satisfied, i.e. does not accomplish its goals, the satisfaction of the system may be reduced. However, if all agents at the lower level fulfil their goals, then the satisfaction of the system should be maximal. This is a (useful) tautology because the goals of agents are described by observers according to the desired function of the system. Notice that this is different from implying that increasing the satisfaction of agents at one level will always lead to an increase of satisfaction at the higher level. The key difference lies in the mediation of conflictive goals to increase satisfaction. Similarly, we can speak about the negative (opposite) of friction, or *synergy* [Haken, 1981]<sup>2</sup>, where the behaviour of one agent increases the satisfaction of another agent. Certainly, not only friction should be minimized, but also synergy maximized. Different ways in which this can be achieved are discussed in detail in [Gershenson, 2007, Ch. 4]. Within practical limits, friction reduction and synergy promotion will be always useful, since all actors, internal or external, will be benefited.<sup>3</sup> For example, if it becomes easier to pay taxes, more people will be more inclined to pay them. This

---

<sup>2</sup>Friction and synergy are defined here as complementary.

<sup>3</sup>Note that high costs cause friction (cost reduction itself increases satisfaction), so self-organization should also take cost into account.

benefits the state (more money collected) and taxpayers (less time lost). Certainly, as in most social systems, a problem will remain when it comes to measuring satisfaction. This will be discussed in Section 4.

The goals of a firm can be easily related to its profits [Van Zandt, 1998]. However, the goals of a bureaucracy are related to its particular function. This function can be codetermined by the state, by the public, and by the bureaucracy itself. Thus, there is yet no general way to measure the performance of a bureaucracy. Efficiency could be a way of evaluating a bureaucracy, but there is the same measurement problem with efficiency: it will differ accordingly to the particular bureaucracy. Still, a lack of explicit descriptions of function, efficiency, or satisfaction are not a limit for speaking about the goals of a bureaucracy. It should just be considered that these can change depending on the behaviour and context of the bureaucracy itself.

In order to adapt to unpredictable changes, bureaucracies require a certain flexibility. Changes should be made, but the function needs to be preserved. Robustness is required [Jen, 2005], so that adaptive changes do not prevent the bureaucracy from reaching its goals. The main idea to guide changes is the following: First, detect how each agent affects satisfaction of others. Then, implement changes to minimize friction and promote synergy. This can be achieved by reinforcement: behaviours that have proven themselves inefficient should be avoided, and beneficial ones should be promoted. This can be seen as a particular case of *learning* [Levitt and March, 1988; March, 1991].

Before implementing radical changes that might lead to unexpected outcomes, computer simulations should be used [Axelrod, 2005]. These will be useful for detecting possible flaws in the changes planned, or simply to improve them. Moreover, the changes themselves can be explored with the aid of computer simulations, since it is not obvious in every case what should be done, as the complexity of bureaucracies exceeds our predictive capabilities. Simulations are important because the behaviour of highly complex systems cannot be predicted beforehand, but it should be observed. Simulations allow observation without potential risks.

The changes could be introduced gradually and with a certain redundancy to compare the benefits and disadvantages of the new methods with the previous ones, with the possibility to return to previous situations.

In the following sections diverse aspects of bureaucracies are explored, suggesting different possible improvements within each domain. Each improvement would contribute to the performance of a bureaucracy in a different manner, so their application could be considered combined or in isolation. Certainly, these sections are not exhaustive. There are many other aspects of bureaucracies that could be improved.

### 3. The Role of Communication

Communication between agents can be classified in two categories: synchronous and asynchronous [DeSanctis and Monge, 1999].

*Synchronous* communication occurs when the agents involved in the process are responding at the same time. There is immediate feedback between speaker and listener, so that a dialogue can be established continuously. Examples of this are verbal communication, telephone conversations, video conferencing, and IRC (Internet relay chat). The advantage of this mode of communication is that dialogues can be resolved without interruption. The disadvantage is that all participants need to coordinate to participate in the process.

*Asynchronous* communication occurs when the agents involved do not participate simultaneously in the process. There is a delayed feedback between agents, so that dialogues are interrupted depending on the length of the transmission delay. Examples include post, telegraph, telex, fax, e-mail, and instant messaging. The tradeoffs of this mode of communication complement those of synchronous communication: on the positive side asynchronous communication allows exchanges without coordination required, but on the negative side the communication can be delayed. Technological development has reduced transmission delay, enabling asynchronous communication to depend only on the constraints of the agents. Moreover, asynchronous communication allows for certain decision time, whereas in synchronous communication most responses should be immediate, reducing the possibility of digesting information properly.

Technology seems to lead to a convergence of synchronous and asynchronous communication into *semisynchronous* communication, where media can be used synchronously or asynchronously depending on the circumstances. This is the case with e-mail, instant messaging, IRC, and SMS: instant replies can lead to a conversation, but the information persists in case one of the parties is not able to respond immediately.

In a bureaucracy, different agents need to communicate to satisfy the goals of the system. Thus, communication delays can be seen as a type of friction between agents. The faster the communication takes place, the better it will be for the system. Thus, synchronous communication might be preferred to enable quick responses. However, this would imply a great coordination effort, since agents usually perform other activities apart from communicating. It could be quite possible that an agent would be busy with other matters to have a synchronous exchange. Then it seems that asynchronous communication would be preferred, since one agent can send a message and keep on working on other matters while the response arrives. Then the question would be: how can asynchronous communication be facilitated?

As mentioned above, one great improvement is given by technology. Being able to send documents electronically instead of postally reduces the delay of message transmission from the scale of days to the scale of seconds. Certainly, organizations have exploited this opportunity, and worries about security have been solved with digital signatures. Still, it is a common practice in several bureaucracies to handle paper documents, even when they must be sent across continents, as it is the case e.g. with the Mexican foreign services.

However, the adoption of electronic means of communication can do much more than reducing the transmission delay of messages. Analyzing the times when a message is sent and when it is replied can provide very useful information, namely that of response delay (see Figure 1). This can be used to detect bottlenecks: If one agent (individual or department) takes too long in replying requests, the work of other agents might be delayed as well, as in a production chain. Resources could then be reassigned in real time to overcome the bottleneck, giving priority to the agent with the response delay. In fact, we can say that a delay in response causes friction to other agents, since they need the feedback to reach their goals. Thus, a bureaucracy could self-organize by modifying in real time its own structure, once it is known where friction is coming from. Solutions can vary depending on the precise nature of the delay: assign more individuals to a department, replace individual(s), or reorganize departments. Like this, efficiency of the bureaucracy would be improved, adapting without needing to wait until somebody complains about its failures and only then consider doing something about it. It would be self-organizing, because the changes

are dictated by the behaviour of the bureaucracy itself. The changes would imply *learning* in the organization from its experience, while enabling it to adapt constantly to changes of its environment.

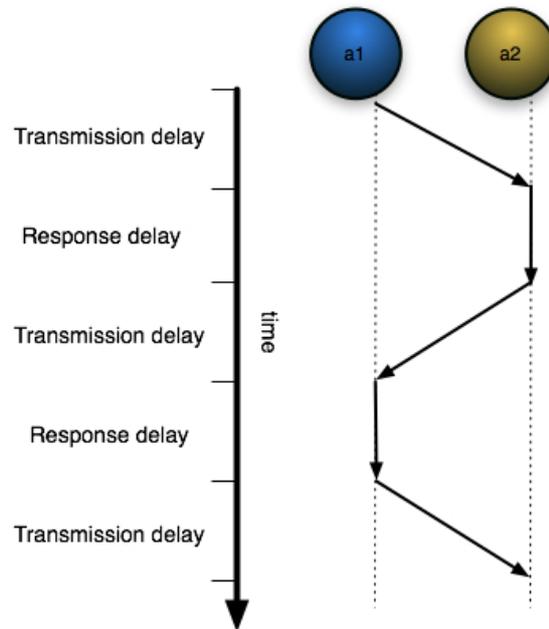


Figure 1. Asynchronous communication. Technology has reduced transmission delays, and can help to detect and decrease response delays.

The response delay would depend on several factors: decision delay (the time it takes the agent to respond), delay from previous tasks (the time it takes an agent to start making a decision), and delay from other responses (the time it takes other agents to respond to the agent's requirements) (see Figure 2). Each of these delays should be taken into account while modifying the bureaucracy.

