TOWARDS MAS REORGANIZATION THEORY

Jeroen Heijmans
16 december 2005

Master's thesis Cognitive Artificial Intelligence
Department of Philosophy, Utrecht University

First coordinator
dr. F.P.M. Dignum
Department of Information and Computing Sciences
Utrecht University

Second coordinator
dr. M.V. Dignum
Department of Information and Computing Sciences
Utrecht University

Third coordinator
dr. V. van Oostrom
Department of Philosophy
Utrecht University

TOWARDS MAS REORGANIZATION THEORY

IN MEMORIAM

Het leven is inderdaad een omweg waard.

(Life's indeed worth a detour.)

Short Table of Contents

FULL TABLE OF CONTENTS

Preface

Abstract

MAS organizations are becoming more and more useful as they come closer to- as well as more alike human organizations. We can improve theory on MAS organizations by looking at human organizations and trying to simulate their behavior. This thesis specifically explores how human organizations are reorganized and how theories on these can be useful in MAS. To know when to reorganize any organization we need to be able to judge the state of affairs in an organization. To be able to judge upon these, we need criteria on these states of affairs. I explore how these criteria are defined and formed in human organizations. Then I try to abstract from this, and translate it to the field of MAS. Finally, I try to apply this abstraction in a simulation, to see whether it's a useful abstraction.

Thesis Information

This thesis was written to obtain the doctorandus degree (MSc) in Cognitive Artificial Science at Utrecht University in The Netherlands. The exam takes place on December 16th, 2005, 12.00 hours at Utrecht University. It was written between June and December 2005. It has a total count of around 19000 words, spread over a total of 63 pages.

The cover illustration is a personal adaptation of the Penrose triangle. Though various interpretations of this illustration are possible and might (or even should) be considered, the most obvious interpretation is analogous to having multiple views on organizations. In this case the eyes represent the views we can take, the triangle represents reality, and the circle shows that all this is interconnected.

Acknowledgements

A thesis like this one is quite a difficult project. Things were extra difficult as circumstances forced me to take a long break about halfway through. All these difficulties, trivial or not, were overcome with the help of other people.

Many thanks go to Frank Dignum and Virginia Dignum, who together gave me enough room and time to cope with many difficulties. They also did a terrific job at stimulating at times when it was most needed. Together with Vincent van Oostrom, all three of my supervisors took plenty of time and effort to judge this thesis in a rather short amount of time.

Moving on, I only got this far in the first place because of my family. They have all supported me throughout my life, and also very much during the last few months. Special thanks go to my uncle Ger, who would always be the first to respond to my e-mails about my work. To use a nice (but very applicable!) Hollywood cliché: Ger was like a father to me. This also brings me to my dad, whose strength and cunning inspired me from the beginning until the very end.

Finally, I want to thank my private muse for inspiring me, keeping faith in me and sticking with me even though much of my time went into writing this thesis.

I very much enjoyed this thesis project, and am therefore very grateful for all the abovementioned folks as well as others who have helped me through this project. Thank you.

1. Introduction

We need to move towards theories on MAS (MultiAgent Systems), and specifically on MAS reorganization. Why, and how can we do this? MAS are becoming more and more widespread, and work together with humans and human organizations. For example internet auctions often already include bots (the agents) negotiating a deal for users. We must study MAS theories to improve this interaction as well as the regular functioning of MAS. Because of the interaction between MAS and human organizations, there will be some general principles that do or might govern both types of organization.

MAS theory on general principles governing issues such as reorganization is very limited. On the other hand, such issues have already been studied for a long time considering human organizations. The idea here is that we might be able to extract some general principles about human organizations, and use them somehow to build MAS theory. A first step in this process is the extraction of general principles governing issues such as reorganization of human organizations. This is the subject of this thesis.

1.1 Why do we do simulation?

Multiagent systems are useful for a wide range of applications. In this thesis I will be looking at using MAS for building simulations of real settings. This field, called multiagent based simulations (MABS), is a growing field of research. It promises a contribution to informatics as well as to other disciplines where the simulations could be used; these disciplines include (but are not limited to) psychology, biology, and economics.

But why do we do simulation? In [12], titled "Having Fun Being Useful", the authors start with this question. In the first part they claim MABS should be aimed at actually contributing something to the mentioned other disciplines. It is no coincidence this introduction is titled as it is: I also consider it to be important to (eventually) be able to build simulations that are of some use in disciplines besides informatics. I even think this is the main goal.

To contribute to the multidisciplinary nature of MABS, this thesis will be aimed at researching the parallel between real situations and realistic simulations. This will hopefully also provide some useful insights on the relation between MAS- and traditional organizations. In this introductory chapter I will start by exploring some unanswered questions on this topic. After that I will explain the structure of this thesis.

1.2 Leading up to the Question

MAS is a relatively young field of research. Many fundamental questions still require (fundamental) answers. Obviously, not every question can be answered in one piece of work; you have to focus on a particular question instead. I will examine a sub-field of MAS research, and then draw some unanswered questions from research already done in this area. I will refine this to pose the questions I will be trying to answer in this thesis.

1.2.1 Reorganization

A MAS can be used to simulate just about any real chain of events. Doing this requires us to construct a mapping from real situations that occur during this chain to a computer simulation. Work in this area has already provided some good results. However, much work has still to be done in making software agents respond to changes in the environment in a more human way: as outlined in [5], "reorganization is the answer to change in the environment". Thus, if we are going to create simulations where software agents respond to changes realistically, we will need to implement some form of reorganization. This holds not only for human organizations, but also for other kinds of organizations including MAS.

Reorganization comes down to changing the way agents are organized. This can be behavioral change (the way the agents act, what roles they have, and so on) or structural change (adding or removing new/different agents). This is extensively described in [5]. Reorganization has several interesting properties. We can investigate these different properties individually, giving a more fundamental insight in the concept of reorganization. Researchers engaged in researching organizations and reorganizations already recognize the need for answers to these fundamental questions. In [10] some basic questions about reorganization are posed using the five W's of the English language:

Quote: from [10], p. 53:

- What are the aspects of an organization that will be reorganized?
- Who has authority to take reorganization decisions, and how are they taken?
- When should reorganization occur?
- Why should we reorganize, what are the strategic reasons for reorganization?
- Whether we should reorganize, i.e. what is the threshold for reorganization, when is reorganization likely to be beneficial?

As these questions are important for future research on multiagent based simulations (MABS), let us look more closely at some of them.

1.2.2 Research Questions

In this thesis, I will be looking more closely at issues related to the question: “When to reorganize?”. Answering this question has traditionally been a task for a human, and professional business managers have been occupied with this question for at least several decades. Managers have to make a lot of decisions related to reorganization. The most important problem here is that the manager cannot *try out* several options: he has to make a decision with no way back. This is where MABS comes in. Simulations can be used in a decision-support tool, where the manager actually can do something that comes close to *trying out* his options: he tries different options in different simulation runs. The simulation then shows the result of the manager's decision. Such a decision-support tool can even implement a software agent to simulate a manager, whose results can then be compared to the results of the human manager, to improve the strategy of the manager (the human- or software agent) that's performing the worst.

In [11] such a decision-support tool has already been applied successfully in the domain of police patrol routes. Also, a large-scale project is now running at the Dutch TNO DECIS-lab called ICIS (see [21] and [18]), where researchers use multiagent based simulations of crisis situations to assist human authorities in their decision making process. However, these projects seem to be successful mostly because they work in specific domains. The properties of the reorganizations (criteria for when to reorganize, what is being reorganized, etc) are drawn directly from the domain. This is not so much bad practice (as it gets research started), but it sustains the current lack of general principles governing reorganization. So, as opposed to these practices, I would like to know more about reorganization on a level more abstract than this. This leads to several research questions focusing on certain properties of reorganization:

Research Questions:

1. What are the criteria for evaluating the state of affairs in an organization?

    A. What is the level of domain-dependence of these criteria? How do we get these criteria?

    B. How can these criteria facilitate formal models that allow the specification of dynamic reorganization of agent societies?

The main question is there to find out what the triggers for reorganization are. The two sub questions allow me to investigate respectively the domain-dependency and uses of the answer to the main question. All the knowledge gained by trying to answer these questions can be applied not only to natural organizations (human organizations, animal organizations, and so on), but also to some extent to MAS themselves. The extent to which this can be done depends on the level of correspondence and connection between MAS and human organizations. This correspondence and connection is subject of research such as [3].

Reorganization decisions have to be made based on the evaluation of an organizational state of affairs, just as this is being done in real situations; real situations that we'll be trying to simulate.

For example

Suppose a pharaoh is building a pyramid, and its progress evaluates to be too slow, the pharaoh might choose to reorganize by capturing extra slaves. The pharaoh evaluates the progress to be too slow by estimating that the finishing time is too late. If we are simulating the pharaoh and his organization, we would need criteria for evaluating the state of affairs to get a good simulation that includes well-timed reorganizations.

The pharaoh-example shows that the main research question is important in building simulations including reorganization. To answer this question, I will look at research from MABS as well as organizational theory (OT). Based on this I will propose an answer as to how we could evaluate the state of affairs in any organization: natural organizations, simulated organizations and MAS organizations alike. While searching for an answer to the main question, two related questions will come up (1.A and 1.B, above). Sub question 1.A actually comes down to stating how generic the answer to the main question is; in other words how much domain knowledge is needed to build a criterion for evaluating the state of affairs in an organization. The second sub question 1.B was proposed as future research in the conclusion of [5]. In the context of this thesis, answering this question actually comes down to finding a way to incorporate the answer to the main question in a simulation.

If and when I have found a possible answer to the main question, I will demonstrate how this works in a simulation. This simulation will then show that evaluating the state of affairs in an organization is important for reorganization, and that it can be done in a way true to the real situation being simulated.

1.3 Thesis Context

When reading this thesis it's important (as it is with every scientific essay or paper) to keep the context in mind. This thesis is all about multiagent based simulations. MABS is a field that brings cognitive- and social sciences together with informatics (Figure 1).

