

University of Groningen

Self-organising processes of task allocation

Zoethout, K.

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version

Publisher's PDF, also known as Version of record

Publication date:

2006

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Zoethout, K. (2006). *Self-organising processes of task allocation: a multi-agent simulation study*. [Thesis fully internal (DIV), University of Groningen]. s.n.

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Chapter 4

Simulating the Emergence of Task Rotation⁶

Abstract

In work groups, task rotation may decrease the negative consequences of boredom and lead to a better task performance. In this chapter we use multi agent simulation to study several organisation types in which task rotation may or may not emerge. By looking at the development of expertise and motivation of the different agents and their performance as a function of self-organisation, boredom, and task rotation frequency, we describe the dynamics of task rotation. The results show that systems in which task rotation emerges perform better than systems in which the agents merely specialise in one skill. Furthermore, we found that under certain circumstances, a task that leads to a high degree of boredom was performed better than a task causing a low level of boredom.

4.1 Introduction

According to the principle of minimal critical specification (Herbst 1974), an organisation should only offer constraints that are necessary to fix critical issues. For the rest, the workers should be free to self-organise their way of performing a task. Self-organisation refers to the process within a system in which a general order is created without the presence of another system dictating this order (e.g. Dalenoort 1989; 1995; Heylighen 1997). An important organising mechanism is the allocation of tasks. Task allocation refers to the way workers split up tasks into subtasks and divide them among each other. This process depends on task components, psychological components and the constraints the organisation has prescribed. These components not only have an impact on the self-organising process of task allocation, but also influence to what extent the task allocation can be maintained or should be changed. A change in task allocation can be considered as task rotation. In the literature the benefits of task rotation have been discussed extensively (for instance Emery & Trist 1960; Van den Beukel 2003). Workers become experienced in all of the skills that are required to perform a task, which creates multi-availability of team members and therefore leads to a redundancy of functions (Morgan 1986; Kuipers 1989; Van den Beukel 2003). This makes a team more flexible to adapt to changes, either within the team (e.g. illness,

⁶ This chapter is published as: Zoethout, K., Jager, W., & Molleman, E., (2006), Simulating the Emergence of Task Rotation, *Journal of Artificial Societies and Social Simulation* vol. 9-1, <http://jasss.soc.surrey.ac.uk/9/1/5.html>

turnover), or within its environment (e.g. changes in product demand). Another benefit of task rotation is that it may prevent the individual worker to become physically or mentally overburdened due to the repetitiveness of the performance of a single operation (Van den Beukel 2003). In addition, workers may start rotating a task as a consequence of boredom.

Because of its benefits, task rotation has been implemented in various settings, often related to self-managing teams (e.g. Atkinson 1984). However, although the outcome of task rotation has been studied intensively, studies on the emergence of task rotation have not been conducted. Emergent properties refer to properties of a system that cannot be reduced to the properties of the elements, i.e. the agents, which the system consists of (e.g. Heylighen 1997). In the case of a popular concept such as task rotation, research into its emergence will certainly contribute to a better understanding of the concept.

In this chapter, we describe three classes of components that could influence the process of task rotation: organisation components, referring to the constraints on the self-organising process, system components, referring to the psychological characteristics of the workers, and task components, referring to the characteristics of the task. On the basis of these components we conducted three series of experiments. First, we studied the performance and the development of the expertise of agents in a system where they were not allowed to rotate tasks. Second, we studied the performance and the expertise of agents who were free to rotate tasks whenever they liked to. Third, we studied task rotation in a setting in which the agents could only rotate at fixed points in time.

4.2 The model

4.2.1 The task

Despite all the different definitions and perceptions of tasks, it is quite clear that a task requires one or more skills. These skills are related to task-actions, i.e., the parts a task consists of. Theories on tasks indicate that a task can be split up into task-actions that can be related to skills in a large number of ways (Hunt 1976; Weick 1979; Tschan & Von Cranach 1996). According to our definition, a task consists of actions in such a way that for every action exactly one skill is required to perform this action. In order to perform the whole task, all of the actions that a task consists of have to be performed one or more times. The number of times the actions have to be performed, we call cycles. In this way we describe a task as a two-dimensional matrix of actions and cycles.

4.2.2 The agents

Important components that determine group performance are expertise and motivation (Wilke & Meertens 1994; see also Steiner 1972). These components can be considered as the characteristics of skills; workers, i.e. task performing humans, perform particular

skills better or worse than other skills (expertise), and prefer using some skills more or less to using other ones (motivation). Workers are formalised as agents with properties necessary to perform the tasks (for an elaborate description, see Zoethout, Jager, and Molleman 2004). These properties are represented as a set of skills, with each skill consisting of two variable components, expertise and motivation, and thresholds. These thresholds determine whether the expertise and motivation are sufficient to perform the task.

Skills can be considered as being either active or passive. Passive skills refer to the total range of skills of the agents. Active skills concern the skills that are actually required to perform a specific task. As soon as agents have to perform a task, a subset of the set of passive skills is activated, which corresponds to the actions the task requires. After the skill has been used and the task has been completed, the skill becomes passive again.

If the agents use the same skill over and over again, their expertise will improve, but their motivation may decrease (Hackman & Oldham 1980; see also Hackman & Morris 1975). This may affect the performance in a number of ways. First, an increase in expertise will lead to an increase in performance. Second, the decrease in motivation may lead to a decrease in performance. Third, the decrease in motivation may lead to re-allocation of the task, i.e. task rotation.