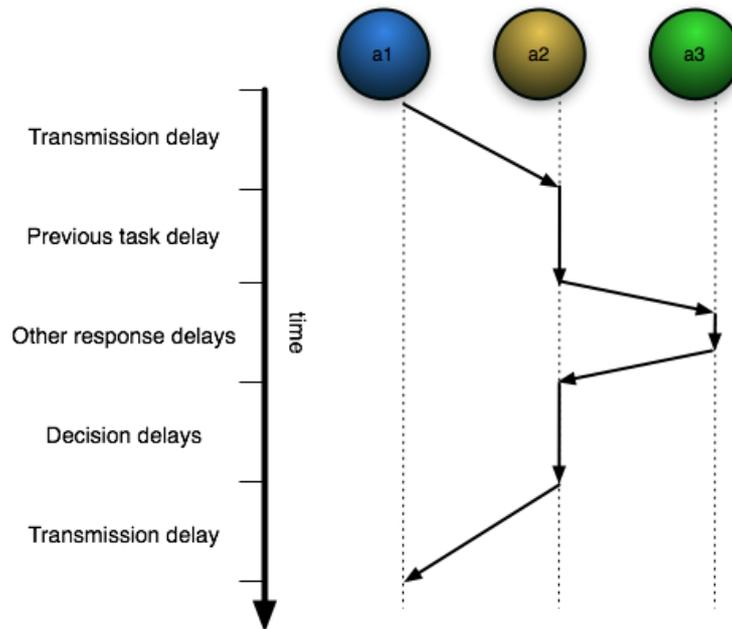


Figure 2. Response delay can be decomposed in previous task delay, other response delay, and decision delay.

Another benefit of electronic logs is that they can be used to provide accountability of decisions, as it was the case for the Enron e-mail dataset<sup>4</sup>. The workload of individual agents or departments could also be measured by the number of requests they send or receive, considering the decision delays required by each request. Finally, a visualization of the interactions within the bureaucracy (who communicates with who) could provide insights to improve its design, for example detecting redundant agents or interactions, or creating “shortcuts” between agents that communicate frequently via other agents.

### 3.1. Decision Delays

Technology has also aided in the reduction of decision delays, which lead to friction. Electronic databases provide instant information, while in a physical repository a clerk has to search an archive for documents. The role of the clerk is taken by software in an electronic database. Similarly, monotonous decisions can be taken by computer systems near instantly, reducing decision delays. An example of this can be seen with bank credit evaluation [Hand and Henley, 1997], where a computer system can give instant decisions on whether to give credit to an applicant or not. Similar methods could be used to make instant decisions to judge e.g. visa applicants or prospective students. Turning decisions into computer systems will reduce decision delays, thus reducing friction, improving communication, and increasing the bureaucracy efficiency. This is precisely one of the directions agent technology is taking [Luck et al., 2005], using notions of negotiation, trust, and reputation to facilitate the coordination of electronic decision makers. Also research and technology applied to eGovernment [Layne and Lee, 2001] and automated [Bowman, 1963; Dawes and Corrigan 1974] or computer aided decision making [Turban and Aronson, 1997; Stahl, 2006] will improve the performance of bureaucracies.

<sup>4</sup>See <http://www.cs.cmu.edu/~enron/>

Such a hybrid scenario, where humans and agents interact in an organization, has been described with the term "cognitive stigmergy" [Ricci et al., 2006]. Stigmergy is used to describe systems, such as insect colonies, that exploit their environment to communicate and coordinate [Theraulaz and Bonabeau, 1999]. In a similar way, computer systems can be used as an environment to facilitate the communication, coordination, and decision of agents.

## 4. The Role of Sensors

In the previous section, reduction of friction within the bureaucracy was discussed. In this section, reduction of friction between the bureaucracy and its environment, namely the public, will be discussed.

Much research has been made in decision making, e.g. [Simon, 1976]. It is clear that without proper sensors there will not be enough information to make proper decisions. Still, even with simple sensors, a system can obtain much information from its surroundings. An example can be seen with people who perceive their environment with a walking stick, sensing by pressure only one point in space. Integrating information in time, they are able to obtain relevant information to navigate through complex areas. Nevertheless, complex sensors can reduce the complexity of a decision making process, by "digesting" relevant information. Therefore, bureaucracies should aim at developing fine sensors to perceive and digest information relevant to their goals. In any case, without proper sensors, no self-organization nor adaptation can take place.

One element that facilitates the sensing process is public participation, since people themselves digest and feed the information to the bureaucracy. However, many people are reluctant to participate in such processes, since they do not see any benefit for it, while it takes some of their time. An alternative would be to reconstruct public opinion from a limited set of the population, as polls have been doing, and novel methods could improve, e.g. [Rodriguez and Steinbock, 2004]. Here, only improvement of sensors that do not require public participation will be discussed.

In order to measure the efficiency of a bureaucracy, sensors should be used. A popular variable related to this efficiency is public satisfaction: if the public is happy with the services provided by the bureaucracy, then its efficiency can be assumed. Polls have been used to measure public satisfaction, but demand a certain effort from public and resources to design and analyze them. Also, they cannot measure all possible mishaps.

Thus, bureaucracies should develop sensors for public satisfaction that do not require public participation. This could be done measuring the public attention delay, which would be the sum of the waiting delay (how much time a person needs to queue) and the procedure delay (how much time a person needs to interact with the bureaucracy). Another indicator would be the frequency of interaction, namely how many times the same person needs to interact with the bureaucracy. These delays can be considered as friction between the bureaucracy and the public, and should be minimized. Both the public and the bureaucracy will be satisfied if they need to interact with each other as few as possible (low interaction frequency), and each of these interactions takes as little time as possible (low public attention delay). Like this, the precise places where bottlenecks arise, and for which cases, can be detected, and measures can be taken. For example, if a procedure for a special type of license takes consistently more time than others, this procedure should be revised and adapted.

One could argue that bureaucracies do not need to care for the public, since they can be considered as monopolies. But the tendency towards improving bureaucratic services refutes this argument. For example, eGovernment would practically eliminate waiting delays. It is beneficial for political parties in office to improve bureaucratic services to increase public satisfaction, and thus more votes for the next election. Natural selection will give better chances of survival to political parties that satisfy public demands. Certainly, this can only happen in countries with a certain degree of political diversity. Otherwise, indeed the state would be a monopoly.

## 5. The Role of Hierarchies

Hierarchies are certainly useful for organizations [Helbing et al., 2006]. A problem might arise when these are too rigid and changes are necessary for adaptation. Moreover, when several aspects should be dealt by a bureaucracy, it might be that one agent should be above another in some aspect (e.g. logistics), whereas in a different aspect the opposite might be the case (e.g. legal advice).

Ashby's law of requisite variety [Ashby, 1956] tells us that a system needs to have enough variety to respond to the variety of its environment (the word variety here could be substituted for the word complexity). A hierarchy could also be necessary for coping with environmental complexity [Aulin, 1979]. Multiscale analysis [Bar-Yam, 2005] is a formal tool that can be used to determine when a hierarchy is required. Basically, if the complexity of an environment cannot be coped by individual agents, these need to aggregate and coordinate to cope collectively. The organizational relations between agents lead naturally to hierarchies.

To visualize hierarchies, bureaucracies can be represented as networks [Strogatz, 2001; Newman, 2003], where each node represents an agent (at a particular scale), and edges represent interactions between agents. Certainly, a network representing a bureaucracy will not be homogenous, since different roles are taken by different agents. A hierarchical bureaucracy can also be represented as a network (Figure 3a). As the complexity (variety) demanded by the bureaucracy's environment increases, the diversity of roles and interactions also augments. A solution would be to increase the number of agents, but this would lead to longer communication delays. A better alternative would be to increase the interaction types between the existing agents, to avoid the introduction of new actors while coping with the required complexity. These new interactions might change the strictly hierarchical nature of the bureaucracy. However, even when the bureaucracy is highly distributed, a certain hierarchy will always be found, simply because the network is not random nor homogeneous, i.e. there will always be agents with more weight in the network's function than others (Figure 3b).