This thesis is about multiagent based simulations, but obviously the research questions I posed in the previous section are more specific. This is also reflected by the literature (see references). Next to that, I draw knowledge from other sciences, including (but not limited to) economy and social sciences. These fields of research are generally a lot older than MABS, and very well researched. This is both an advantage as well as a potential pitfall! The advantage is that using rich and well-founded theories gives us many options in simulating the real phenomenon. The pitfall is that we slipstream these theories

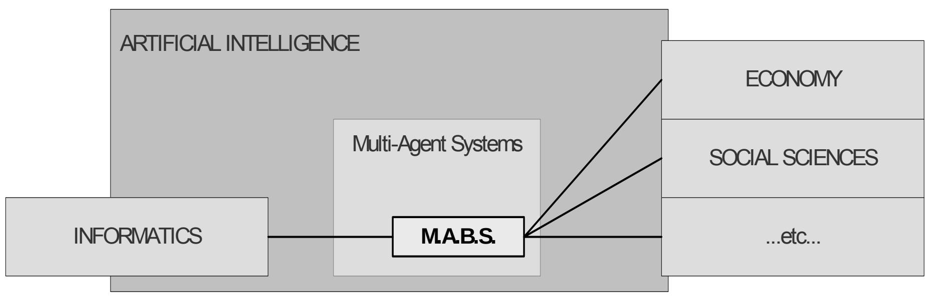

Figure 1

into our simulations too much too quickly. This would make a mess of simulation experiments, because too much complexity was added at once. Because the advantage and the pitfall could respectively lead to success and disaster, I follow two guidelines that are important in using existing theories from other disciplines when researching MABS:

Guidelines

1. Take advantage of the richness of existing theories from other disciplines.
2. Beware of the pitfall of using existing theories too much too quickly.

These guidelines are important because they try to get the greatest effect out of existing theories. The guidelines are obviously conflicting. To take advantage of the existing theories, I will examine them and try to incorporate their knowledge into my answer to the main research question. To avoid the pitfall, I propose that the answer will be implemented in simulations only one step at a time. The simulation in this thesis will therefore only be the first step, showing that the answer is viable. Integrating things even further is left to future research.

To summarize, the answer to the research questions will be constructed by combining insights from disciplines (mostly) outside the field of Artificial Intelligence with current simulation approaches. These answers will be (partially) demonstrated in an example simulation. Finally, all this knowledge will hopefully allow us to draw some conclusions on natural-, simulated-, *and* MAS organizations, and the links between those organizations.

1.4 Thesis Layout

In this thesis I describe a search for answers to the research questions. In this first chapter I have introduced these questions and their context. To answer these questions, I will first examine literature about natural organizations. This will be the subject of *chapter 2*. In that chapter I shall try to pose some preliminary answers. In *chapter 3*, these answers will be viewed from the context of informatics, as computer simulations and MAS are sub fields of this field. In this chapter, I move towards using these answers in an actual simulation. This simulation is the subject of *chapter 4*. A small economic setting is introduced to test the answers to the research questions, and examine the links between simulated organizations and natural & MAS organizations. Finally, in *chapter 5*, I summarize the answers to the research questions and results from the simulation experiments. From this I'll draw several conclusions about reorganizing organizations and the links between the different types of organizations.

At the end of this thesis, two appendices are included. First, *appendix A* can be used with the simulation as it explains the organization parameters in detail. Second, *appendix B* shows how to use and interpret the output of the simulation experiments.

2. Organizations

To build simulations, we must have a clear view of the things we want to simulate. With MABS, our subject is any organization of agents, for example a society of animals or a group of friends. In this chapter I examine a well-studied type of organization: the human organization. A lot of literature about human organizations exists, which can be useful for answering the research questions posed in the previous chapter.

I will start by introducing terminology I wish to use when talking about theories on human organizations meeting MABS research. Then I move to the theories themselves: *what* theories can be used and *how* can they be used? At the end of this chapter I extract the useful information from theories about human organizations so that it facilitates answers to the research questions.

2.1 Terminology

Before going into more details of the subject matter, I will first introduce some terminology. I don't wish to build my own complete agent terminology: some very talented researchers have tried this and are still trying. Building clear and well-defined terms is a general problem for current agent research. For example, we do not even yet have a unique, widely accepted definition of the term 'agent'. However, some researchers still do a very good job summarizing things and presenting overviews. For a clear overview of terminology used in agent research refer to [19]. For basic terms such as "agent" I will adhere to definitions from that book.

So I use basic terms from [19]. Some terms I wish to define, though. In my main research question ("What are the criteria for evaluating the state of affairs in an organization?") a very central concept is mentioned: the *state of affairs*. Intuitively, the state of affairs is what an organization *is like* at any point in time: it's a snapshot. This leads to the following definition for the organization's state of affairs.

Definition: State of affairs
The entire set of properties of the organization and related entities, at a certain point in time (shorthand: "properties of an organization").

In this definition 'related entities' include agents, relations with other organizations and so on. What the 'properties' are depends on what type of organization you're looking at. Any property can be either quantitative or qualitative. Evaluating these properties can then be done in any quantitative or qualitative fashion, respectively (I will come back to this in chapter 3). Not all properties need to be used in evaluating the state of affairs. For example, in Artificial Life simulations an organization might be evaluated solely on the average health, leaving out other properties such as the total number of creatures. The evaluation of properties is defined in the following fashion:

Definition: evaluating (a state of affairs)
Measuring upon the state of affairs of an organization in any qualitative or quantitative way.

With these terms defined, let's look once more at the main research question "What are the criteria for evaluating the state of affairs in an organization?". Phrased using the above definitions, this literally comes down to the following:

Main research question, rephrased
"What are the criteria by which we can measure upon the properties of an organization (and related entities), either in a quantitative or qualitative way?"

To find these criteria and measuring methods several approaches are possible. In the next section I will explain what approach I will be taking.

2.2 Approaches with MABS

2.2.1 Different Approaches

As stated in the research question, we are searching for criteria and measuring methods for the properties of an organization. To find these there are two approaches, which I call the descriptive and prescriptive approach. The first approach usually comes down to examining actual cases (such as a business manager occupied with the research question in some form or another), and then tailoring an answer to the research question accordingly. In the second approach, we construct a very generic answer to the research question by using domain independent tools (for example logic). Before discussing both methods, I will explain where the difference of these approaches comes from and how it is relevant to the research questions of this thesis.

The distinction between the different approaches is best described analogous to a very heated discussion in the field of philosophy of science. Scientific realists and empiricists debate on the way science should be conducted. To sketch this debate I will use somewhat simplified definitions of realism and empiricism, based on papers by Van Fraassen [9] and Musgrave [13]. First, an empiricist claims science has to describe what reality is like; a theory is true or false depending on whether it is "empirically adequate": whether our observations can be explained by our theories. The realists on the other hand suppose that science studies subjects that really can (and sometimes even really do) exist. Our scientific theories will then be true or false depending on the real properties (and existence) of these subjects.

The debate between empiricists and realists is a very fundamental one. It is about how we *should* be conducting science. However, in fact history determined for a large part whether science is conducted the empiricist- or realist way. For example, the geocentric model[1] was considered true for a long time, as it was empirically adequate considering measurements from early science. As a different example, recent history has shown the development of advanced mathematical techniques, which are useful for posing scientific (realist) theories about things such as subatomic particles: things that might even *never* be measurable.

So, how is this debate from philosophy of science useful in choosing an approach in answering the research question of this thesis? As an empiricist MABS researcher, we would use the descriptive approach. To find the criteria by which we can measure the properties of an organization, we should indeed consider several real cases. Researching these cases then allows us to "describe what the criteria are". The quality of the criteria is then based on how empirically adequate they are. On the other hand, if we are building a MABS as a realist researcher, we would use the prescriptive approach. We would try to find a generic way to formulate the criteria: we would use our research to "prescribe what the criteria are".

Now that we know the difference between empiricist and realist MABS researchers, all we've got left to do is to pick a side. But this raises a new question: do you need to discover whether you are a realist or empiricist before choosing an approach, or do you just pick one approach and accept the philosophy of science that comes with it? Finally, starting with just picking an approach, we could also adopt a fusion approach: try to take the best of both. As I already mentioned in the introduction, I agree with the authors of [12] that we must choose a course of action leading to usable results. A logical choice is then actually to try and adopt the best of both the descriptive and prescriptive approach. Note that while using such a fusion approach we are in danger of merging parts of both approaches that can't be merged. However, as long as we evade this danger by using only one approach at a time (though we might use different approaches at different times), this fusion approach has the best chance of getting highly usable results.

Let's now turn to the next two sections, to see what parts of these approaches are useful in answering the research questions of this thesis.

2.2.2 Descriptive Approach

To recapitulate, this approach comes down to trying to describe what the criteria are that are being used to evaluate the properties of an organization, by considering real cases. This approach is the most domain-dependent type. Establishing criteria to be used in a MABS will most likely include an in-depth analysis of a particular domain and how domain experts would evaluate the properties of an organization. Getting this knowledge from experts and formalizing it can be done using a wide array of techniques such as the CommonKADS methodology [17].

[1] Paradigm placing the earth at the center of the universe.

For example

Suppose we are to build a MABS of a predator/prey situation. We could then use protocol analysis, giving the biologist an example and a task such as “determine whether a specific situation is stable”. Using the protocol analysis technique we can then try to extract what criteria the biologist uses to complete the task. These criteria can then be generalized, checked with the expert and formalized to fit into a MABS.

There are several advantages to this approach. While we are extracting the criteria from the domain (expert) through techniques such as protocol analysis, it is usually also easy to find out how these criteria can be *used*. A perhaps even greater advantage is that this approach allows us to draw upon the immense body of expert knowledge that is usually available when building a MABS.

Several nasty disadvantages also come with this approach. First, we must consider the ‘pitfall’ already mentioned in section 1.3: don’t use too much knowledge from other disciplines too quickly, because then the complexity will get out of hand. This problem is even bigger when you consider the fact that domain experts don’t always have well-defined and/or well-structured knowledge of their domain [17]. Because of this you sometimes even cannot really avoid stepping into the pitfall. The second disadvantage is that you are automatically using a very domain dependent solution. Starting with very specific instances of criteria, you can only generalize up to a certain hight. This makes your simulation and techniques developed for building future MABS also domain dependent.