4.2.3 The model

Figure 13a and 13b describe the different components of our model:

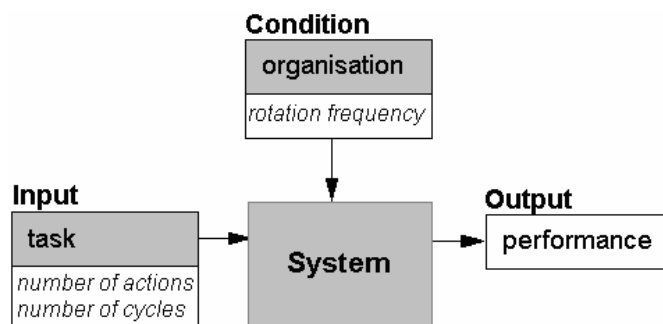


Figure 13 a: The system and its environment

The model is a simple input-process-output model (see Figure 13a). The input is a task. As we stated above, a task consists of a number of actions and a number of cycles. Some organisations offer a higher degree of freedom in the self-organising allocation process than others. For instance, in some organisations the agents are allowed to rotate tasks whenever they want, in others only once a day, or not at all. In the model this is represented as rotation frequency, i.e. the number of possible rotations of one task. The output of the system is the task that has been performed. The performance of the system, which is the dependent variables, indicates how well the agents have performed the task. The performance is based on expertise and motivation, the two

mediating variables, and the way in which the task has been allocated. It is expressed as the total time.

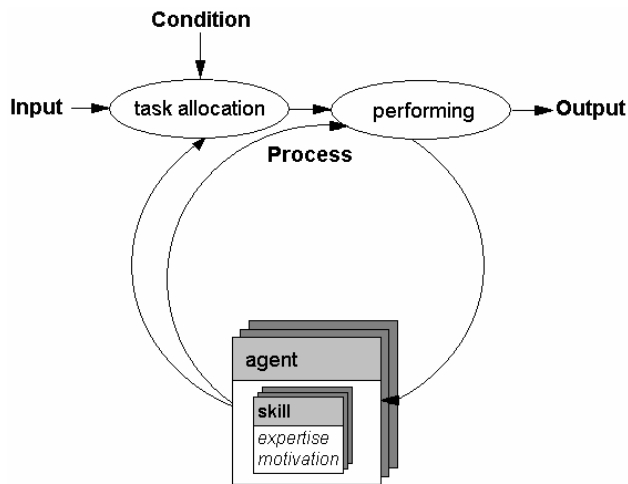


Figure 13b: The System

The system that has to perform the task consists of agents (see Figure 13b). Every agent has skills that correspond to the action that the task consists of. Each skill has two components, expertise (how good?) and motivation (how nice?). On the basis of these components, the agents allocate the task and subsequently perform it. The performance process influences the expertise of the agents: the agents learn by carrying out the task and forget by not carrying out the task. Motivation is subject to fluctuation because the agents can become bored and can again recover from that. When there is a change in an agent's expertise and/or motivation, task rotation may emerge. Whether or not the process of task allocation is self-organising depends on the specifications of the organisation (see also Herbst 1974). Input, condition, and output refer to the concepts as proposed in Figure 13a.

4.2.4 The allocation process

On the basis of the expertise and motivation components of the active skills, each agent determines individually what actions he does or does not want to perform by applying the following rule:

IF Expertise > ExpertiseThreshold AND Motivation > MotivationThreshold THEN I DO ELSE YOU DO

With every action the agents choose whether to perform it (I do) or leave it to others (You do). This choice is represented as two nodes, I and YOU, whereby there is a tension based on the values of the expertise and motivation. The choice forms the basis of influencing other agents. The influence process starts from the assumption that every agent persists in his choice and will try to influence other agents to make a choice complementary to it (see also Kerckhoff & Davis 1962). For example, five agents have

to sail a ship. Sailing a ship is a task consisting of five actions: handle the helm, the main sail, the jib, the swords, and watch the area. Based on his expertise and motivation, agent one only wants to handle the helm. Therefore, he tries to influence the other ship members not to handle the helm, but to handle the main sail, the jib, etc. instead. Simultaneously, the other agents will try to influence others in the same way. We represent the influence process by using excitatory, i.e. tension increasing, and inhibitory, i.e. tension decreasing connections (see Figure 14):

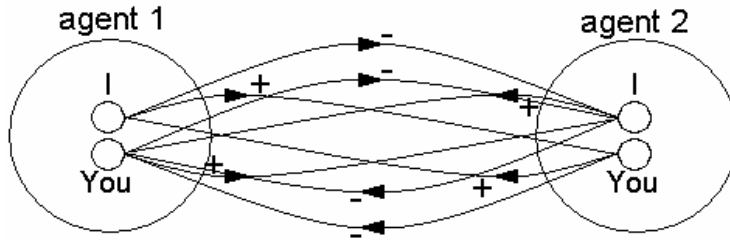


Figure 14: Excitation (+) and inhibition (-) of two agents

With every action, the I-node of one agent inhibits the I-node of the other(s) and excites the You-node of the other(s). The You-node of one agent inhibits the You-node of the other(s) and excites the I-node of the other(s). If the I-node of an agent is bigger than the You-node, the choice of the agent is ‘I Do’, otherwise it is ‘You-do’.

The excitation and inhibition functions are based on the values of I and You of both agents and the differences of their nodes. The differences of the nodes are included into the functions in order to prevent possible oscillations when the nodes have about the same value (see also Zoethout, Jager, and Molleman, 2006b):

$$I_2 := I_2 - \alpha I_1 I_2 \text{Diff}_{II} \quad I_1 \text{ inhibits } I_2 \quad (11a)$$

$$You_2 := You_2 - \alpha You_1 You_2 \text{Diff}_{YouYou} \quad You_1 \text{ inhibits } You_2 \quad (11b)$$

$$You_2 := You_2 + \epsilon (1 - You_2) I_1 \text{Diff}_{II} \quad I_1 \text{ excites } You_2 \quad (11c)$$

$$I_2 := I_2 + \epsilon You_1 (1 - I_2) \text{Diff}_{YouYou} \quad You_1 \text{ excites } I_2 \quad (11d)$$

I_1 , I_2 , You_1 , and You_2 correspond to Figure 14. Diff_{II} equals the absolute value of $(I_1 - I_2)$ and Diff_{YouYou} equals the absolute value of $(You_1 - You_2)$.