In a system where too many agents need to interact at once, such as the European Union with its current twenty five members, the complexity of the interactions may be too difficult to manage. *Modularity* can help in coping with the complexity [Simon, 1996]. Following the EU example, it will be less complicated if some decisions are made (locally) e.g. by five groups of five countries, and then these five groups discuss a final decision, than having all members discussing at once (globally). This is because the decision delays of each agent add up, since agents (in theory) need to listen to other agents before making a decision. More agents interacting imply more potential friction. In the modular case, discussions can go in parallel, so it would take five decision delays for the first round, and five for the second round, ten in total. In the plain case, it would take twenty five decision delays to have a discussion. Certainly, too many modules would also create delays. A balance should be sought

where agents can make decisions and interact as efficiently as possible. What is important is to note that modularity in a network also implies a certain hierarchy (Figure 3c). The size of modules will also be limited by the cognitive abilities of the agents [Miller, 1956], so in principle each human agent should not have more than seven interactions. Common sense would tell us that other types of agent should also keep the number of interactions low.

A desirable property of bureaucratic networks will be that they have a "small world" topology [Watts and Strogatz, 1998]. This means that most interactions between agents will not need many intermediaries. This is important for information transfer, again, to reduce communication delays. Simply ensuring that messages do not need to pass through several agents before reaching their destination will result in a small world effect, because like this the agents that need to be connected will be connected. If the same message needs to pass from agent A to agent B, to finally reach agent C, it might be worthwhile to simplify and do a shortcut from A to C (Figure 3d). This same idea was posed by [Bollen and Heylighen, 1996] to improve website navigation by dynamically creating direct links between pages that users reached via other pages.

In this scenario, a bureaucratic hierarchy is dynamic and changing when necessary, adapting to changes of its environment directed by friction reduction. Again, these changes can be said to be self-organizing, because the restructuring comes from within the institution, directed by its own dynamics.

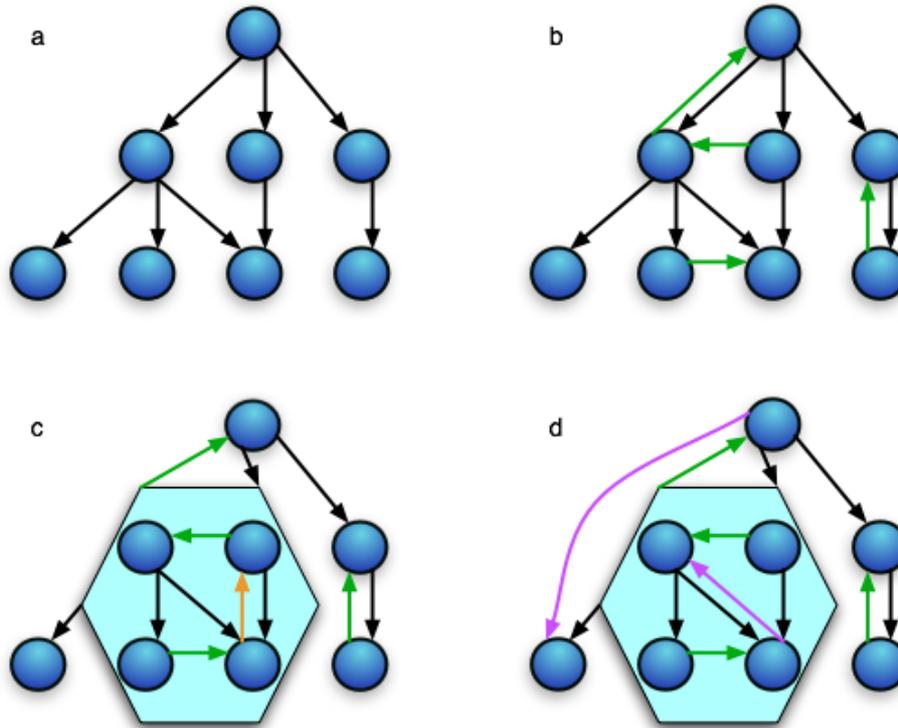


Figure 3. Hierarchy represented as a network. a) Strict hierarchical network. Black arrows show hierarchical dependencies. b) As interactions and dependencies increase, strict hierarchy is broken, but still far from homogeneous. Green arrows show non-hierarchical dependencies. c) Modules can be created when too many interactions cause delays. d) “Shortcuts” can be made to avoid intermediaries, shown in purple. Links can also be removed or modified, as the network adapts, e.g. orange arrow in c).

## 6. The Role of Context

In order to cope with complexity, any organization will try to simplify procedures. Abstracting from several instances, details can be omitted, and a uniform approach can be taken to deal with new instances, internal and external. For example, the public tends to be treated uniformly. This makes sense in cases when there are not many differences between individuals, e.g. to apply for a passport: every citizen has a name, address, etc. However, when a uniform approach is used for cases where there is diversity in the public, difficulties may arise. A single template cannot predict beforehand all cases, and usually makes simple cases complicated. An example can be seen with certain tax declarations, that include very specific sections that only few people must fill in, but are delivered to everybody, even if they just need to sign and declare that they had no income. And changes in the taxation policies can make it complicated enough to pay somebody to fill in a declaration. Rather than including all possible cases in a single form, a more reasonable approach would be to *contextualize* the situation, providing individual solutions for specific cases. Electronic media make this feasible, by generating instant options depending on the current circumstances.

Contextualizing interactions will reduce frictions, internal and external, because both agents involved in the interaction would be benefited if delays are reduced by

removing considerations that do not apply for the current situation. Certainly, too much contextualization can be counterproductive, since agents need to learn how to deal with each case. If every case requires new decisions, then expertise will not be able to improve the performance of agents.

What could be done is to categorize contexts into common occurring categories, by using one of many well known techniques for automatic classification or clustering. Returning to the tax declaration example, people who filled in similar parts of the form can be automatically classified into a contextual category, such as pensioners or unemployed. Like this, a system can find automatically which contexts are common, and what measures should be taken only for those contexts. Since this would be a continuous process, new contextual categories can arise and old ones may disappear. The advantage is that these changes are lead by the usage of the bureaucracy itself, satisfying its demands.

## 7.A Toy Model: Random Agent Networks

To have a better feeling of the usefulness of the ideas described so far, a simple computational model can be used to measure the performance of abstract bureaucracies, represented as “random agent networks” (RANs). This model, partly inspired by random Boolean networks [Kauffman, 1969; Kauffman, 1993; Wuensche, 1998; Aldana-González et al., 2003; Gershenson, 2004b], tries to make as few assumptions as possible about the structures of bureaucracies.

A RAN consists of  $N$  nodes. A node represents an agent, which could represent a person, a department, or a ministry. Each agent  $i$  solves a task. To do so, it sends requests to  $K_i$  other agents, which can be called “dependencies” or connections. The dependencies of every agent are chosen randomly at the beginning of a simulation, not to assume any structure. Once the agent receives a response from all its dependencies, the task is complete. However, the dependencies might receive several requests from several agents. Thus, they store requests in a queue, which they attend in a first-come, first-served basis.<sup>5</sup> Time is also abstracted, so agents take one time step to send requests (transmission delay), one time step to answer one request from the queue (decision delay), and one step to integrate the responses and complete a task (decision delay). Once a task is complete, agents start a new task. Agents respond requests from their queue only when they are expecting responses from their own dependencies.

We can say that the goal of agents is to perform their tasks as fast as possible. The efficiency of the network can be measured by the number of tasks it is able to complete in a given timeframe. Thus, the time “wasted” by agents while waiting for responses from their dependencies (response delay) and having an empty queue should be minimized, since this will be friction.

Now, there are many possible ways of randomly assigning dependencies to agents. The simplest would be **homogeneous**, where each agent has exactly  $K$  dependencies chosen randomly. Following a **normal** probability distribution, every agent will have on average  $K$  dependencies, so some agents will have more and some will have less. A more natural distribution would be **scale-free** [Aldana, 2003], where

---

<sup>5</sup>For simplicity, in the model dependencies do not propagate: requests from queues are answered in one time step, not requiring requests to the dependencies of dependencies. In real bureaucracies, some of these requests might require further requests to further dependencies. Also, requests to other agents could be directed towards persons, groups, or institutions. This would imply different time scales of interaction. Again for simplicity, the time scale is considered to be the same for all agents.

few agents have several dependencies, and most agents have not so many<sup>6</sup>. Intuitively, a special topology where every agent receives the same number of tasks should obtain the best performance, so that workloads are distributed equitably, not allowing request queues to grow for particular agents. A (non random) topology where every agent connects to  $K$  neighbours, similar to cellular automata, fulfils this requirements. This topology can be called **symmetric**. Figure 4 shows examples of the different topologies.