2.2.3 Prescriptive Approach

This approach is the one taken by the realist MABS researcher, and comes down to searching for a generic phrasing of criteria and measurements on properties of an organization. This approach is less domain dependent than the previous one. In fact, it tries to abstract from specific cases, and be as domain independent as possible. Considering MABS, there is a wide variety of methods already available when taking this approach. Whereas the descriptive approach is more likely to be used by business people in need of concrete solutions, the prescriptive approach is more commonly taken by scientific MABS researchers.

Two prominent methods are good examples of this approach. First, various different kinds of logic, such as modal and deontic logic are being used for modeling (organizations of) agents and their actions. This is because researchers coming from the classical approach of AI already used logic to model agent reasoning: it was only a small and perhaps even logical step to extend these kinds of logic to be able to model features introduced along with the notion of agent organizations. The second method used within the prescriptive approach are MAS methodologies such as GAIA [20], Prometheus [14] and OperA [3], offering a way to specify agents, multiagent systems and so on.

The advantages of this aproach are for the large part the opposites of the descriptive aproach's disadvantages. The prescriptive approach is very useful (and even aimed at) building general solutions for MAS. Usually you don't start with tailoring your solution (for example a new kind of logic) for a specific domain. MAS methodologies are also usually built with no specific application in mind. Another advantage of this approach is that although the solutions are general, it's easy to build more specific instances from these solutions to tailor some application. To put this in terms borrowed from (database) modeling[2]: it is easier to specialize from a general prescriptive solution, than to generalize from a descriptive solution.

Of course, some disadvantages with this approach also exist. A practical objection to this approach is that it can take quite some time to perfect (or even build) the solution. Relating this to the philosophy of science analogy: it is usually a lot easier to determine whether a solution is empirically adequate, than whether it truthfully describes the world as it is. Another disadvantage is that it takes an extra step to come to an actual solution: first a method (such as logic) has to be developed and perfected, then this method still also has to be implemented in an actual solution (such as a simulation).

2.2.4 Fusion Approach

As said before, I will try to use the best of both approaches when trying to answer the research questions in this thesis. This thesis and its research questions are obviously not directly aimed at building a particular simulation or tool. Starting with the prescriptive approach in mind, I will try to answer the research questions in a generic fashion. This answer can then be specialized to tailor a simulation or tool for a specific application.

After I have a generic specification of the criteria by which an organization can be evaluated, I will show how this generic specification can be specialized to tailor a simulation. In chapter 4 I describe a small economic setting and use aid of the descriptive approach to form domain-specific criteria. These criteria will be constructed in a way conforming to the more generic (prescriptive) specification of criteria. At the end of chapter 4 experiments with this simulation will show whether the generic specification was good enough to be of practical use in a real simulation.

As a final note on this approach, I need to point at the discussion between Edmonds, and Dignum & Sonenberg in [6], [2] and [7]. Jumping directly to the concluding article [7], we can see that there are quite some concerns with overusing the prescriptive approach, ending up with formalisms that are of no real use. This relates directly to my second sub question: "How can these criteria facilitate formal models that allow the specification of dynamic reorganization of agent societies?". Although I would like my answer to the main research question to facilitate the specificiation of formal models (possibly using logic), the conclusion of [7] warns that (some of) these formal models must be useful. This is then another good reason for trying to apply the

[2] In fact, the issue that specialization seems to be easier than generalization is found in many different disciplines. For example this issue is also found in lambda calculus.

criteria for evaluating organizations in a real simulation (as is the subject of chapter 4), because we can then see whether any formalism built with answers to the research questions would possibly be useful. However, before going on to a specific simulation, I will first specify criteria for evaluating an organization in a generic fashion. This is the subject of the rest of this chapter.

2.3 Organization Theory

MAS can be used to model any type of organization. This is largely because an agent is useful for modeling many autonomous entities, including animals and humans. In addition to this, agents can also be autonomous pieces of software, not really used to model some other entity. Organizations of these three different types of agents are all studied in different disciplines, respectively biology (animals), sociology and economy (humans) and informatics (software agents). The first two organizations fall in the category "natural organizations", the latter is a MAS itself. MABS can be built for all these types of organizations, or even hybrid types. Especially with human organizations there exists a vast body of literature on how to evaluate the state of affairs. Following the first guideline from section 1.3, I would like to take advantage of this literature. Therefore, the remainder of this thesis will be about human organizations, and correspondingly organization theory.

Studying organization theory (OT) is useful within the prescriptive approach, because it tells us (in a generic fashion) what human organizations are like. We can use this information to find a way MABS *should be built*. Moreover, OT is about the kind of natural organization (human organizations) that is most likely to interact with MAS organizations. OT is then most likely to give results useful for studying that interaction. In the next subsection I explore how OT proposes to evaluate the state of affairs in an organization. After that I will discuss how well this information can be used to answer the research questions.

2.3.1 Evaluating a State of Affairs

A very extensive overview of the field of human organizations is given in [16], where many popular and/or prominent theories are presented. Chapter 15, about organizational goals and effectiveness, is related directly to the question of how to evaluate the state of affairs in an organization. I will summarize parts of this chapter that are important regarding my thesis. As a starting point, the definition of goals in [16] is taken from Etzioni [8]:

Definition: organizational goals
A desired state of affairs which an organization attempts to realise.

This fits nicely with the terminology introduced in section 2.1, as it uses the term "state of affairs". This definition still has a minor problem, as it is not quite clear how an organization can "attempt" to realise something. In case of this thesis, the intuitive notion of "attempts" will suffice. Note however, that in other research it might be necessary to define this term more precisely. In OT, several other concepts related to "goal" exist, such as "objectives", "mission" and "strategy". As these don't relate directly to the question of how to evaluate the state of affairs, I will leave them out of this thesis. What *does* relate directly to this evaluation, is the term effectiveness[3]. I will adopt (what is called in [16] the best-known approach) the "Goal Approach" for measuring effectiveness:

Definition: effectiveness (Goal Approach)
The extent to which an organization achieves its goals.

Let's go back to goals for a bit. We already have a proper, widely used definition. For goals however, at some point we are going to need instances of goals (conforming to this definition). Ever since it has been clearly stated in [15], it has been accepted that just about any organization will have more than one goal (moreover, these goals can usually not be joined into *one*, more general goal). This is for a large part because goals can (and often will) be conflicting. Consider the following.

For example
Suppose we are looking at a factory building weapons such as firearms. This organization wants to produce their weapons as cheap as possible (goal 1). Also, they need to make sure that the weapons are of sufficient quality (goal 2). Lastly, they need to build these weapons in a safe way: the organization doesn't want employees to get injured while building weapons (goal 3). Goals 2 and 3 usually increase costs, which is counter to goal 1. Now suppose also that the higher the needed quality is, the more dangerous it becomes to build the weapons. Then goals 2 and 3 are also conflicting.

The above example demonstrates that organizations with multiple goals face conflicts in choosing their actions. These different goals come from different 'stakeholders': people who want something from the organization. In the above example goals 1, 2 and 3 are posed by respectively the organization itself, the customer, and the employees. These stakeholders can be classified according to a variety of typologies [15]. What is *important* here though, is that the organization can, and most likely *will* have several different (possibly conflicting) goals, because different stakeholders approach the organization with different perspectives. These perspectives should then also be included in a simulation of a human organization.

With the definitions of goals and effectiveness, and the notion of perspectives, I will now review the research questions once more. Before going on to this topic a final note on human organizations has to be made. With human organizations it is common practice (as is done in the first chapter of [16]) to view an organization as an artificial entity brought into existence to serve a purpose. In agent research on the other hand, this restriction usually doesn't exist: an organization can represent just about any collection of interacting agents. The thesis presented here is about human organizations. It might be generalized to reason with and about organizations in general, but special care

---

[3] Effectiveness is easily (and therefore often) confused with the term effeciency. Generally speaking, the first term refers to whether the organizational behavior is appropriate, the second term is a measure for resource usage. This confusion of terms also exists in computational organizational theory (COT), as is discussed in [1].

has to be taken when this is done. In this thesis, I will stick with human organizations, leaving any generalization to future work.

2.3.2 Research Questions

So, once more: what are the criteria for evaluating the state of affairs in an organization? Obviously, this directly relates to the organizational goals. Because there are multiple perspectives, there will be multiple ways of evaluating the state of affairs. A goal is formed by a desired state of affairs. We can then evaluate a current goal by comparing the corresponding desired state of affairs with the current state of affairs. The 'distance' between these states of affairs is then a proper evaluation of an organizational goal.

The criteria by which this measuring can be done depends on the specific goal. To see this, we must look at the use of the word 'distance', above. Although distance usually refers to a quantitative measurement, the use of this term is to be taken more abstract here: the distance between states of affairs can be quantitative as well as qualitative, and it can in principle even be hard or impossible to measure. The very word 'distance' however, does imply there is a certain metric. I here assume that there is such a metric. Distance is then defined as:

Definition: distance
The distance between different states of affairs is a value for a qualitative or quantitative measurement function upon pairs of values of organizational properties.

The criteria by which a state of affairs can be evaluated is the specification of how to measure this disctance between states of affairs. This unfortunately means that we have only a vague definition of what criteria are: we have only moved the problem to having to define what makes up the function determining the 'distance' between states of affiars. Only if we take a particular (domain specific!) organizational goal, we can get rid of this vagueness. To put criteria it in the form of a definition:

Definition: criteria (to evaluate a state of affairs)
Specification of how to measure the 'distance' between two states of affairs.

Although this definition actually only moves the problem, it at least gives an intuitive notion of how criteria can be specified. Also, it provides me with a way to construct a preliminary answer to the main research question, on what the criteria for evaluating the state of affairs of an organization are:

Preliminary answer: question 1
A specification of how we can measure the 'distance' between the current properties of the organization and the desired properties of the organization.

This is then also the end of using the prescriptive approach to answering the research questions. If we want to know what 'distance' really means we need to start using the descriptive approach: some kind of domain analysis is required to answer the research questions more specifically. Note that all this leads directly to an answer to the first sub question: "What is the level of domain-dependence of these criteria? How do we get these criteria?". The domain-dependence of the actual criteria now turns out to be very high: although we know that any criterion must be a specification of how to measure distance, the actual metric for doing this will be domain-dependent. The answer to the second part of the sub question then follows automatically: we can get these criteria by analyzing the domain.