This process ends as soon as the agents reach a complementary situation. This means that after the allocation process, with respect to a particular action, there is only one agent who finally chooses ‘I Do’. All the others choose ‘You Do’. Only when this complementary situation has been reached with all actions, the agents start performing the task.

4.2.5 Performance, learning and boredom

We consider the performance of the system as a function of expertise and motivation (see also Wilke & Meertens 1994). Both expertise and motivation are defined in terms of the time it takes to perform a task: the higher the degree of expertise or motivation, the sooner the task will be finished. Furthermore, we define a minimal time to complete an action, t_{action_i} , which is equal to the actual time it takes to perform the action at a maximal rate of expertise and motivation. The actual performance time of a single agent, $t_{perf.}$, can therefore be expressed as:

$$t_{perf.} = \sum_{i=1}^n \frac{t_{action_i}}{\lambda \frac{e_i}{e_{max}} + (1-\lambda) \frac{m_i}{m_{max}}} \quad (12a)$$

λ represents a parameter that determines the balance between expertise and motivation (0.5 in all of our experiments). In the present study, the agents perform the actions simultaneously. This means that the time it takes to perform the total task, $t_{perf.}$, is determined by the slowest agent.

During the performance process, the expertise and the motivation of the agents change. When agents perform a task, their expertise increases, and when they do not, it decreases. According to Nembhart (2000), forgetting can be described as ‘following the way back’ on a learning curve. Therefore we use the inverse function of learning to describe forgetting. Furthermore, an important characteristic of most learning curves is that they reach a maximum asymptotically (Nembhart 2000). Therefore, we define learning by means of the relations among expertise (e) at a certain time (t), expertise in the future ($t+1$), the maximum expertise (e_{max}), and a parameter (λ , [0,1]) that determines the learning speed:

$$e_{(t+1)} = e_t + \lambda \frac{e_{max} - e_t}{e_{max}} \quad (13a)$$

Forgetting is the inverse of learning, therefore:

$$e_{(t+1)} = \frac{(e_t - \mu)e_{max}}{e_{max} - \mu} \quad (13b)$$

whereby μ [0,1] determines the forget speed.

Motivation curves can be described by applying the same characteristics: a maximum that is reached asymptotically, and recovery as the inverse of boredom. This means that formula (13b) describes the motivational decrease related to boredom and formula (13a) represents the motivational increase related to the recovery from boredom. In this case the parameters μ and ν respectively describe the recovery and the boredom speed.

By using the parameters λ and μ we can set the degree of boredom/recovery and learning/forgetting of different tasks. For instance, repetitive tasks, i.e. tasks with a large number of cycles, may lead to an increase in expertise and a decrease in motivation. As the failure of the Taylor-approach indicates, this may finally lead to a decrease in performance (see also Wilke & Meertens 1994). On the other hand, not using skills for a period of time may lead to a decrease in expertise. This suggests that there must be a balance between expertise and motivation. Task rotation may help to reach this balance.

4.3 Experiments

All experiments were conducted with WORKMATE II, a computer simulation program that we developed in DELPHI6⁷. We studied a task performing system using three types of organisations. In accordance with the first type, self-organisation, the organisation does not constrain the allocation process. This means that the agents are free to allocate tasks by themselves and change the task allocation whenever they want to. Here, the allocation process takes place in the way we described it in the former section. The second type we call semi-self-organisation. This type differs from the former type in three ways: first, the agents are still free to allocate the tasks by themselves, i.e. the task is not yet assigned by the organisation. However, agents are only allowed to re-allocate the task, i.e. rotate tasks, after a fixed number of cycles. Thus, during this fixed number of cycles, the agents have to maintain the same task allocation. Second, the agents are only allowed to use one skill per cycle. Third, the allocation process is not based on interaction between agents, but on the following rule:

$$\text{IF } (I_{agent_i, skill_k} - You_{agent_i, skill_k}) > (I_{agent_j, skill_k} - You_{agent_j, skill_k}) \text{ THEN } agent_i \text{ performs } skill_k$$

This rule implies that as regards a certain skill, the agents with the highest difference of the I- and You-node will perform that skill. For example: agent 1 has an I-node of 0.5 and a You-node of 0.2, and agent 2 has an I-node of 0.5 and a You-node of 0.4. The difference of the nodes of agent 1 is $0.5 - 0.2 = 0.3$. To agent 2 applies: $0.5 - 0.4 = 0.1$. Therefore, agent 1 will perform the action. The use of this rule will make no difference in the actual allocation. The only difference is that here the outcome, i.e. the final allocation, is calculated on the basis of the initial choice of the agents rather than on the

⁷ The program can be downloaded at <http://jasss.soc.surrey.ac.uk/9/1/5.html>

interaction among the agents⁸. In the third type of no self-organisation, the agents cannot rotate tasks at all, but have to keep performing the same action that they initially started with. The agents choose that particular action in the same way as they do in case of the former type.

For each of these organisation types we studied tasks with degrees of demotivation/recovery that were high, moderate and low. Regarding the semi-self-organisation types, we studied three settings of rotation frequency. Rotation frequency is defined as the inverse of the number of cycles in which the agents are forced to use the same skills consecutively. As regards the other types, the rotation frequency has not been manipulated: the ‘no self-organisation’ type does not involve task rotation and in case of self-organisation, the agents decide for themselves whether or not to rotate. Table 4 summarises the experimental design by showing the independent variables.