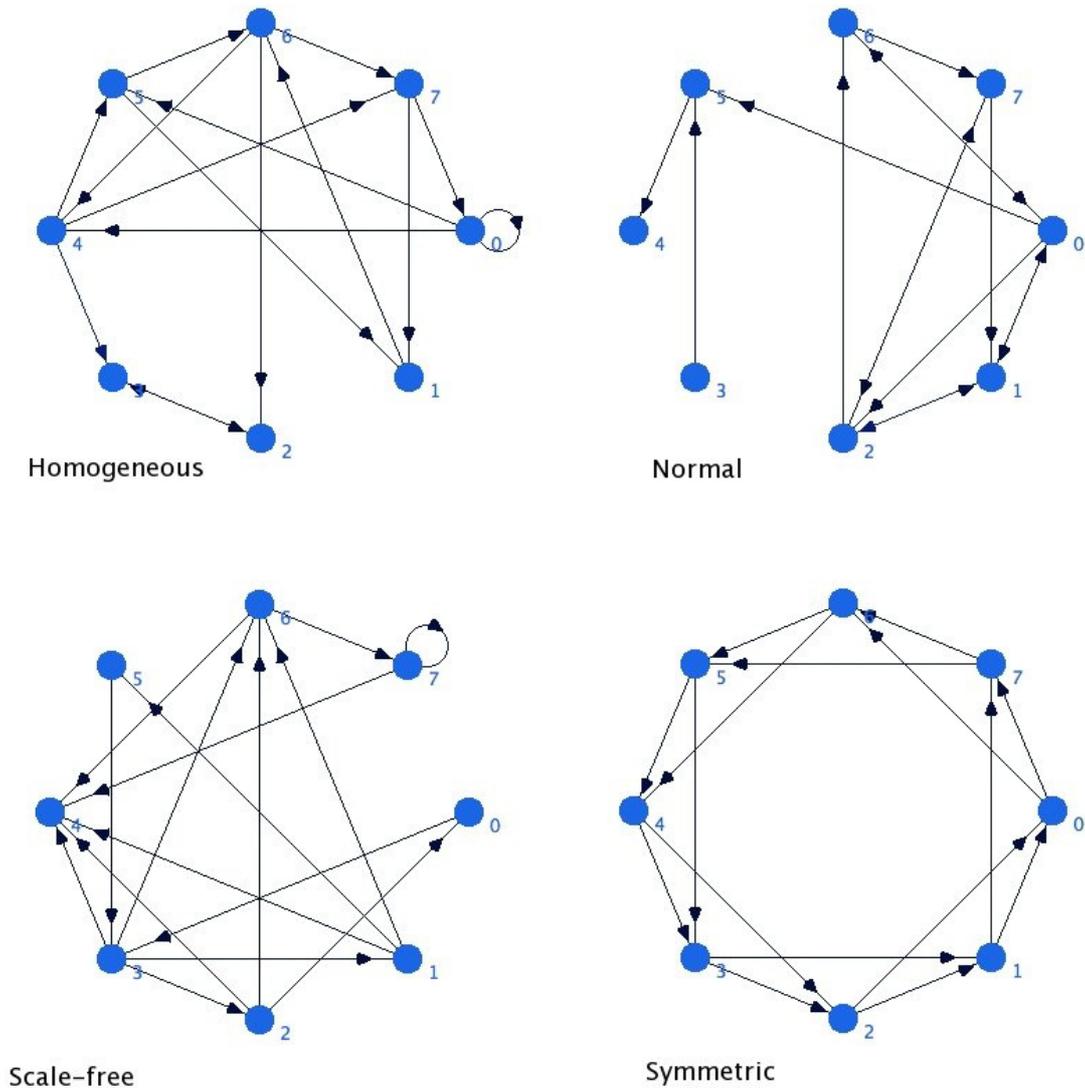


Figure 4. Different RAN topologies for  $N=8$  and  $K=2$ , i.e. eight agents with two dependencies on average. The homogeneous RAN has *exactly* two dependencies chosen randomly. The normal RAN has *on average* two dependencies, some agents having three, some one. The scale-free RAN has most agents with few dependencies, few with many (agents four and six in this case). The symmetric RAN has exactly two dependencies per agent, assigned deterministically, namely each agent

<sup>6</sup>More precisely, the number of dependencies for each agent is generated with the probability distribution  $P(x)=(\gamma-1)x^{-\gamma}$

depends on two neighbours, leading to an equal distribution of dependencies *and* dependents.

The RAN model was implemented in a computer simulation written in the Java programming language. The reader can try the simulation and download the source code via the website <http://rans.sourceforge.net> .

As an initial state, all agents send requests to their dependencies. Afterwards, agents are updated sequentially each time step, i.e. there are  $N$  updates per time step. The behaviour of the network converges to a periodic or quasi-periodic pattern, i.e. an attractor. Interesting parameters to observe are the response delays (how long it takes an agent to complete a task, which is determined by how quickly its dependencies are able to process its requests) and queue lengths (how many requests an agent has yet to process). An example of the dynamics of these parameters can be seen in Figure 5.

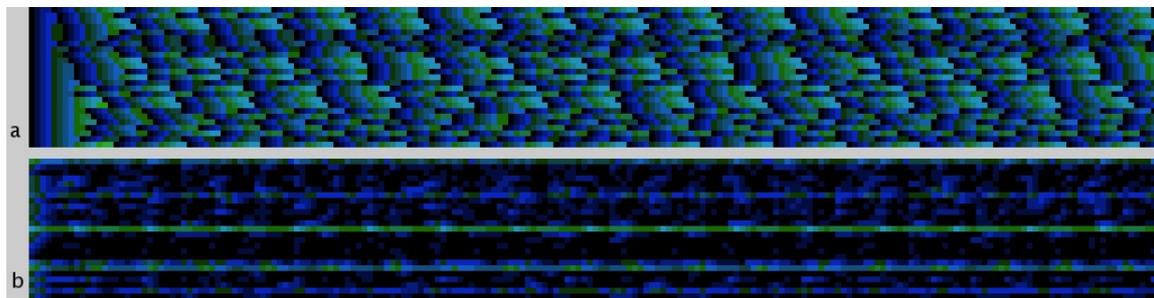


Figure 5. Dynamics for a random agent network of  $N=25$ ,  $K=5$  with homogeneous topology for 200 time steps. a) Response delays. b) Queue lengths. Lighter colours indicate higher values. The initial state is the leftmost column, and the subsequent columns show the temporal evolution of the RAN.

Except for the symmetric topology, for any values of  $N$  and  $K$ , the task queues seem to converge typically to a power law distribution: there are few long ones (bottlenecks), and many short ones.

### 7.1. Using self-organization to improve performance

If we see the satisfaction of agents in terms of the tasks they are able to complete, the satisfaction will be lower when the response delay is higher. Agents with longer queues cause more friction than others, because they cause a high response delay on the agents that are dependent on them, i.e. they have a high previous task delay. Thus, a natural way of reducing friction would be by restructuring the network in such a way to reduce the longest queues.

A very simple criterion achieves this. To restructure a RAN, the agent with the maximum queue average ( $A$ ) is detected. Then, the agent with a maximum response delay ( $B$ ) that has as a dependent the agent with the longest queue changes its dependency to the agent with the shortest queue ( $C$ ). In a real bureaucracy, electronic logs could be used to obtain these measures. This mechanism is illustrated in Figure 6. In many cases, the agent with highest response delay ( $A$ ) has the agent with longest queue ( $B$ ) as its dependency. This is natural, since  $B$  will be able to complete its task only when  $A$  reaches the request after processing its queue. It is obvious that changing

the dependency of  $B$  from  $A$  to  $C$  will reduce the response delay. What is not obvious is the precise effect that this will have on the global performance, i.e. the impact of the change.