2.4 Chapter Summary

In this chapter I have started by clearly stating what terms and definitions I use. After that, I have explained two approaches to answering the research questions. The descriptive- and prescriptive way, corresponding to respectively the empiricist's philosophy of science, and the realist's philosophy of science. The approach taken in the rest of this chapter is a fusion approach taking the best of both, starting with mostly the prescriptive approach.

Using this fusion approach I tried to gain insights from organization theory, where a lot of generic theories about organizations already exist. These theories lead to a very clear definition of organizational goals and organizational effectiveness. Finally, while trying to find a definition of the criteria by which we can evaluate the state of affairs, we stumbled upon the boundry of domain-dependence. Criteria turn out to be the way to measure the 'distance' between states of affairs. The actual criteria are very domain dependent, however. When these are needed for a simulation, domain analysis will have to reveal the true nature of the criteria by which we can evaluate the state of affairs in an organization.

In the next chapter I will move on by translating the knowledge summarized here to the field of MABS. I will also be looking at the second sub question: how can all this help in building formal models?

3. Agent Organizations

In the previous chapter I posed some preliminary answers to the research questions, based on OT. In this chapter I will be describing how these preliminary answers fit into the field of MABS.

This chapter is organized as follows. I will first discuss how information from OT from the previous chapter will be translated to informatics, such that it will be usable in the field of MABS. I will then go on to discuss MAS methodologies. This relates to the second sub question: how this thesis facilitates building of formal models. After doing this I will discuss MABS itself, how it is used and how insights from this thesis can be used to make MABS more effective. I will also discuss the use of MABS in decision-support tools. At the end of this chapter I discuss how the insights from this thesis could be used in a real simulation. This forms a bridge to the next chapter.

3.1 Translation to Informatics

As a subfield of Artificial Intelligence (see section 1.3), MABS is a very multidisciplinary field of research. Different disciplines connected with MABS use their own terminology, however. This is also the case with Organization Theory. The information posed in the previous chapter needs to be translated to be usable in a MAS context. This translation has to be done in such a way that everything fits nicely into existing MAS theories and methods.

I will first discuss the difficulties with translating terminology to informatics in general. Then I will go on to the actual translations that help construct final answers to the research questions.

3.1.1 Translation Process

Translating real world concepts from fields of research such as economy to informatics is well studied. Researchers have developed methods over the past 50 years to represent properties of objects such as organizations explicitly. In the field of statistics this is also called operationalizing: making some real world item measurable. If we take the word 'measurable' to be either quantitative or qualitative, translating actually comes down to operationalizing. This technique of operationalizing is used widely in making surveys, just as in the following example:

For example
Suppose an organization holds a survey amongst its employees. A question could be included to investigate whether employees use time at work to do private things, such as checking their private e-mail. To investigate this, a question can

be included like: “How much time do you spend on private things during working hours: ‘< 15 minutes’, ‘15-30 minutes’ or ‘> 30 minutes’?”. To get an indication of employee happiness, a second question could be included like: “On a scale from 1 to 10, indicate how satisfied you are with your current job?”. In this example the properties ‘private business’ and ‘employee happiness’ are operationalized, i.e. made measurable.

Of course, there are many problems with operationalizing. The biggest problem is one that haunts any kind of translation: you are always loosing information, and you must *choose* which information to retain. With the first survey question you are obviously loosing information, because the scale of the survey answer differs from the scale of the property: the survey answers are on ordinal scale, while the actual property (time spent on private business) is on ratio scale. A second problem with operationalizing properties is illustrated by the second question. This question actually changes the property type: employee happiness is a qualitative property, but the answers of the survey are quantitative.

All this is basic knowledge from the field of statistics. How does it relate to MAS? Every part of the real world that is simulated undergoes a transformation resembling the abovementioned operationalization. Quantitative properties are easiest to translate to informatics: they can be represented by some numeric variable. On the other hand, qualitative properties pose a few problems. First, a representation is needed. Second, the translation is only valuable if the representations correspond to the real-world property in a natural way.

The first problem (representation) is addressed by a wide array of techniques. The most ‘straightforward fashion’ to translate a qualitative value is by representing it with a numeric variable anyway. This is done in the above survey example: ‘employee happiness’ is represented by a variable on a scale from 1 to 10. In a simulation of the business in the above example, ‘employee happiness’ could likewise be represented by a numeric variable. The second problem can now be reduced to the problem of creating a natural scale for the variable. The scale in the example seems ‘reasonable’, but there’s no *definite* solution for this problem. Finding a natural correspondence between a property and the variable will remain tricky business. To summarize, this approach can be posed as follows:

Definition: Translating in a ‘straightforward fashion’
Representing real-world organization properties (quantitative as well as qualitative) in a simulation by numeric variables.

Many other techniques (including for example logic) for representing real-world properties are available. These techniques themselves form a whole field of research, and thus fall outside the scope of this thesis as a subject matter. I will be using one of these methods in chapter 4, though: the ‘straightforward fashion’ of translating properties to variables in a simulation will be the approach I take in this thesis. In the next section I will discuss how this affects the answers to the research questions.

3.1.2 Research Questions

This section is about rephrasing the answers to the research question given in section 2.3.2, such that they will fit in the context of MAS. To do this I will use the 'straightforward fashion' of translating terms, posed in the previous section. Let's recapitulate the answer to the main research question given at the end of the previous chapter:

> **Preliminary answer: question 1**
> A specification of how we can measure the 'distance' between the current properties of the organization and the desired properties of the organization.

This answers the question how we can evaluate an organization. The question gets a different form if we're building a MABS. In that case, we are interested in the *simulated* organization. This also means we want to measure the 'distance' between *simulated* properties: the variables representing those properties. As I proposed in the previous section, these variables will be some numeric representation of the real-world properties. Measuring the 'distance' between two different values of the same variable then becomes a simple piece of algebra!

So, the question was what the criteria are to evaluate the state of affairs in an organization? In the context of MABS, we can extend the preliminary answer with a specification of how to measure the 'distance'. In other words, when we commit ourselves to the 'straightforward fashion' of translating properties of an organization, we get the following answer to the research question:

> **Answer: question 1 ('straightforward fashion')**
> The criteria by which the state of affairs of an organization can be evaluated are the distances between the numeric variables representing properties of the current- and desired state of affairs.

Note that this answer corresponds to the 'straightforward fashion' of representing the organization in a simulation. If you don't want to commit to a certain approach of translating reality to a simulation, you are stuck with the preliminary answer. To put it this way: the specification of how to measure 'distance' between different values of a state of affairs' property is completely dependent on the way you represent those properties.

So the criteria are dependent on the way organizational properties are represented. Note however, that at the end of the previous chapter I concluded that the specific criteria are dependent on the domain of the organization. Though this seems to be contradictory, it makes perfect sense: if you are trying to be useful while having fun[4], the choice of the 'translating fashion' depends on the domain. For example, if you are simulating an economic situation (as I will be demonstrating in chapter 4), usage of the 'straightforward fashion' is very clear, because the properties of the organization are for the largest part already represented as numeric variables. The dependences between various concepts are demonstrated in Figure 2.

---

[4] This refers to [12] (see also section 1.1): we need to choose our methods and techniques in a way that they will actually contribute something useful to science.

Figure 2

In fact, figure graphically answers the first part of the research sub question "What is the level of domain-dependence of these criteria?": it is very domain dependent. This answer, as well as the answer to the second part of this sub question, essentially remains unchanged since chapter 2:

Answer: question 1.A
The specific criteria (specifications of how to measure 'distance' between different values of a state of affairs' property) are completely dependent on the domain of the organization. These criteria can be found by analyzing the domain of the organization.

To summarize, I have posed answers to the main research question and the first sub question. The rest of this chapter will be about finding an answer to the second sub question, and also a way to use and test the answers.

3.2 Methodology

In the previous section I have shown what the criteria for evaluating the state of affairs of an organization are, and how we can get specific instances of these criteria by analyzing the organization domain. Now suppose that you need to build a specific MABS: how can answers to the research questions be helpful? In other words: how can these criteria facilitate formal models that include dynamic reorganization of agent societies (research sub question 1.B)? This sub question is the subject of this section.

To answer this question, we must consider two steps necessary for creating a formal model. First, we must analyze the domain of the organization we want to simulate with the MABS. Second, we must build the model specifying the environment, agents, simulated organization, and so on. Several methodologies exist for taking these steps, including GAIA [20] and OperA [3].

The answers to the first two research questions can facilitate formal models built with GAIA or OperA, because they tell us that models must include specifications of how to measure 'distance' between different values of the organization's properties. These specifications then enable us to view the organization from several different perspectives (corresponding to different organizational goals). Such a perspective is then defined as follows:

Definition: perspective
A view of an organization by looking at a certain combination of one or more 'distances' between two values of organizational properties.

This enables us to include dynamic reorganization in our simulation. Let's recapitulate the second sub question: "How can these criteria facilitate formal models that allow the specification of dynamic reorganization of agent societies?”. This question can now be formally answered in the following way:

Answer: question 1.B
Specifications of ways to measure ‘distance’ between different values of organizational properties can be included in (formal) models used as a basis for MABS, because they offer us different perspectives. The organization can then be dynamically reorganized by using multiple perspectives on the organization.

Formal models already allowed dynamic reorganization, as was demonstrated in [4]. However, the simulation in [4] only used one perspective that only used one ‘distance’ measure on the property ‘survival rate’. I would like to contribute to this is then the knowledge that you need to include multiple perspectives in realistic simulations. I would like to do this because organizations virtually always have multiple goals (see section 2.3.1).

Methodologies for specifying formal models used as a basis for simulations must then allow us to specify these perspectives. As a final note on methodologies, it is important to realize that any methodology must include some mechanisms for specifying organizations and organizational goals. As mentioned in [3], not all current MAS-methodologies have these mechanisms. When building a MABS that will include multiple perspectives on the organization, either you need to extend your methodology with these mechanisms (as would be the case when using -for example- GAIA [20]), or you need to use a methodology that already has these mechanisms (as would be the case when using -for example- OperA [3]).