Table 4: Experimental design

Organisation type	Boredom/Recovery	Rotation Frequency
No Self-organisation	High (100/100)	NONE
	Moderate (50/50)	
	Low (29/29)	
Semi Self-organisation	High	1/5, 1/25, 1/50, 1/75
	Moderate	1/5, 1/25, 1/50, 1/75
	Low	1/5, 1/25, 1/50, 1/75
Self-Organisation	High	NOT MANIPULATED
	Moderate	
	Low	

A rotation frequency of 1/5 implies that the agents are allowed to rotate every 5 cycles. Expertise, motivation and performance are considered as the dependent variables.

We conducted the experiments by using the following parameter values:

A system consisted of 5 agents, each with 5 skills

The agents had to perform one task of 200 cycles and 5 actions

The maxima of both motivation and expertise were set at 25

The motivation –and expertise thresholds were set at 10

⁸ We maintained the interaction model to be able to conduct experiments by using the variable *interaction time*, which we did not use in the study we described in this paper.

The learning speed was 100, the forget speed 10

After conducting numerous trials, we found out that these values appropriately marked the parameter space in which the processes occurred that we wanted to study (see also Zoethout et al, 2006b).

In all experiments, we used the following initial skill values of the agents (see Table 5).

Table 5: Initial setting of the agents in all experiments

	Agent 1	Agent 2	Agent 3	Agent 4	Agent 5
Skill 1	11	12	13	14	15
Skill 2	12	13	14	15	11
Skill 3	13	14	15	11	12
Skill 4	14	15	11	12	13
Skill 5	15	11	12	13	14

We see that all the agents have the same pattern of values, but they are assigned to different skills. Initially, to each skill applied that the expertise value was the same as the motivation value. Therefore, in Figure 15, with each skill we mention only one value, representing both the initial expertise and the initial motivation. Although we start from these values, under the influence of expertise and motivational processes, the agents may decide to start rotating their tasks.

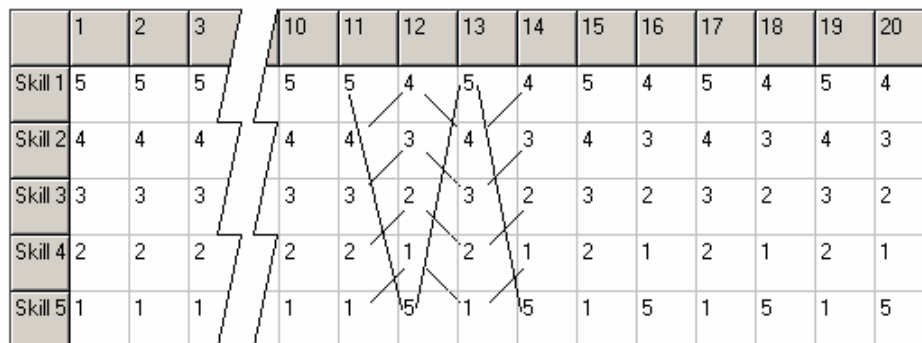


Figure 15: An example of task allocation

Figure 15 depicts an example of task allocation. The x-axis represents the cycles and the y-axis shows the skills. The numbers in the table refer to the agents using a particular skill. We see that the agents start in accordance with the initial values as described in table 5, i.e. agent 1 starts with skill 5, agent 2 starts with skill 4, etc. At the 12th cycle, the agents rotate for the first time, i.e. agent 1 rotates from skill 5 to skill 4, agent 2 rotates from skill 4 to skill 3, etc. This is represented by means of the small

lines. During the next cycle, the agents rotate back, etc. This example depicts the actual task allocation as it occurs by using certain parameter values. Since we will give an elaborate description of expertise, motivation, and performance in the next section, in the actual allocation tables no information has been added. Therefore, we did not include them in the result section.

4.4 Results

In this section we have chosen not to depict all the results in detail, but to focus on the most interesting phenomena instead. Therefore, we will only present the results that show the most important aspects of the allocation processes and their outcomes.

4.4.1 Organisation type

In this section we will discuss the influence of the three organisation types, self-organisation, semi self-organisation, and no self-organisation, in a setting with a high degree of boredom/recovery. For the semi self-organisation type we used a rotation frequency of 1/5. We will discuss the expertise, motivation, and performance.

4.4.1.1 The influence of self-organisation on expertise development

In accordance with the self-organisation type, the agents rotate between two skills (see Figure 16a).

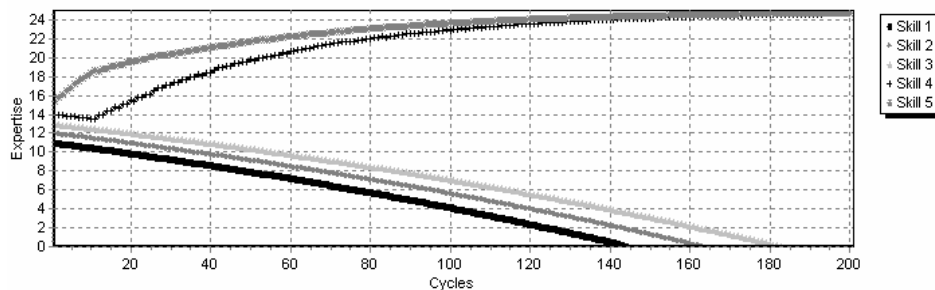


Figure 16a: Expertise in case of self-organisation

Figure 16a shows the development of expertise with respect to the five skills of agent 1. The x-axis represents the cycles and the y-axis shows the expertise. Since the agents are similar in the sense that they will all become specialised in the two skills with the highest initial value, and will forget the other three, we only show the results of agent 1. We see that the expertise in the second best skill of the agents increases after they have started rotating their tasks at the 12th cycle. The first time the agents rotate is determined by boredom. After that, they rotate after every cycle. As a result, the agents do not have time to forget their skills, but are able to increase their expertise. The semi self-organisation type shows about the same development of the expertise. No self-

organisation leads to a situation in which the agents specialise in one particular skill while forgetting the other skills.