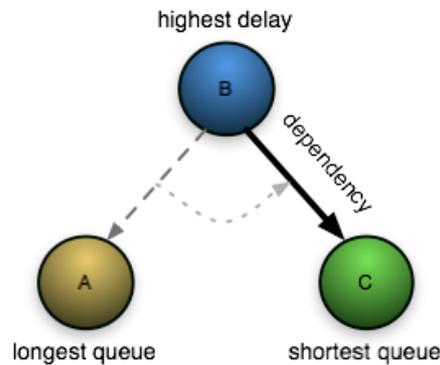


Figure 6. Self-organization mechanism: The agent with highest delay ( $B$ ) restructures its dependency from the agent with longest queue ( $A$ ) to the one with shortest ( $C$ ).

This mechanism can be called self-organizing because the change is local. Certainly, one requires global information to obtain the agents with highest delay, longest, and shortest queues, but this information can be seen as shared in a common environment. There is no central agent dictating the reconfigurations: these are determined by the actual demand, interactions, and dynamics of the RAN. The RAN organizes itself, not by an external force, nor by a central authority.

## 7.2. Simulation Results

To compare the scale-free topology with the others, it was normalized to have a total number of dependencies in the network very close to  $N*K$ .<sup>7</sup> The networks with normal topology also were checked to have a comparable number of dependencies, since networks with less dependencies are able to process more tasks.

Simulation runs were performed for different parameter values. For a network size of  $N=15$ , for each topology, 1000 RANs were created, and their response rate (average tasks completed per time step) was plotted as the self-organization was iteratively applied once each 1000 time steps, i.e. there was a single change in topology every 1000 time steps. For  $K=1$  (Figure 7), the scale free topology performs even better than the symmetric one. This is because there is usually only one or few nodes with lots of dependencies, and many nodes with few or no dependencies. The latter ones are able to complete their tasks quickly, since they have little or no interference from the delays of other nodes, and this enables them to respond quickly to the demands of the former ones. Note that there is already a certain hierarchy and modularity inherent in this configuration. As self-organization restructures the RANs, the non-symmetric topologies increase their performance, and after few reorganizations, the homogeneous and normal topologies also perform better than the symmetric. These two have initially bad performance because randomly some nodes are dependencies of more than one node, while others are dependency of none. This

<sup>7</sup>More precisely,  $\gamma=2.48$  was used, since networks with this value have similar properties to  $K=2$  [Aldana, 2003]. Then, the probability was multiplied by  $K/2$ , to normalize.

creates queues that affect all the nodes that share busy dependencies. By changing the dependency to idle agents, the idle agents are still able to complete their own tasks, while serving as dependencies of other nodes, and reducing the overall friction in the network. Still, self-organization is not able to improve the performance of the symmetric networks, which seem to be trapped in a local optimum.

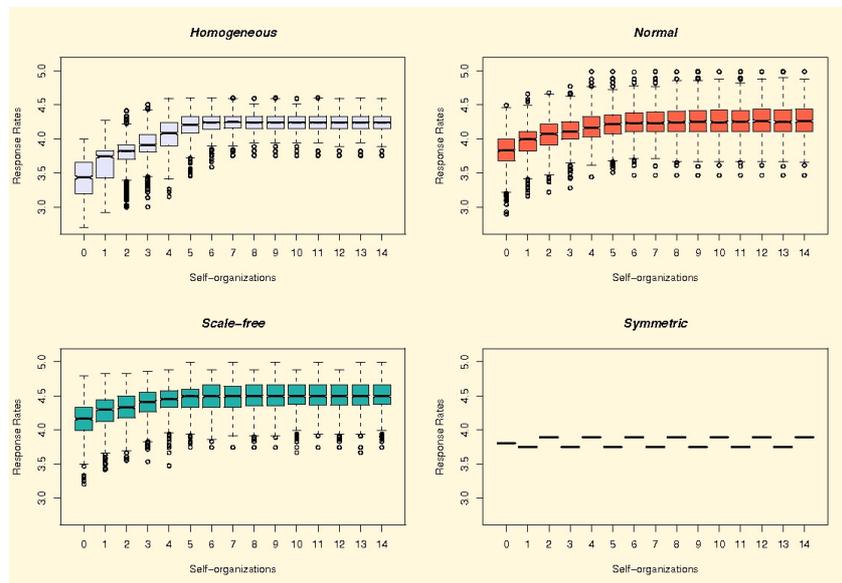


Figure 7. Results for  $N=15$ ,  $K=1$ .

For  $K=2$  (Figure 8), before self-organization, the symmetric topology performs the best. After few self-organizations, the homogeneous topology achieves the same performance, while the normal and scale free surpass it. This is because of the same reason explained above: if some nodes have several dependencies while these have few, overall they are able to process more tasks than if every node has the same number of dependencies: the nodes with few dependencies are able to process their own tasks quickly, and to respond to the agents with many dependencies promptly, creating certain hierarchy and modularity at the same time. A similar case is seen for the case when  $K=5$  (Figure 9).

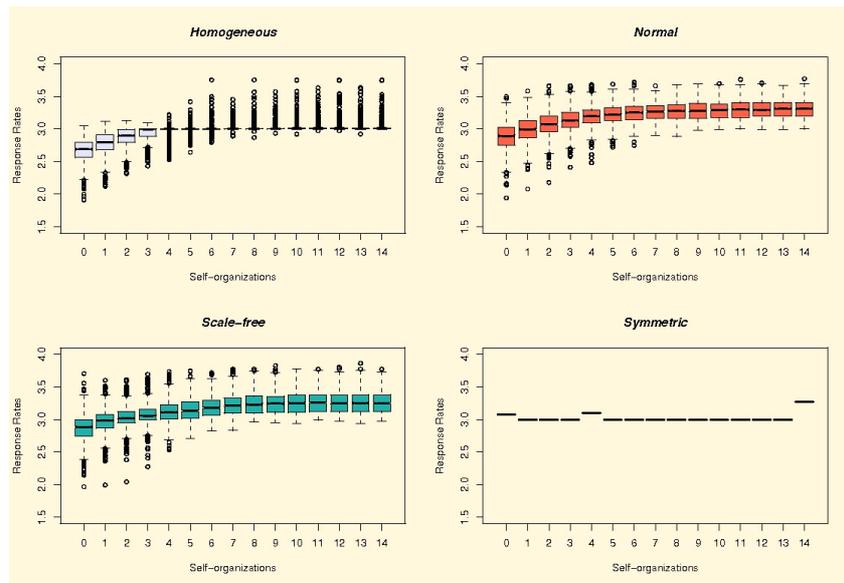


Figure 8. Results for  $N=15$ ,  $K=2$ .

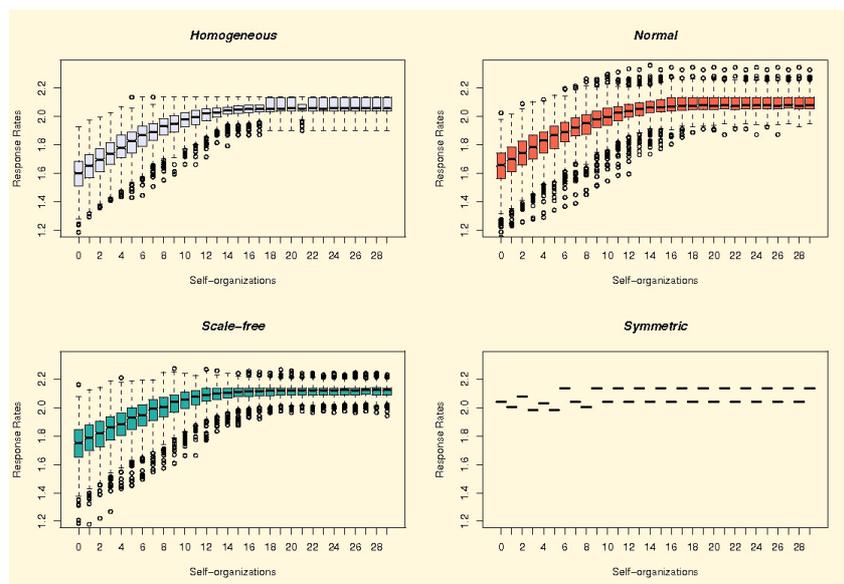


Figure 9. Results for  $N=15$ ,  $K=5$ .

As  $K$  increases, it becomes more difficult to benefit from having several dependencies, since these will have also several dependencies. A high degree of connectivity also implies less hierarchy and less modularity. Thus, initially all topologies perform worse than symmetric, but through self-organization, they tend to reach a similar performance. This can be seen for the extreme case when  $K=15$  (Figure 10).