3.3 MABS and Decision-Support Tools

We have gathered information from OT, and from this we know how to build criteria to evaluate the state of affairs of an organization. We also know how this facilitates the formal models we’ll use as a basis for building a MABS. But, as I promised in the first chapter, we need to be useful while having fun finding answers to the research questions. So before jumping to conclusions, we must first see if the information from OT and its translation to informatics can be useful. In this section I explore where ‘perspectives’ on simulated organizations are useful, and where they aren’t.

Generally, MABS are very useful for building decision-support tools for managers. Such a tool could for example include a simulation of a specific organization. The manager of that organization could use the tool to test and train his ability to dynamically (re)organize his organization.

For example
Consider once more the example of the pharaoh from the first chapter, who wants to build a pyramid before he’s a late pharaoh, before he’s a stiff[5]. However, it must also become the most beautiful pyramid ever built. For this he needs slaves doing the dirty work (lifting, dragging, etc), and artists. Suppose that the pharaoh

[5] “… before he’s a late pharaoh, before he’s a stiff …” refers to the fabulous dead-parrot sketch (and probably the most famous sketch) by Monty Python.

has a computer and he instructs one slave to build a decision-support tool. The tool will include a simulation of the different groups of slaves building the pyramid, the environment, resources and so on: a MABS. The pharaoh can then (re)organize his organization in the simulation and find out what kind of risks come with certain types of (re)organizations. This *supports* his real-life decision-making process.

For the slave building the decision support tool with a MABS in the above example, it's important to use a methodology that offers a way of specifying an organization and multiple perspectives. In fact, this holds not only for the above example, but also for any decision-support tool with a simulated organization. Because the organization can be looked at by the manager from multiple perspectives, the simulated organization in the decision-support tool must also offer these perspectives to the manager training his (re)organizing abilities.

Multiple perspectives using specifications of measuring 'distance' between values of a certain property aren't always useful. Although most organizations have multiple goals, not all organizations do. Or perhaps sometimes you would like to focus on one particular goal. This is the case if for example you are building a decision support tool solely directed at training managers to achieve one goal, for example making profit.

Also, remember from that information on organizations in the previous chapter was drawn from OT. This field of research is about human organizations, which are artefacts created with some purpose [16]. The thesis presented here can therefore not be used when working with other kinds of organization; at least not without (considering) modifications.

3.4 Towards a Simulation

I have claimed information from OT can be translated to informatics and that the resulting answers to research questions are useful in building MABS, which can serve as the basis for decision-support tools. Such a claim would be hollow without some kind of proof. To show that multiple perspectives on an organization can play an important role in a MABS, I will construct a simulation. Experiments with this simulation will then show the implication of including multiple perspectives. Before going on to the actual simulation in the next chapter, I will discuss limitations of the bridge between the research questions and such a simulation.

First notice a limitation already mentioned in the previous section: you can only simulate human organizations, as the answers to the research questions were constructed using OT. With some extra work, perspectives in one form or another can contribute to non-human organizations, such as for example artificial life simulations. However, this research is left to future work.

Second, the bridge (between the research questions and a simulation) itself has some limitations. The simulation in the next chapter will only demonstrate how the answers to the research questions *can* be used, not how they *must* be used.

In between lies the methodology. How this methodology is used to implement multiple perspectives is left implicit in this thesis. The simulation presented in chapter 4 will therefore not build the bridge in the *best* possible way. The best possible way can only be found by researching additions and methods for existing methodologies that explicitly state how multiple perspectives (from which to view organizations) can be implemented.

I wish to conclude with what good a simulation then *would do*. While only being about *human* organizations, a simulation can demonstrate how multiple perspectives *might* be used to make improve MABS. I will show this in the next chapter.

4. Simulation

4.1 Simulation Introduction

In this chapter I will show how the ideas from the previous chapters can be used to improve simulations and decision support tools. A small but realistic scenario will be simulated. It is then shown by a series of experiments how reorganization might require us to evaluate the state of affairs of the organization from different perspectives.

The simulation will be implemented as a small discrete state machine. A typical run consists of several steps that transform the current state into a new state. This can be represented by the following pseudo-code:

Pseudo-code simulation run

1. set up model
2. WHILE not bankrupt DO
    a. let the employee agents generate capacity
    b. calculate new values of organizational properties
    c. let the manager reorganize the organization

A continuous setting can be approximated by a discrete setting [19]. This is exactly what I will do here.

4.1.1 Scenario

Let's start with the scenario. We are looking at a store in a small town. This store provides many products for the people in the town, but also for people living around the area. The exact type of product is not really important, and left to the reader's imagination. The amount of people in town, and the way demand for products comes from the area or from town, and so on are variable parameters which can be set different for every simulation (an overview of all variable parameters is given later this chapter, as well as in appendix B).

The amount of people in town, as well as the amount of people that's employed somewhere else is kept constant. What might change however is the amount of people employed by the store. New employees can be hired, and current employees can get fired. As a consequence, the amount of unemployed people depends totally on the amount of people that work at the store. Part of the human employees at the start of a simulation run can be 'old' employees: people that have worked at the store for a long time.

People in town will spend part of their money at the store, which is possibly less if they are unemployed. The total of spending by townsfolk is first part of the demand. The second part of the demand comes from the area surrounding the town, and follows a certain economic cycle (with some randomized variations). The exact cycle, as well as the proportion of demand from inside or outside of town can be set different for each simulation.

So, the store gets a certain demand for products[6]. The store also has a certain amount of costs, composed of costs for the products made, costs for the employees and constant costs (which steadily increase over time). At the end of each timestep (for example a month or a year), the amount of profit or loss equals the total demand minus the total costs. For example if total demand is 1500 and total costs are 1200 profit is 300.

To be able to satisfy the total demand, the store has employees. Each employee generates some 'capacity' for the store. The total capacity is the maximum demand the store can satisfy. For example, if the store has 10 employees each generating capacity 110, total capacity is 1100. If the capacity is below the demand, the demand will actually be cut off to match the capacity. If the above two examples are combined, demand would shrink to 1100, as there are only enough employees to sell that much products. This would result in a netto result of 1100 minus 1200: a loss of -100.

Next to human employees, the store might also have machines to generate capacity. In this scenario machines are also considered as employees generating capacity. Machines may have a cost/capacity ratio different from humans, perhaps one that's more profitable. In any case, the ratio between machine- and human employees is fixed (per simulation), as only a certain maximum amount of machines can be operated with a given amount of human employees. For example if the machine/human ratio is 0.3, then with 10 human employees you can hire 3 machines, with 20 humans you can hire 6 machines, and so on.

The total amount of assets the store has is called the bank value. The netto result is added to this value at the end of each timestep. If the bank value ever reaches below a minimum (which can be set for each simulation) the store goes bankrupt.

To summarize, this scenario encompasses a small economic situation with a store selling products, and hiring employees and machines. Demand for products comes from people in town (including employees), and also from outside town (following a simple economic cycle). With this scenario, we can now continue to set some objectives.

4.1.2 Objectives

The scenario is perfect for studying reorganization. To be able to focus on this, we introduce a manager, called the HEAD agent. Depending on the economic situation and the details of the scenario, the manager can decide to hire and fire

[6] Note that all variables and parameters such as demand, cost, capacity, etc. are expressed in an amount of money, never in an amount of products.

both humans and machines[7]. For example, to take advantage of increasing demand, the manager needs to increase capacity correspondingly by hiring machines and/or humans. However, this specific type of reorganization was already studied in depth by [4]. In that simulation (called VILLA), the effects of changes in society typology on the organizational goal (surviving) were studied. As promised, the simulation presented here implements ideas from the previous chapters: we will need to look at the organization from different perspectives.

To incorporate these different perspectives we need to extend the scenario. Remember that in [4], the only perspective is formed by the single utility function indicating the organization's survival rate. In the scenario presented here, we have a similar perspective formed by the organizational goal of making profit. Because the scenario is quite realistic, it is also quite easy to find different perspectives. Though many organizational goals (as for example keeping profit stable over a certain period of time) could form new perspectives, only one such perspective is implemented in this simulation[8]. The new perspective is added as an organizational goal. The organizational goals are now:

1. Make Profit.
2. Employ Townsfolk.

The rationale behind the second goal is that the store has some social responsibility towards townsfolk. Either goal might require different reorganization strategies considering different settings.

To summarize, the different organizational goals form different perspectives from which reorganization decisions can be made. We will look more closely at reorganization strategies when discussing the experiments in section 4.3. But first I will specify how the scenario will be implemented as a simulation.

4.2 Simulation Specification

In this section, the simulation details will be laid out. The simulation consists of a model and agents. The model represents the environment. This environment is specified by parameters, sets of agents occupying the environment, and two types of functions. The parameters determine the values for a range of items, from the amount of machines available to the starting assets of the store. The model also specifies what agents are in the environment. Last, two types of functions are implemented. There are functions for determining the value of 'parameters' whose values can be calculated using other parameter values. Then there are so-called 'timestep

[7] For convenience, and to keep the text from becoming clottered with too many terms, I use "hiring/firing" machines instead of "buying/selling".

[8] The scenario here allows us to study *multiple* perspectives when considering reorganization. This is studied most easily with only two perspectives,. When appropriate, many more perspectives can be incorporated in a fashion similar to the approach taken with this simulation.

functions', which are calculated once at the end of each timestep, to determine the new value of certain variable parameters. These three parts form the model, which will be explained first.

After I've done the environment, the agents are specified. The human- and machine agent behavior is specified right after the model, a separate section is dedicated to the HEAD agent later this chapter.

The simulation was implemented using the Repast[9] modeling toolkit. The version used for this simulation is Repast J, version 3.1. Repast is freely available and can be downloaded from [22]. The simulation is available from [23].

4.2.1 Simulation Model Parameters

The economic environment as described earlier this chapter is implemented as a series of variables, also called parameters. The user can set these parameters before the simulation is run. Each set of parameters will result in a different run, with a different outcome. The parameters are grouped together in semantic groups[10]. Note that parameters with a * are set for the first timestep, but they might change over time. The other parameters remain constant during a run.