4.4.1.2 The influence of self-organisation on motivation

In all three organisation types, initially, the motivation of the agent for using the best skill, i.e. the skill that is initially used, decreases while their motivation for the other skills remains the same. In case of no self-organisation, the motivation of the best skill simply decreases to zero. In case of self-organisation, after the first rotation at the 12th cycle, the agents start using their second best skill. As a result, the motivation for the best skill increases, whereas the motivation for the second best skill decreases until the next rotation. After that, the whole process is repeated. In the long run, we see that the motivation for both skills stabilises at 12. In case of semi-self-organisation, we observe about the same periodic changes, although here the frequency is 10 cycles instead of 2, as is the case with the self-organisation type (see Figure 16b):

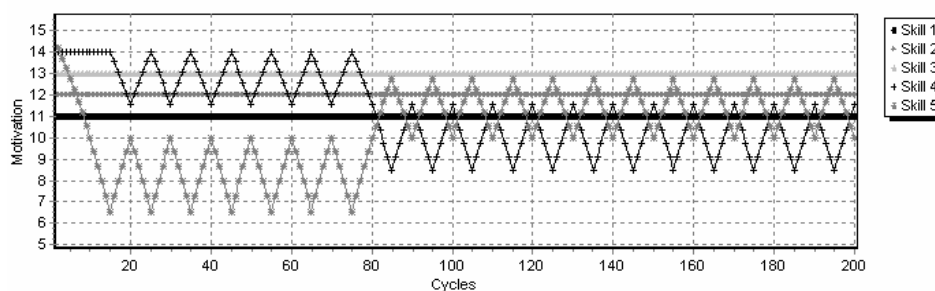


Figure 16b: Motivation in case of semi self-organisation

Figure 16b shows the motivation for all skills. The flat lines represent the motivation level of the skills that are not used. The motivation for the skill that was initially performed the best, skill 5, initially decreases and then starts fluctuating. At the 80th cycle, the agents did not rotate, although they had the opportunity. This led to a further recovery of the best skill, and a higher degree of boredom with respect to the second best. The reason for this anomaly can be found in the expertise development. The expertise levels regarding the skills used to perform the task are inclined to increase in a similar way, because both skills will eventually reach the same maximum (see also figure 16a). This means that as the task proceeds, expertise will have less influence on the decision what to do than motivation has. This implies that the agents will only rotate the task after the lowest and the highest motivation rates have been swapped. From the 85th cycle onward, we see that this is precisely what happens.

4.4.1.3 The influence of self-organisation on performance

In the case of self-organisation, the agents rotate between two skills. Because of the initial increase in the boredom, the performance time initially increases until the first task rotation. After that, the agents rotate after every cycle. Therefore, the agents do not forget or get bored. However, they do learn. This implies that the performance time is

solely determined by the increase in expertise. Finally, when the expertise reaches its maximum, the performance time ends in a flat line at the value of 14.

With respect to the performance time the semi self-organisation type shows about the same development (see Figure 16c).

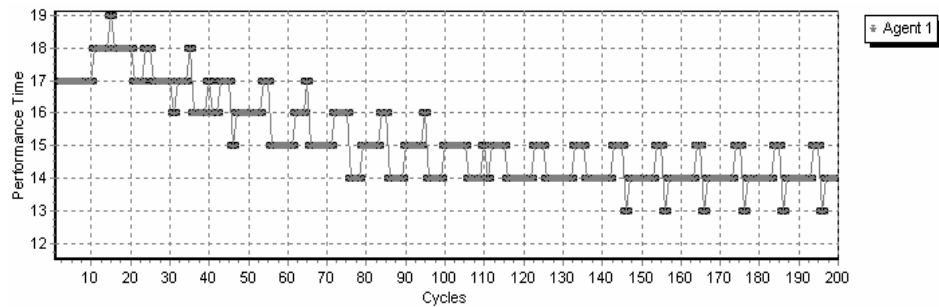


Figure 16c: Performance time in case of semi-self-organisation

Because the agents work synchronously and start with the same initial values, their performance time is identical. This is why we only show the performance of one agent. Here the rotation frequency of 1/5 causes some small fluctuations because during the time period of 5 cycles there is some time to learn, to forget, to get bored and to recover from it. Apart from these fluctuations we see the same development as with the former type, including the final performance time of 14.

The type of no self-organisation shows about the same development. First, we see an increase in performance time due to boredom. However, this increase does not stop at the 12th cycle as in the former types, but continues until the agents have lost all of their motivation at the 22nd cycle. After that, the increase in expertise causes a small decrease in the performance time until it ends in a flat line at value 19. Thus, without self-organisation the performance is a great deal worse.

4.4.2 Boredom

In this section we will discuss the influence of high, moderate and low levels of boredom in an organisation type with self-organisation. We will discuss the impact on expertise, motivation, and performance respectively.

4.4.2.1 The influence of boredom on expertise development

The first setting in which there is a high degree of boredom is identical to the first organisation type we discussed in the former section: the agents rotate between two skills from the 12th cycle (see also Figure 16a). From that point on, both of these skills improve, whereas the other skills are forgotten. In case of a moderate boredom rate, the task rotation starts at the 34th cycle. This implies that the expertise of the skill that the agents use after the task rotation has already decreased somewhat before rotation. This

effect is even more clearly visible in the setting in which there is a low level of boredom (see Figure 17a).