Scale-free topologies already imply a hierarchy, since few agents have several dependencies. Actually, these "central" nodes dictate the rhythm of the network, simply because they can quickly propagate changes to most of the network via their dependencies, whereas most agents with few dependencies cannot.

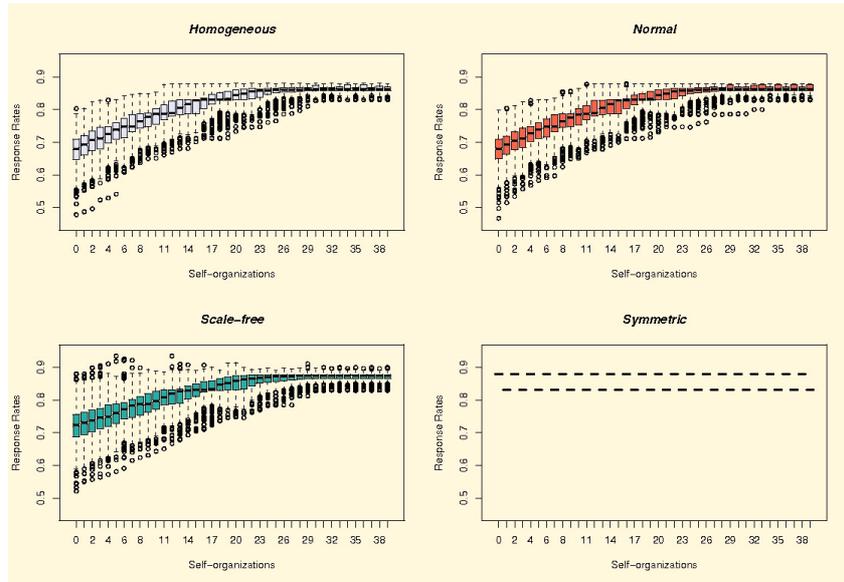


Figure 10. Results for  $N=15$ ,  $K=15$ .

The same behaviour as described above can be seen for larger networks ( $N=100$ ), for  $K=1$  (Figure 11),  $K=2$  (Figure 12),  $K=5$  (Figure 13), and  $K=N$  (Figure 14) (for this last case, only 25 networks were generated).

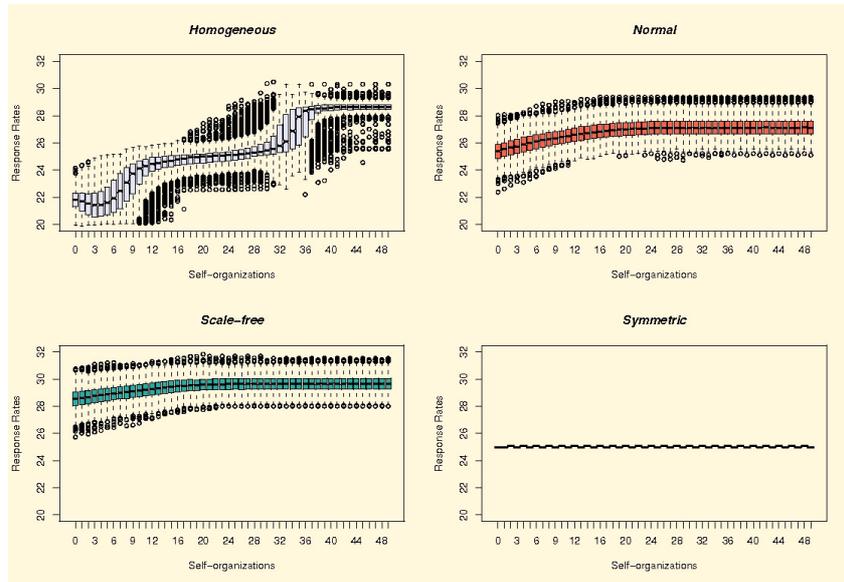


Figure 11. Results for  $N=100$ ,  $K=1$ .

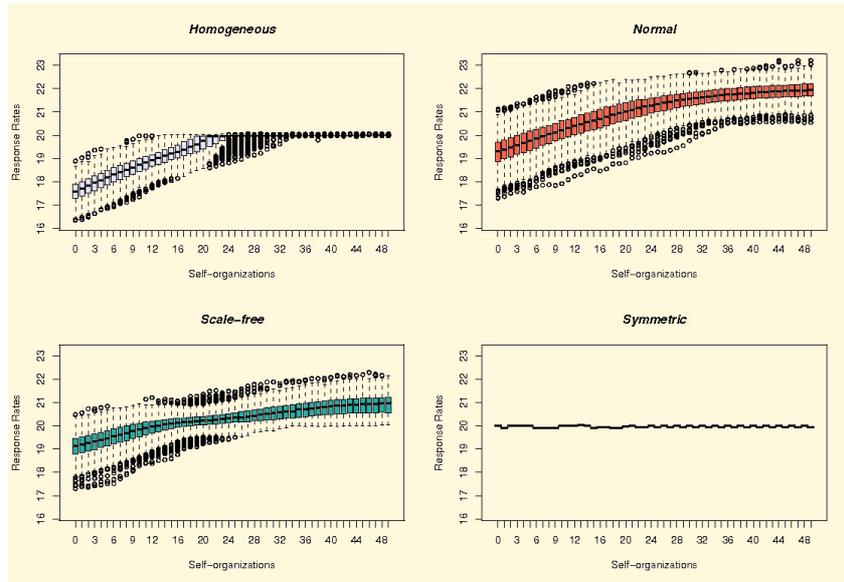


Figure 12. Results for  $N=100$ ,  $K=2$ .

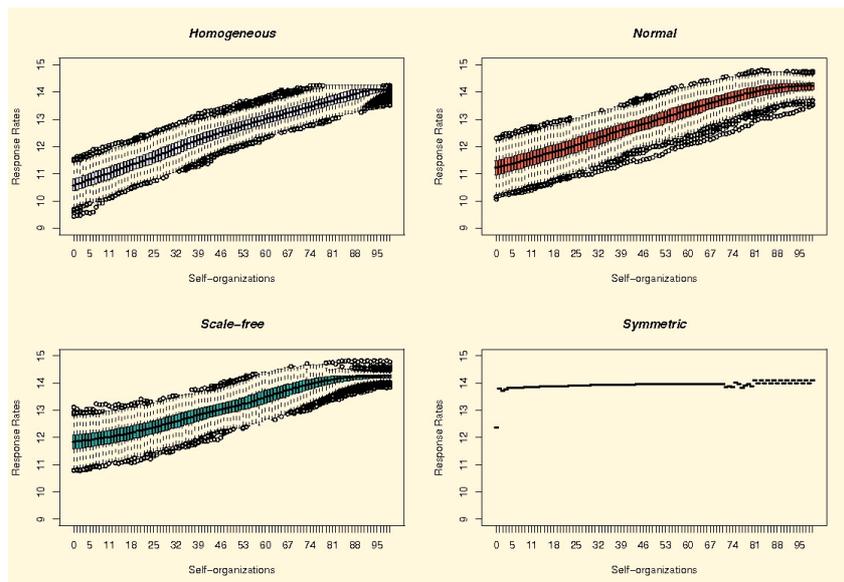


Figure 13. Results for  $N=100$ ,  $K=5$ .

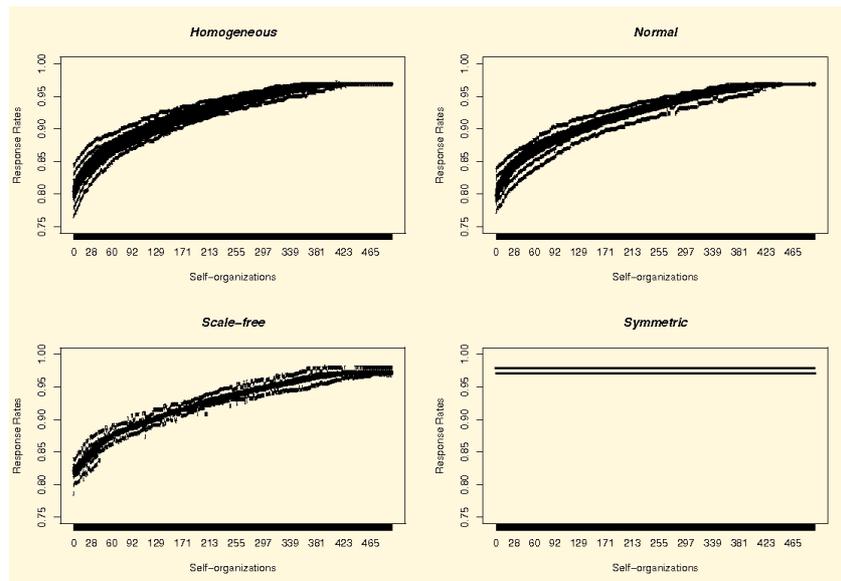


Figure 14. Results for  $N=100$ ,  $K=100$ .