Amounts of Agents

amount of human employees*	$SIZE_T(H)$
amount of machine employees*	$SIZE_T(M)$
amount of humans in town	$SIZE_T(H_\Sigma)$
amount of machines in town	$SIZE_T(M_\Sigma)$
amount of humans employed elsewhere	$SIZE_T(H_{OTHER})$
amount of steady ('old') human employees	$SIZE_T(H_{OLD})$

Costs & Capacity

capacity per human	CAP(H)
capacity per machine	CAP(M)
cost/timestep per human	COST(H)
cost/timestep per machine	COST(M)
constant cost/timestep	COST(CONST)
constand cost trend	C_1

Proportions

maximum proportion of machines	$P_{MAX}(M)$
proportion of product cost	P(PRODUCTS)
proportion of demand from town	P(INT)
proportion of demand from outside town	P(EXT)

---

9 Recursive Porus Agent Simulation Toolkit.

10 This grouping of parameters is only for reading convenience, in the simulation they are all handled alike.

Bank Values

minimum bank value	$BANK_{MIN}$
bank value at a certain time*	$BANK_T$

Demand constants

demand/timestep per human with job	D_{JOB}
demand/timestep per human without job	$D_{NO\text{-}JOB}$
external-demand constant (base demand)	D_1
external-demand constant (amplitude)	D_2
external-demand constant (frequency)	D_3
external-demand constant (trend)	D_4
margin for random variations	D_5

All these parameters are described in relation to the actual simulation software in appendix A.

4.2.2 Simulation Model Functions

At any timestep, we have a set of parameters. Using the values of these parameters, we can calculate the value of certain variables that will be used in the timestep functions (next section). The value of these variables can be calculated with the following numeric functions.

Variable Calculations

$SIZE_T(H_\Sigma)$ - $SIZE_T(H \,\&\, H_{OTHER})$	$= SIZE_T(H_{NO\text{-}JOB})$
$CAP(H) \cdot SIZE_T(H) + CAP(M) \cdot SIZE_T(M)$	$= CAP_{\Sigma,T}$
$COST(H) \cdot SIZE_T(H) + COST(M) \cdot SIZE_T(M)$	$= COST_{E,T}$
$P(PRODUCTS) \cdot CAP_{\Sigma,T}$	$= COST_{PRODUCTS,T}$
$COST_{E,T} + COST_{PRODUCTS,T} + COST_{CONST}$	$= COST_{\Sigma,T}$
$SIZE(H_{OTHER} \,\&\, H) \cdot D_{JOB} + SIZE(H_{NO\text{-}JOB}) \cdot D_{NO\text{-}JOB}$	$= D_{INT,T}$
$D_1 + D_2 \cdot SIN(D_3 \cdot T) + D_4 \cdot T$	$= D_{EXT,T}$
$P(INT) \cdot D_{INT,T} + P(EXT) \cdot D_{EXT,T}$	$= D_{\Sigma,T}$
IF $(D_{\Sigma,T} > CAP_{\Sigma,T})$ THEN $CAP_{\Sigma,T} - COST_{\Sigma,T}$	$= R_T$
IF $(D_{\Sigma,T} \leq CAP_{\Sigma,T})$ THEN $D_{\Sigma,T} - COST_{\Sigma,T}$	$= R_T$

4.2.3 Simulation Model Timestep Functions

At the end of every timestep the model is updated. This means the value of several parameters might change. First, several parameters change as a consequence of the environment. These are the environmental timestep functions. The second type is the agent timestep functions, which depend on the HEAD agent. The following items are updated every timestep.

Environment Timestep Functions

$BANK_T + R_T$	$= BANK_{T+1}$
IF $BANK_T < BANK_{MIN}$	THEN $BANKRUPT_T$

HEAD Agent Timestep Functions

$SIZE_T(H) + HIRE_T(H) - FIRE_T(H)$	$= SIZE(H)_{T+1}$
$SIZE_T(M) + HIRE_T(M) - FIRE_T(M)$	$= SIZE(M)_{T+1}$

4.2.4 Simulation Agents: Machines & Humans

Below are the agent specifications for humans and machines. The environment (i.e. the parameters and outcome of variable calculations) is completely visible to all agents. The actions of machine and human agents are deterministic given their local environment (employed or unemployed). This behavior is kept simple and constant, as I want to focus on the HEAD agent's behavior. Note however, that any kind of behavior could technically be implemented, for example to study the effects of strikes or shopping sprees during Christmas.

Let's first examine machine agents. These agents correspond to machinery that a store might use to do work that's otherwise done by humans. Their behavior is the simplest of all agents. It's a function of the local environment. When employed by the store, they will generate a certain capacity. When unemployed they won't do anything of importance. This is summed up in the following table.

Local Environment	**Action**
employed	generate capacity CAP(M) for the store
unemployed	DO NOTHING

Human agents correspond to humans that might be employed by the store. Their behavior has some extra complexity, but is still a function of the local environment. Human agents can have three different local environments. They can be employed by the store, employed elsewhere or unemployed. Depending on this state, they choose their actions.

Local Environment	**Action**
employed by store	generate capacity CAP(H) for the store demand D_{JOB} FROM STORE
employed elsewhere	DEMAND D_{JOB} FROM STORE
unemployed	DEMAND $D_{NO\text{-}JOB}$ FROM STORE

4.2.5 Simulation Agents: HEAD Agent

This agent is the main factor of interest, and he will need some reasoning to be able to make decisions. First, let's recall that two timestep functions are tailored for the HEAD agent: each timestep he must decide how many machines and humans are hired or fired. These values form the reorganization strategy, and determine how the society typology is changed. The HEAD can have several different reorganization strategies, to determine how to fire or hire employees.

To see what type of strategies might be relevant, we must first realize why the HEAD wants to hire or fire employees. The answer is that it's the perfect way to influence attributes determining how organizational goals are evaluated. In the following table, our two goals are linked to these attributes:

Goal	Attribute	Desired Effect
Make Profit	$BANK_T$	BANKT → ∞
Employ Townsfolk	$SIZE(H_{NO-JOB})$	$SIZE(H_{NO-JOB}) \rightarrow 0$

So, how are these goals evaluated? We need some scale and a function to put values of the above attributes on that scale. For a scale a range from 0 to 100 will be used. Then, the function that evaluates the above goals is defined:

Goal evaluation-function
EVAL(G) → [0 … 100], where G is the goal.

This definition is still very generic, and requires a more fine-grained definition for each possible input. The following definition shows how to get from the attributes to the scale:

Goal evaluation-function (goal 1 and 2)
If $BANKRUPT_T$ then
eval(1) = eval(2) = 0
else

$$eval(1) = ((r_{t} / d_{\Sigma,t}) + 0,5) * 100$$
$$eval(2) = size(h) / (size(h) + size(h_{no-job})) * 100$$

These functions transform an attribute and its range to a [0…100] scale, and have a rationale for the way this is done. The first goal (make profit) is achieved by maximizing the BANK value. This is done by making as much profit as possible. In economics, profit is often given as a percentage of the total sales. We use the same index, although a profit of 0% is lifted to the index of 50: this gives the range [0…50] for loss and [50…100] for profit.

Note that technically it's possible to get a loss greater than –50%, or a profit greater than +50%, which would both push the goal evaluation out of the [0…100] range. In the simulation these extreme values are therefore transformed back to 0 and 100 respectively. This theoretically limits the extend to which the perspectives (given by these evaluations of the goals) can be used. In simulations where you face the problem that values of goal evaluations cannot be *practically* presented on a finite scale a different way of representing the perspectives has to be found.

The second goal also transforms certain attributes to a scale of [0…100]. The goal will evaluate to 0 if there are no humans employed at the store. It reaches 100 when every single human available in town has been hired.

Both goals can be evaluated quantitatively (even on a finite scale), which is good news. In the previous chapter, it is claimed that the criteria for evaluating the state of affairs in an organization must be used in different perspectives. The qualitative evaluations of the organizational goals offer exactly these perspectives. The HEAD agent can use these perspectives to improve organizational structure with reorganizations triggered by criteria on these goal evaluations. A human manager could train his reorganization skills with these newly offered perspectives. Moreover, a software HEAD agent can be implemented to test whether usage of these perspectives has a desired effect. Such a software agent is subject of a series of experiments, which will be handled in the next section.

4.3 Simulation Experiments

The simulation was built to demonstrate that multiple organizational goals can offer different perspectives, which can help improve reorganization decisions of the HEAD agent. To test whether the simulation usefully implements multiple perspectives, a series of experiments with the simulation will be run. First, I will state the general idea of these experiments. Second, I give the actual parameters used to run the experiments. Finally, the reorganization strategies are given. In the next section (4.4) results of the experiments described here will be analysed.

4.3.1 General Idea

In all experiments the parameters describing the 'economic situation', (the amount of townsfolk, the capacity of humans and machines, and so on) remain constant. The parameters are chosen such that the experiments simulate a simple business cycle. The variable input of the experiments will be several different reorganization strategies for the HEAD agent. In the first series of experiments these strategies are fixed and don't use the extra perspectives that are offered. In the second series the goal evaluation data is used in the reorganization strategies. Any contrast between the first and the second series will indicate *whether* the extra perspectives are useful. The second series shows *how* these perspectives might be used in forming reorganization strategies.

Graphically, the flow of these experiments can be viewed as in Figure 3. The input is formed by the reorganization strategy and a fixed economic situation. With this information, the simulation is run and outputs economic data as well as goal evaluation data. This output, and more specifically the link between the two different types of output, will be analysed.

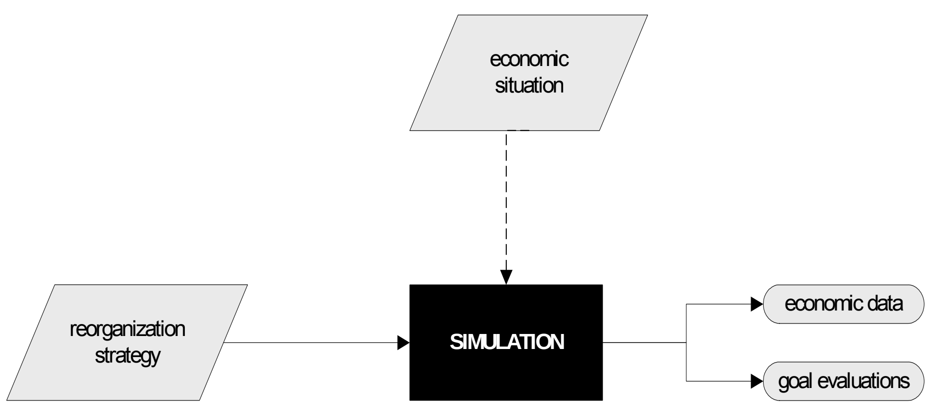

Figure 3

In the next section, I will describe in the parameters determining the economic situation. I go on to the reorganization strategies for the HEAD agent, and finally I describe the output of these experiments.