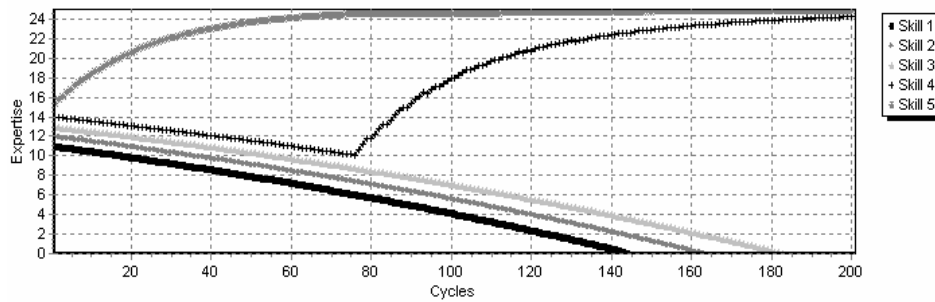


Figure 17a: Expertise development in a situation with a low degree of boredom

Here we see that, initially, the agents seem to specialise in one skill. However, at the 77th cycle, they develop expertise in the second skill. This is ‘right on time’, because the expertise in the second best skill has almost dropped below its threshold of 10. This explains why a boredom-recovery rate < 29 will not result in a task rotation process.

4.4.2.2 The influence of boredom on motivation

It is evident that at a high boredom rate the motivation decreases quicker than it does at a low rate. Figure 17b shows that the curve of the best skill decreases much slower than in the situation of a high boredom level.

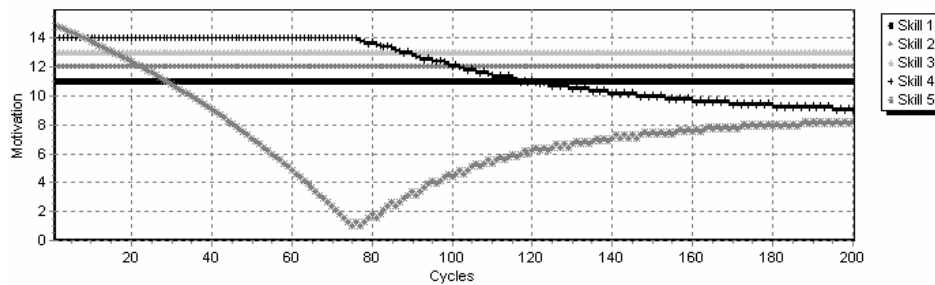


Figure 17b: Motivation in the situation with a low boredom rate

After the first rotation at the 77th cycle, we see that the motivation for the skill that was initially performed the best increases again, whereas the motivation of the second best skill decreases.

4.4.2.3 The influence of boredom on performance

As we discussed in 4.4.1.2, in the first setting of high boredom, the performance time first increased because of the boredom, and subsequently decreased from the first

rotation at the 12th cycle to a flat line at a value of 14. In the second setting of moderate boredom we found that the initial increase in the performance time was less rapid. This implies that the performance time decreases later, i.e. from the first rotation at the 34th cycle, and also ends at the value of 14. The last setting of low boredom starts with a decrease in the performance time (see Figure 17c).

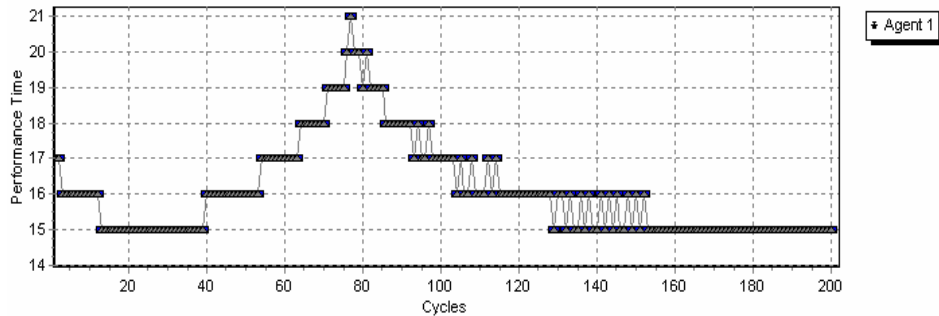


Figure 17c: Performance time in a situation with a low boredom rate

The initial decrease is caused by the increase in expertise, which has a stronger effect than the decrease in motivation. After 30 cycles the performance time increases until the first rotation at the 77th cycle, and then decreases to a value of 15. The value of 15 instead of 14 is caused by the recovery time: In the experiments the values of boredom and recovery were set at the same value. This implies that the setting of low boredom includes low recovery. Apparently the recovery process takes place slower than the learning process. With respect to the allocation process this means that, if the recovery rate decreases, the role of the expertise of the agents becomes more important. Figure 17c depicts only one agent because the performance time of all the agents is identical, apart from one exception which the figure does not depict: from the 78th cycle, the agents sometimes perform 2 actions in one cycle and none in the other. This indicates that the performance time either suddenly increases to twice as high or decreases to zero. We consider this phenomenon as an anomaly that does not contribute to an understanding of the process.

4.4.3 Rotation frequency

In this section we will discuss the influence of four different rotation frequencies on the semi-self-organisation type. Rotation frequency only occurs with semi self-organisation: In case of self-organisation, the rotation frequency is determined by the agents themselves, and no self-organisation does not involve task rotation. We studied the rotation frequency in a high boredom setting.

4.4.3.1 The influence of rotation frequency on expertise development

In the first situation with a rotation frequency of one possible rotation occurring every 5th cycle, the agents develop two skills while forgetting the other skills. If we lower the frequency to once every 25 cycles, the effect remains the same, although the second skill is learned later (see Figure 18a).