### 7.3. RAN Discussion

Many open questions remain in this model. However, the main goal was to illustrate the benefits of self-organization, namely by reducing delays (friction) in communication. The simulation results showed that only a few modifications of the network topology are required to increase performance to near optimal levels. For higher values of  $K$ , more self-organizations are required, simply because there are more dependencies in the network. Still, an interesting outcome of this model is that only a small fraction of the total number of dependencies needs to be reconfigured to enable a *random* network to achieve good performance. Note, however, that the cost of the changes is not taken into account by the present model. Still, the model showed how few changes can lead to relevant improvements.

A symmetric topology might be desirable, as it gives good results, but it is not realistic. This is because not all agents can cope with the same number of tasks and the requests are not spread uniformly over a bureaucracy. The topology of a real bureaucratic network will be determined more by its demands than by its designers.

Self-organization does not ensure optimality, but adaptability. For example, if the changes in demand of a bureaucracy change the decision time for a task, this can reconfigure itself to accommodate the change robustly. This can be useful for automatic detection of malfunctions and initial response to them: if an agent “breaks down”, its queue would grow, but the agents that have it as a dependency could rearrange their connectivity towards agents working properly. Notice that the presented model assumes that all agents have equal decision and transmission delays. However, weights could be used to model diversity in the delays of agents. Also, if in a real bureaucracy some dependencies cannot be changed, self-organization will find its way with the available flexible dependencies.

One remaining question is: when to stop self-organizing? In principle, self-organization can continue without degrading the performance of the network, while ensuring its adaptability. However, for the simple case presented here, there is no need of adaptation, so sooner or later the self-organization will take the RAN to a previously visited configuration. Some changes actually decrease slightly the performance, but the long run tendency is towards the highest possible performance

for a particular network. If it is known what is the desired maximum performance, then that can be a criterion to stop the self-organization, but it is not obvious to know beforehand the maximum performance.

RAN-like models might also be useful to study organizational robustness, i.e. how well an organization can respond to node failure, using sensitivity analysis. For example, redundancy of nodes can be useful to ensure functionality of key or problematic tasks [Gershenson et al., 2006].

## 8. Conclusions

This paper presented suggestions to use self-organizing techniques to improve the efficiency of different aspects of bureaucracies. All the improvements mentioned decrease different delays within a bureaucracy, reducing frictions and leading to efficient adaptability and robustness. This is because increasing speeds of reaction and decision will allow a bureaucracy to adapt quickly to unexpected changes, while preserving its functionality. In consequence, the "satisfactions" of agents, whether internal or external, will be increased. Moreover, the ideas presented in this paper could be helpful to increase accountability, continuity, and depersonalization of decisions.

It should be noted that even when self-organization can suggest changes to improve efficiency, the human factor also needs to be taken into account, since it is natural to have a resistance to changes. Because of this, any actual implementation of the ideas presented here needs to be developed together with the members of the bureaucracy, to explore which changes are viable and which ones are not.

Standards and digital signatures certainly could be used to comply with the formalities of bureaucracies. Using electronic media is not an impediment for this. The adoption of these media is already underway, and it might bring in more benefits than just decreasing transmission delays. They can effectively support different types of self-organization within bureaucracies. The real value of self-organization will only be appreciated once it is applied in these organizations. But the ideas presented here are encouraging enough to try. Similar approaches could also be useful for other types of organizations: if frictions are reduced, satisfactions will increase.

The model presented, random agent networks, is very abstract indeed. It would be interesting to study more realistic versions of the model, e.g. considering costs of changing the topology, non-homogeneous delays, and domain expertise of agents. Its relationship with fitness landscapes [Levinthal, 1997; Gavetti and Levinthal, 2000] is also worth studying.

## Acknowledgements

I should like to thank Antonio Gershenson, Francis Heylighen, Mixel Kiemen, Diana Mangalagiu, Marko Rodriguez, and anonymous referees for useful discussions and relevant comments. This research was partly supported by the Consejo Nacional de Ciencia y Tecnología (CONACyT) of México and by the Research Foundation - Flanders (FWO).

## References

- Aldana, M. (2003). "Boolean dynamics of networks with scale-free topology", in *Physica D*, vol. 185, issue 1, pp. 45–66.
- Aldana-González, M., Coppersmith, S., and Kadanoff, L. P. (2003). "Boolean dynamics with random couplings", in Kaplan, E., Marsden, J. E., and Sreenivasan, K. R. (Eds.), *Perspectives and*

*Problems in Nonlinear Science. A Celebratory Volume in Honor of Lawrence Sirovich*: Springer Applied Mathematical Sciences Series.

- Anderson, P. (1999). "Complexity theory and organization science", in *Organization Science*, vol. 10, no. 3, pp. 216–232.
- Anderson, P., Meyer, A., Eisenhardt, K., Carley, K., and Pettigrew, A. (1999). "Introduction to the special issue: Applications of complexity theory to organization science", in *Organization Science*, vol. 10, no. 3, pp. 233–236.
- Ashby, W. R. (1956). *An Introduction to Cybernetics*. London: Chapman & Hall.
- Aulin, A. Y. (1979). "The law of requisite hierarchy", in *Kybernetes*, vol. 8, pp. 259–266.
- Axelrod, R. (2005). "Advancing the art of simulation in the social sciences", in Rennard, J.-P. (Ed.), *Handbook of Research on Nature Inspired Computing for Economy and Management*. Idea Group.
- Axelrod, R. and Cohen, M. D. (2000). *Harnessing Complexity: Organizational Implications of a Scientific Frontier*. New York: Free Press.
- Bar-Yam, Y. (2005). "About engineering complex systems: Multiscale analysis and evolutionary engineering", in Brueckner, S., Serugendo-Di Marzo, G., Karageorgos, A., and Nagpal, R. (Eds.), *Engineering Self Organising Systems: Methodologies and Applications*, volume 3464 of *Lecture Notes in Artificial Intelligence* (pp. 16–31). Springer.
- Beer, S. (1966). *Decision and Control: The Meaning of Operational Research and Management Cybernetics*. New York: John Wiley and Sons.
- Bollen, J. and Heylighen, F. (1996). "Algorithms for the self-organisation of distributed, multi-user networks. possible application to the future world wide web", in Trappl, R. (Ed.), *Cybernetics and Systems '96*, volume 911-916: Austrian Society for Cybernetics.
- Bowman, E. H. (1963). "Consistency and optimality in managerial decision making", in *Management Science*, vol. 9, no. 2, pp. 310–321.
- Carley, K. and Prietula, M. J. (Eds.) (1994). *Computational Organizational Theory*. Hillsdale, NJ: Lawrence Erlbaum.
- Carley, K. M. (1997). "Organizational adaptation", in *Annals of Operations Research*, vol. 75, pp. 25–47.
- Comfort, L. K. (1994). "Self-organization in complex systems", in *Journal of Public Administration Research and Theory: J-PART*, vol. 4, no. 3, pp. 393–410.
- Cyert, R. M. and March, J. G. (1992). *A Behavioral Theory of the Firm*. Blackwell Publishing, 2nd edition.
- Dawes, R. M. and Corrigan, B. (1974). "Linear models in decision making", in *Psychological Bulletin*, vol. 81, no. 2, pp. 95–106.
- DeCanio, S. J. and Watkins, W. E. (1998). "Information processing and organizational structure", in *Journal of Economic Behaviour and Organization*, vol. 36, issue 3, pp. 275–294.
- DeSanctis, G. and Monge, P. (1999). "Introduction to the special issue: Communication processes for virtual organizations", in *Organization Science*, vol. 10, no. 6, pp. 693–703.
- Epstein, J. M. and Axtell, R. L. (1996). *Growing Artificial Societies: Social Science from the Bottom Up*. Brookings Institution Press MIT Press.
- Gavetti, G. & Levinthal, D. (2000). "Looking forward and looking backward: Cognitive and experiential search", in *Administrative Science Quarterly*, vol. 45, no. 1, pp. 113–137.
- Gershenson, C. (2004a). "Cognitive paradigms: Which one is the best?", in *Cognitive Systems Research*, vol. 5, no. 2, pp. 135–156.
- Gershenson, C. (2004b). "Introduction to random Boolean networks", in Bedau, M., Husbands, P., Hutton, T., Kumar, S., and Suzukim H. (Eds.), *Workshop and Tutorial Proceedings, Ninth International Conference on the Simulation and Synthesis of Living Systems (ALife IX)* (pp. 160–173). Boston, MA.