4.3.2 Parameters

The Repast environment offers many ways to do a batch of runs, where each run can have its own parameters. In the experiments the economic parameters stay the same. The following parameters were used:

Amounts of Agents

$SIZE_T(H)$	= 100
$SIZE_T(M)$	= 20
$SIZE_T(H_\Sigma)$	= 500
$SIZE_T(M_\Sigma)$	= 250
$SIZE_T(H_{OTHER})$	= 100
$SIZE_T(H_{OLD})$	= 75

Costs & Capacity

CAP(H)	= 100
CAP(M)	= 125
COST(H)	= 50
COST(M)	= 25
COST(CONST)	= 5300
C_1	= 0.5

Bank Values

$BANK_{MIN}$	= -250,000
$BANK_T$	= 0

Proportions

$P_{MAX}(M)$	= 0.5
P(PRODUCTS)	= 0.1
P(INT)	= 0.5
P(EXT)	= 0.5

Demand Constants

D_{JOB}	= 75
$D_{NO\text{-}JOB}$	= 20
D_1	= 12000
D_2	= 2000
D_3	= 0.01
D_4	= 2.0
D_5	= 0.15

For all six reorganization strategies, 100 runs are computed of each 2500 timesteps, using the above parameters. These strategies will be discussed in the following section.

4.3.3 Reorganization strategies

First Series

As stated before, a first series of experiments will be run to determine whether the goal evaluation data can be used in combination with reorganization strategies. In these experiments, several fixed strategies are tried, combining them with the economic situation. If the goal evaluation data is useful, a link between the chosen strategy and the goal evaluation data should become clear. The fixed strategies that will be used are the following:

1. **No Reorganization**: the situation is not changed during the experiment.
2. **Hire Cheapest**: when capacity must be lowered, the most expensive employees are fired first. When capacity must be increased, cheapest employees are hired first.
3. **Hire Humans First**: when must be increased, humans are hired first. When capacity must be lowered, machines are fired first.

Note that all these strategies talk about situations where capacity must be increased or decreased. This is because the simulation at hand imposes a constraint: for both goals it is essential the store doesn't go bankrupt. If capacity is too low, you miss out on a portion 0f the demand and will almost always go bankrupt. If the capacity is too high the cost (for employees) is also too high and you will go bankrupt as well. In other words: the capacity must be at level with the demand value at all times, regardless of the organizational goals.

At a set interval the difference between the demand and capacity will be taken as a basis for the reorganization. The above strategies will then be followed to level demand and the capacity. It is possible and probably even wise for a real HEAD agent to use additional measures next to this difference between demand and capacity to trigger reorganization. For example, a margin between capacity and demand could be implemented to create a more realistic situation. However, as this would make things even more complex, we would risk loosing an overview of the subject at hand: the link between economic data and the goal evaluation. Therefore, this kind of additional measures is left to follow-up research.

Second Series

In the second series of experiments, we will determine how the goal evaluation data can be used in reorganization strategies. The HEAD agent will now use this goal evaluation data at runtime in his reorganization decisions. We follow the same constraint as with the first series, stating that we always try to keep the capacity at level with the demand, as to avoid bankruptcy.

Now then, the strategies 2 and 3 from the first series can be combined: depending on the goal evaluations, we choose to follow one of these strategies. Obviously, these reorganization strategies correspond directly with respectively organizational goal 1 (Make Profit) and 2 (Employ Townsfolk).

To choose between these reorganization strategies, the HEAD agent must use his two perspectives to assess the priority of achieving either goal. After that the HEAD agent can choose a strategy based on the goal with the greatest priority. This main reasoning structure can be viewed graphically as in Figure 4.

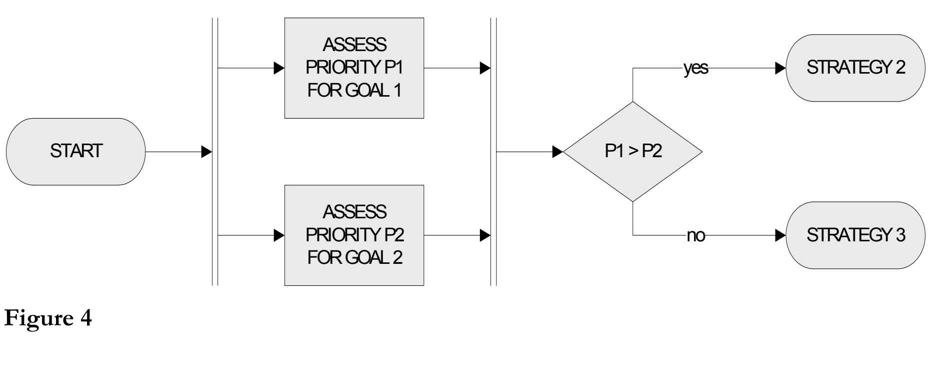

Figure 4

The (P1 > P2) decision is discrete: it either assigns priority to goal/strategy 1 or 2. This is for simplicity's sake. Obviously future work should be directed at using continuous values for P1 and P2, and perhaps also at using strategies in a weighted sense, i.e. do a little of both strategies. For example, if P1 and P2 are both 0.5, you could decide to increase capacity 50% by using the cheapest agents, and 50% by using humans. Note that this assessment can also be done by a human, for example in the case where a human uses this simulation as a training platform for his reorganization skills[11].

All that's left now is a way to specify how the HEAD agent assesses priority for both goals in the experiments. In these experiments, we use a software agent. Such a software agent can assess priority using several triggers. In the following table some possible triggers are given.

#	Priority	Condition	Rationale
1	1 > 2	-	No ratio: blindly follow goal 1.
2	2 > 1	-	No ratio: blindly follow goal 2.
3	1 > 2	BANK < 0	Prevent the store from running on borrowed money.
4	2 > 1	BANK > 0	Employing humans is important if the store has a positive amount of money.
5	2 > 1	SIZE(H) < C	Try to keep at least a constant number of human employees (for example employees who have worked there for a long time).
6	1 > 2	ΔEVAL(1) < 0 & ΔEVAL(2) > 0	If goal 1 is dropping and goal 2 is rising, then give goal 1 priority.
7	1 > 2	ΔEVAL(1) > 0 & ΔEVAL(2) < 0	If goal 1 is rising and goal 2 is dropping, then give goal 2 priority.
8	1 > 2	CAP(M) / COST(M) > 2 * CAP(H) / COST(M)	If the capacity/cost ratio of machines is twice that of humans, goal 1 is given priority.
9	2 > 1	SIZE(H) < SIZE(H_{OTHER})	The store tries to employ at least as many people as there are people employed elsewhere.
10	1 > 2 2 > 1	ODD OR EVEN REORGANIZATION	The HEAD agent will give priority to the goals in an alternating fashion.

Note that triggers 1 and 2 can function as default triggers, as they have no precondition: this is because the can always be applied.

[11] Note that a human *might* use the perspectives offered by the goal evaluations, but that he can also use other resources, for example intuïtion.

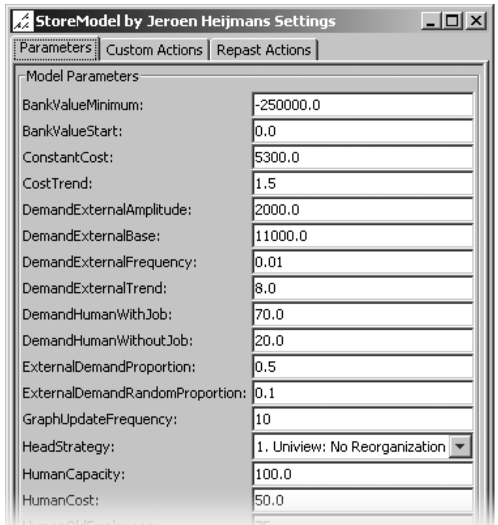

Figure 5

Any trigger from this (by far not even exhaustive) list could in theory be implemented. In these experiments I will only use some of the mentioned triggers. The following two 'flexible' strategies evaluate the situation from different perspectives, then goes through a series of triggers until a reorganization decision can be made.

1. Trigger 6 → Trigger 7 → Trigger 1. ('idealist' strategy)

2. Trigger 5 → Trigger 3 → Trigger 2. ('conservative' strategy)

3. Trigger 10 ('fair' strategy)

To summarize, we now have three different fixed strategies, and three different flexible strategies. Both types of strategies will be run with the same, fixed economic situation. In the next section I describe the results of these experiments.

4.4 Simulation Experiments, Results

4.4.1 Typical Outcome

First I will discuss how a typical, single simulation run goes and what the input and output are. The Repast environment provides 2 main mechanisms for inputting model parameters. First, you can use a parameter file. This is the basis for doing a batch of runs. Second, you can input the variables at runtime in a window such like Figure 5.

Once the parameters are selected for the model of the current run, the simulation is run. The program runs through timesteps, adjusting parameters (such as the bank value) as it goes. The important (changing) attributes are directly shown (at runtime) to the user in graphs. This would look something like in Figure 6.