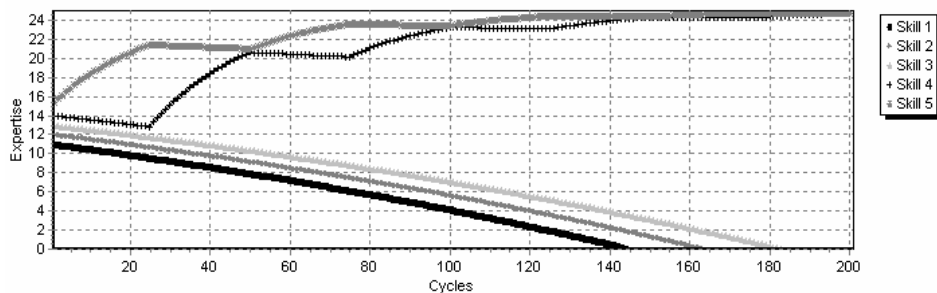


Figure 18a: Expertise development at a rotation frequency of 25

This also applies when we lower the frequency further. Even at a frequency of once every 75 cycles, the agents are still capable of learning two skills (see Figure 18b).

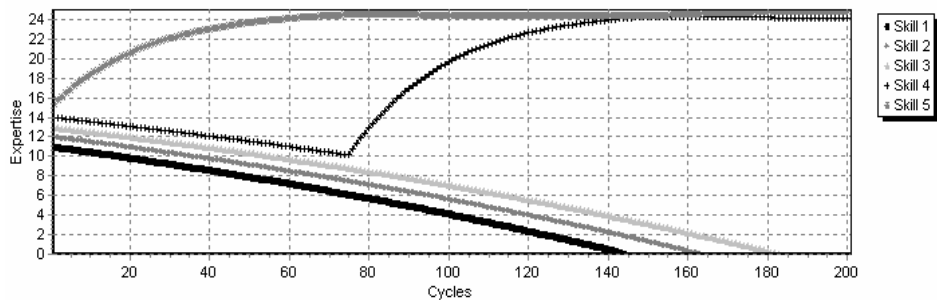


Figure 18b: Expertise development at a rotation frequency of 75

Figure 18b shows a process comparable to the one depicted in Figure 17a. The decrease in the rotation frequency results in the same effect as the decrease in boredom/recovery does: because the first moment of task rotation occurs quite late, either because of a low boredom rate or a low rotation frequency, the agents have more time to specialise in one skill and forget the second. This implies that if we would lower the rotation frequency even further, the agents forget their second best skill at a level below the threshold of 10, and as a result end up specialising in only one skill. As Figure 17a indicates, a setting with a rotation frequency lower than 1/77 would actually lead to this outcome.

4.4.3.2 The influence of rotation frequency on motivation

As soon as an agent uses a skill, the motivation decreases until the skill is no longer used. After that, the motivation increases again, until the skill is used again. The periodicity of this in- and decrease is determined by the rotation frequency. A rotation frequency of 5 implies that it takes 10 cycles from the first use before the motivation has reached its original level again. Figure 18c shows a situation in which there is a rotation frequency of 25.

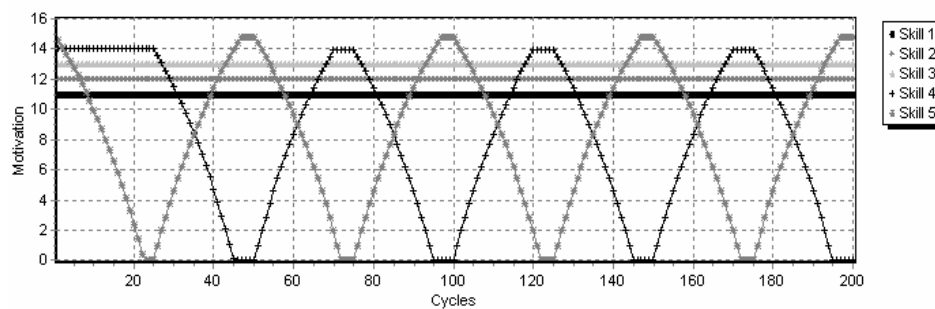


Figure 18c: Motivation at a rotation frequency of 25

We see that the skill with the highest motivation rate (green) starts decreasing, reaches zero at the 22nd cycle and maintains this level until the next rotation at the 26th cycle. Then, the motivation increases again until it reaches its original value at the 47th cycle. Simultaneously, the agent uses its second best skill, following the same periodicity as the highest skill.

4.4.3.3 The influence of rotation frequency on performance

The first situation, in which the rotation frequency takes place once every 5 cycles, is described in section 4.4.1.3 (see also Figure 16c). During the time period of 5 cycles there is some time to learn, to forget, to get bored and to recover from it, which causes small fluctuations. Apart from these fluctuations the performance time slowly decreases due to the increase of the expertise and finally reaches a value of 14. In the second situation, in which there is a rotation frequency of 1/25, we see the development of a cyclic change (see Figure 18d).