- Gershenson, C. (2005). "Self-organizing traffic lights", in *Complex Systems*, vol. 16, no. 1, pp. 29–53.
- Gershenson, C. (2007). *Design and Control of Self-organizing Systems*. Mexico: CopIt Arxivs.
- Gershenson, C. and Heylighen, F. (2003). "When can we call a system self-organizing?", in Banzhaf, W., Christaller, T., Dittrich, P. Kim, J. T., and Ziegler, J. (Eds.), *Advances in Artificial Life, 7th European Conference, ECAL 2003 LNAI 2801* (pp. 606–614). Berlin: Springer.
- Gershenson, C., Kauffman, S. A., and Shmulevich, I. (2006). "The role of redundancy in the robustness of random boolean networks", in Rocha, L. M., Yaeger, L. S., Bedau, M. A., Floreano, D., Goldstone, R. L., and Vespignani, A. (Eds.), *Artificial Life X, Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems*. (pp. 35–42).: MIT Press.
- Haken, H. (1981). "Synergetics and the problem of selforganization", in Roth, G. and Schwegler, H. (Eds.), *Self-Organizing Systems: An Interdisciplinary Approach* (pp. 9–13). New York: Campus Verlag.
- Hand, D. J. and Henley, W. E. (1997). "Statistical classification methods in consumer credit scoring: A review", in *Journal of the Royal Statistical Society. Series A*, vol. 160, no. 3, pp. 523–541.
- Helbing, D., Ammoser, H., and Kühnert, C. (2006). "Information flows in hierarchical networks and the capability of organizations to successfully respond to failures, crises, and disasters" in *Physica A*, vol. 363, issue 1, pp. 141–150.
- Helbing, D. and Vicsek, T. (1999). "Optimal self-organization", in *New Journal of Physics*, vol. 1, no. 13, pp. 13.1–13.17.
- Heylighen, F. (2003). "The science of self-organization and adaptivity" in Kiel, L. D. (Ed.), *The Encyclopedia of Life Support Systems*. Oxford: EOLSS Publishers.
- Hofacker, I. and Vetschera, R. (2001). "Algorithmical approaches to business process design", in *Computers & Operations Research*, vol. 28, issue 13, pp. 1253–1275.
- Holland, J. H. (1975). *Adaptation in natural and artificial systems*. The University of Michigan Press.
- Hutchins, E. (1995). *Cognition in the Wild*. MIT Press.
- Jen, E., Ed. (2005). *Robust Design: A Repertoire of Biological, Ecological, and Engineering Case Studies*. Santa Fe Institute Studies on the Sciences of Complexity. Oxford University Press.
- Kauffman, S. A. (1969). "Metabolic stability and epigenesis in randomly constructed genetic nets", in *Journal of Theoretical Biology*, vol. 22, pp. 437–467.
- Kauffman, S. A. (1993). *The Origins of Order*. Oxford University Press.
- Kauffman, S. A. (2000). *Investigations*. Oxford University Press.
- Kohonen, T. (2000). *Self-Organizing Maps*. Springer, 3rd edition.
- Layne, K. and Lee, J. (2001). "Developing fully functional E-government: A four stage model", in *Government Information Quarterly*, vol. 18, pp. 122–136.
- Levinthal, D. A. (1997). "Adaptation on rugged landscapes", in *Management Science*, vol. 43, no. 7, pp. 934–950.
- Levinthal, D. A. and Warglien, M. (1999). "Landscape design: Designing for local action in complex worlds", in *Organization Science*, vol. 10, no. 3, pp 342-357.
- Levitt, B. and March, J. G. (1988). "Organizational learning", in *Annual Review of Sociology*, vol. 14, pp. 319–338.
- Lissack, M. R. (1999). "Complexity: the science, its vocabulary, and its relation to organizations", in *Emergence*, vol. 1, issue 1, pp. 110-127.
- Luck, M., McBurney, P., Shehory, O., Willmott, S., and the Agentlink Community (2005). *Agent Technology: Computing as Interaction. A Roadmap for Agent Based Computing*. University of Southampton.
- March, J. G. (1978). "Bounded rationality, ambiguity, and the engineering of choice", in *Bell Journal of Economics*, vol. 9, no. 2, pp. 587–608.

- March, J. G. (1991). "Exploration and exploitation in organizational learning", in *Organization Science*, vol. 2, no. 1, pp. 71–87.
- Miller, G. A. (1956). "The magical number seven, plus or minus two: Some limits on our capacity for processing information", in *The Psychological Review*, vol. 63, pp. 81–97.
- Miller Medina, E. (2005). *The State Machine: Politics, Ideology, and Computation in Chile, 1964-1973*. PhD thesis, MIT.
- Morgan, G. (1996). *Images of Organization*. SAGE Publications, 2nd edition.
- Newman, M. E. J. (2003). "The structure and function of complex networks", in *SIAM Review*, vol. 45, pp. 167–256.
- Radner, R. (1993). "The organization of decentralized information processing", in *Econometrica*, vol. 61, no. 5, pp. 1109–1146.
- Ricci, A., Omicini, A., Viroli, M., Gardelli, L., and Oliva, E. (2006). "Cognitive stigmergy: A framework based on agents and artifacts", in *Third International Workshop on Environments for Multiagent Systems (E4MAS 06)*.
- Rodriguez, M. A. and Steinbock, D. (2004). "Societal-scale decision making using social networks", in *North American Association for Computational Social and Organizational Science Conference Proceedings*.
- Simon, H. A. (1964). "On the concept of organizational goal", in *Administrative Science Quarterly*, vol. 9, no. 1, pp. 1-22.
- Simon, H. A. (1976). *Administrative behaviour: a study of decision-making processes in administrative organization*. Free Press, 3rd edition.
- Simon, H. A. (1982). *Models of bounded rationality*. MIT Press.
- Simon, H. A. (1996). *The Sciences of the Artificial*. MIT Press, 3rd edition.
- Stahl, G. (2006). *Group Cognition*. MIT Press.
- Strogatz, S. H. (2001). "Exploring complex networks", in *Nature*, vol. 410, pp. 268–276.
- Theraulaz, G. and Bonabeau, E. (1999). "A brief history of stigmergy", in *Artificial Life*, vol. 5, no. 2, pp. 97–116.
- Turban, E. and Aronson, J. (1997). *Decision Support Systems and Intelligent Systems*. Prentice Hall, 5th edition.
- Van Zandt, T. (1998). "Organizations that process information with an endogenous number of agents", in Majumdar, M. (Ed.), *Organizations with Incomplete Information* chapter 7, (pp. 239–305). Cambridge: Cambridge University Press.
- Watts, D. J. and Strogatz, S. H. (1998). "Collective dynamics of 'small-world' networks", in *Nature*, vol. 393, pp. 440–442.
- Weber, M. (1968). *Economy and Society*. Berkeley: University of California Press.
- Weick, K. E. and Roberts, K. H. (1993). "Collective mind in organizations: Heedful interrelating on flight decks", in *Administrative Science Quarterly*, vol. 38, no. 3, pp. 357–381.
- Wuensche, A. (1998). "Discrete dynamical networks and their attractor basins", in Standish, R., Henry, B., Watt, S., Marks, R., Stocker, R., Green, D., Keen, S., and Bossomaier, T. (Eds.), *Complex Systems '98* (pp. 3–21). University of New South Wales, Sydney, Australia.