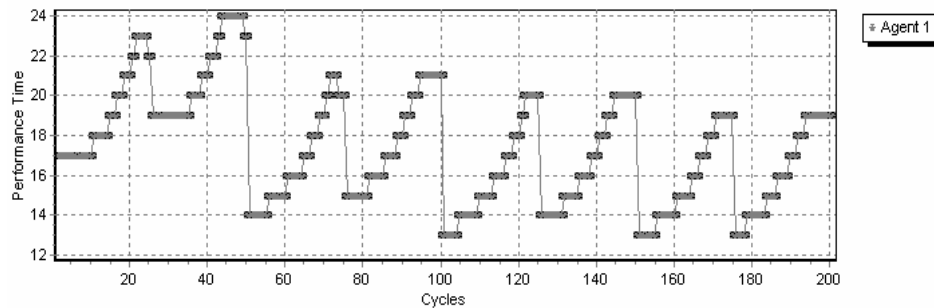


Figure 18d: Performance time at a rotation frequency of once every 25 cycles

During the first period of 25 cycles, expertise is build up in the best skill, but boredom will nevertheless cause an increase in the performance time. After the task rotation at the 25th cycle, the performance time starts at a lower level because now the agent uses its second best skill. During the first 25 cycles the expertise in this skill dropped resulting in a higher initial start level of the performance time and a higher level at the end of the second period at the 50th cycle (see also Figure 18a). Nevertheless, the expertise in both skills has increased over 20, causing the performance time in each period to start from value 14 from the 51st cycle, after which as a result of boredom it increases until it reaches the value of 21. Then the agents rotate, the performance time shifts back to its lowest level and the same process starts all over again. The differences in performance time at the start of every period (1st, 26th, 51st, etc.) decrease because the differences of the best and the second best skill will decrease once they have both reached their maxima (see also Figure 18a).

The figure shows that it takes about 50 cycles, i.e. two rotations, before this periodicity has been adjusted. If we lower the rotation frequency to 1/50, it again takes two rotations, to adjust this periodicity, in which the performance time now shifts from 11 to 19. If we lower the rotation frequency to 1/75, we again see that it takes two rotations, i.e. 150 cycles to adjust the periodicity. When studying the results of a task with more than 200 cycles, we found out that in the latter situation the periodicity of the performance time shifts from 10 to 19. If we compare the different situations, we see that the slowest performance time decreases slightly as a function of the rotation frequency. A lower frequency results in more time to recover from boredom, which leads to a higher motivation.

4.5 Conclusions and Discussion

4.5.1 Conclusions

With regard to self-organisation, we conclude that the influence of self-organisation on expertise and performance shows that in the situations where task rotation emerged, the system delivered a better performance than in case of no self-organisation. This finding

suggests that, when performing a task, workers must have the freedom to rotate tasks whenever they feel bored. As regards boredom/recovery, we have arrived at two conclusions: First, if the boredom/recovery rate decreases, it takes more cycles before the agents start to rotate tasks. Consequently, it takes more time before they develop a second skill. If the boredom/recovery rate drops beneath a certain point, the agents will specialise in only one skill, because the time it takes to decrease the motivation enough to rotate tasks exceeds the time that it takes to forget the second best skill. Second, if the boredom/recovery rate decreases, at a certain point the motivational processes become slower than the expertise processes. This implies that, with respect to a particular skill, at the time the agents are bored, the level of their expertise has become quite high. This results in a situation in which the agents only rotate tasks when their level of motivation is very low and rotate back before they have fully recovered. As regards the task rotation frequency, we conclude that with respect to expertise, the decrease in the rotation frequency has the same effect as the decrease in boredom/recovery. We found no significant effects of the interaction of self-organisation and boredom or boredom and task rotation frequency, except in the situation with a low degree of boredom: This situation is ‘close to the edge’, which means that in the setting with a high task rotation frequency, which we used to manipulate by adjusting the boredom/recovery rate, the agents still rotated tasks. But as soon as we lowered the boredom/recovery rate from 29 to 28, the edge was crossed and task rotation did no longer emerge. The same happens if we lower the task rotation frequency. In the other situation, a frequency of once every 78 cycles still led to task rotation. However, in this particular setting the process of task rotation did no longer occur at a frequency lower than 1/25. Therefore, if an organisation or a work group needs workers that are capable of using multiple skills, for instance to create flexibility, the components described in this study should be taken into account. A low task rotation frequency and a task that is interesting (low level of boredom) more easily lead to the specialisation in one particular skill than a boring task or a high task rotation frequency does.

4.5.2 Discussion

In the present study we used expertise and motivation as components that determine group performance. These components can be considered as elements within a broad range of factors that affect team performance, such as work-related attitudes, team composition, commitment, and team cohesion (e.g. Cohen, Ledford & Spreitzer 1996). Although motivation and expertise are important components that affect performance (Wilke & Meertens, 1994), it is obvious that they do not cover all the other factors. Moreover, we supposed that processes, such as getting bored, were solely influenced by the repetitiveness of tasks, whereas it is likely that several other factors might cause boredom, such as for example, the physical condition of an agent. Factors that we consider to be fruitful extensions of our model are, for example, coordination costs (e.g. Cohen et al., 1996), task interdependence (Van der Vegt & Van de Vliert 2005) team size and team composition (e.g. Molleman 2005). Regarding coordination costs, prior tests have indicated that simulation experiments could only produce plausible outcomes if the interactions among agents led to the emergence of rotation rules and

routines that decrease coordination costs (see Zoethout et al., 2006b). Task interdependence might seriously affect the possibilities to divide tasks among agents, and the composition of teams in terms of team size, the distribution of skills, demographic characteristics and personality traits has proven to influence team functioning (Molleman 2005). Inclusion of such factors will bring our model more close to reality, but will also make the results much more difficult to interpret.

Several studies have elucidated the benefits of job rotation within teams (e.g., Van den Beukel 2003). These studies have led to a so-called design-based view on organisations, i.e., since job rotation has proven to bring forth advantages, management should implement it. As a consequence, job rotation has been implemented by management in various settings, even in work designs that are considered as self-managing teams. This may raise the question whether job rotation that is designed for and implemented in a system by an external party has the same effect as job rotation that has spontaneously emerged from the part of the workers themselves. Of course, the discussion about designing job rotation externally versus spontaneous development within the organisation itself involves a lot more issues than dealt with in the present study. Nevertheless, our use of computer simulation has made it possible to start comparing both approaches. We therefore conclude that simulation studies of this kind contribute to the understanding and analysis of the social dynamics of work groups.