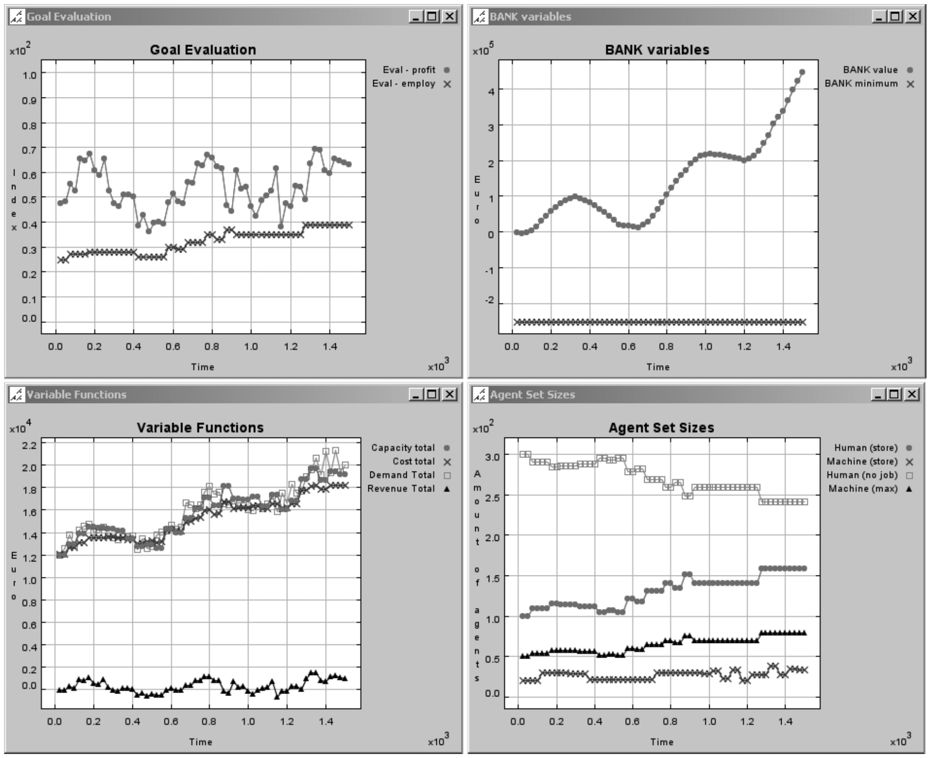

Figure 6

Let's look at these four graphs[12]. All of them plot some parameters over time (measured in timesteps). Top-left the graph titled "Goal Evaluation" shows two lines, one for each goal. The evaluation of this goal is plotted on a scale [0…100]. Top-right the graph "BANK variables" displays how the bank value evolves over time. It also shows a line at the bottom indicating the minimum bank value. If the first line ever drops to the bottom line, the store has gone bankrupt.

On the second row, down-left there's a graph called "Variable Functions". This graph shows the interesting economic variables over time. The line at the bottom that's always around 0 is the store's profit. The three lines at the top are the total demand, total costs and the total capacity. Down-right, the last graph "Agent Set Sizes" displays how employees and other folks are hired and fired. From top to bottom, the particular graph in figure shows the humans without a job, humans working at the store, the maximum amount of machines, and finally the actual amount of machines employed by the store.

The four graphs in the figure form the typical output for a simulation with the parameters for the experiments, using the 'multiview fair' strategy. Starting with the bottom row, we can see that the manager is doing a reasonable job in keeping capacity alongside the demand, thus making profit most of the time. The bottom-right graph shows that the manager looks at his organization from

[12] A short explanation of the graphs is given here. Refer to appendix B for a more detailed, in-depth, 'tutorial-like' description of these graphs.

one or the other perspective. Sometimes he hires (cheap) machines, but sometimes he also hires (expensive) humans. Top-right we can see again that the store is making profit most of the time, as the bank value keeps increasing. Finally, the top-left graph displays how the goals are evaluated, and offers us a practical version of the different perspectives on the organization. The average value of the evaluation line represents how how well a run was, considered from a certain perspectives. In the next section I will discuss how these averages come up if we run a typical simulation in a batch, one batch for each strategy.

4.4.2 Results of Reorganizations

The simulation discussed in the previous section was done in a batch of 600 runs, 100 for each strategy. For each of these 600 cases, the average for both evaluation functions was measured. This forms the basis for the boxplots in Figure 7 and Figure 8.

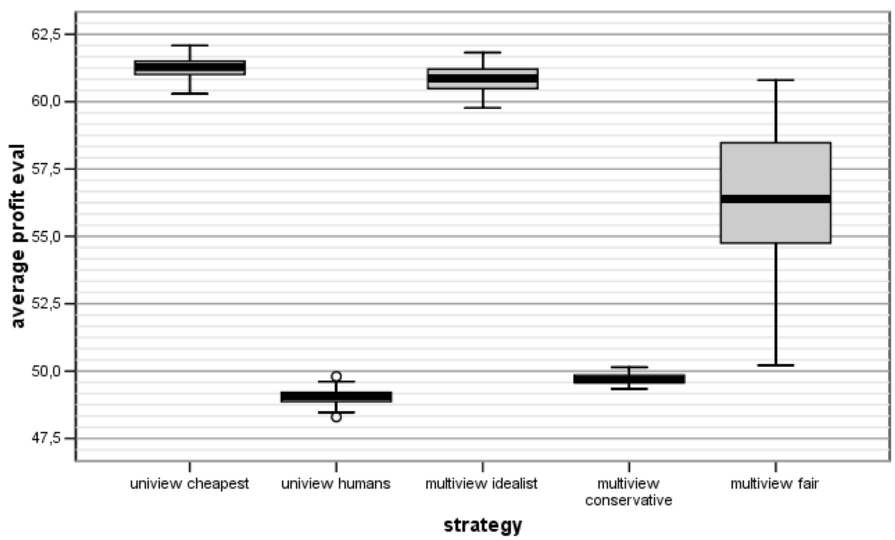

Figure 7

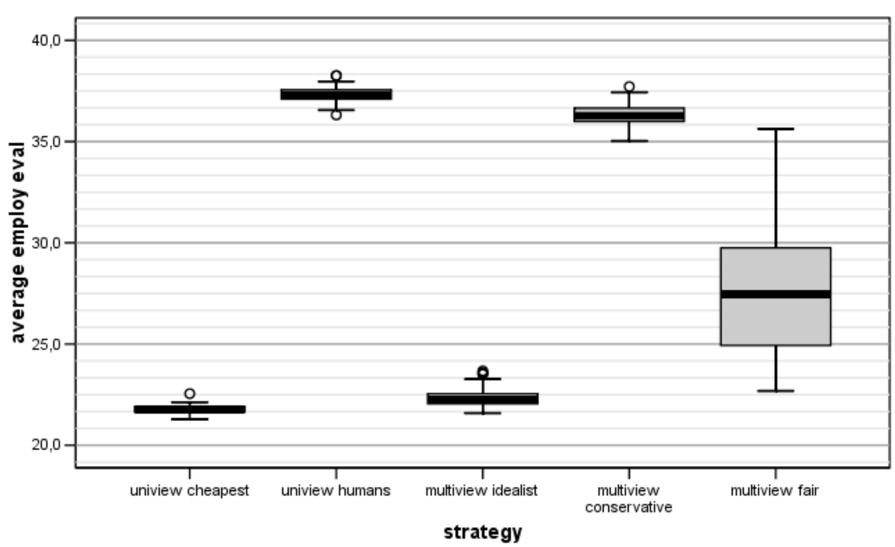

Figure 8

First, notice that one strategy is missing: the very first strategy, that doesn't reorganize at all. As was to be expected, this strategy went bankrupt every single run. This is because the cost and demand follow a trend, but the capacity doesn't because the strategy does not reorganize anything. Although it was therefore not useful to include it in the above graphs, it does allow me to state an obvious result:

Result: no reorganization
Reorganizing nothing at all is not really an option in dynamic open environments such as the one from this economic simulation.

Moving on to the boxplots with the other five strategies, we can immediately make some observations. First, both graphs have a similar format: the average values for the strategies are divided in two high averages, two low averages and one mediocre average. Second, we can immediately see that in both graphs four boxes are very condense and one box is very stretched.

So how can we explain these two observations? Let's start with the first: the difference in box positions. The two *uniview* strategies have the position we would expect. The first strategy focuses on making profit, the second strategy on hiring humans. Both strategies excel very well in their aim, but fail with the other goal (relative to the other strategies). Surprisingly, the first two *multiview* boxplots follow the same pattern as the uniview strategies: the 'idealist' and 'conservative' strategies correspond to the 'cheapest' and 'humans' strategies. This is an unexpected result, as triggers for assessing priority for goals were included in these strategies, to avoid having a 'weakness' (to excel in one goal but fail in accomplishing the other goal). Upon closer inspection, it turns out that this can be explained by the fact that the triggers apparently almost always assess priority to one particular goal. In the form of a result:

Result: uniview and multiview
The uniview strategies behave as expected and excel in accomplishing their goal. The first two multiview strategies can be reduced to the uniview strategies, as they assess priority to goals very one-sidedly.

Finally, the last strategy breaks with the pattern of very high and very low averages. The average evaluation of both goals with the 'fair' strategy falls in between the two high, and two low boxes. This strategy *does* succeed in assessing priority to both goals, which is obvious because it follows directly from the trigger used: alternate priority between the two goals. This can be put as a result:

Result: 'fair' finds the "middle"
The 'fair' strategy succeeds in offering both perspectives a chance, resulting in average scores in between the high and low scores by the uniview strategies.

Second, we observed the first four boxes are very dense, while the last is very stretched. This can be explained by the fact that the 'fair' strategy is more sensitive to the randomness of the demand value. The value of demand can sometimes rise quite a bit due to the random factor. The first four strategies will always respond in a similar manner, under similar conditions. With the 'fair' strategy, these conditions don't have any influence: the combination of a demand spike and a certain goal-priority is also random. This increase in randomness allows the goals to be achieved with various degrees of success.

Result: randomness increased
The 'fair' strategy amplifies the random factor in the demand value, creating more spread out goal evaluations over different runs.

To summarize, we can draw three conclusions from the above experiments. First, the two uniview strategies behave as expected, and excel in the goal they aim to accomplish. Second, the 'idealist' and 'conservative' strategies fail to use the perspectives offered by the goal evaluations, and seem to be reducable to the two uniview strategies (considering this particular economic setting). Finally, we see that the 'fair' strategy uses both perspectives in an alternating way, thus finding the "middle" between the uniview strategies.

4.4.3 Research Questions

At the end of this chapter, let us see how the results from the simulation experiments relate to the research questions. Also will I look at how this relates to the three different types of organizations.

Let's go back once more to the main research question, about the criteria for evaluating an organizational state of affairs. The abstract answer in chapter 2 showed that we needed a "specification of how to measure 'distance' between organizational properties". In the simulation we have a concrete answer to this question. The goal evaluation functions offer us two different specifications of how to measure 'distance' between organizational properties.

In chapter 3 I claim in an answer to sub question 1.A that the actual criteria are completely domain dependent, and can be found by analyzing the domain. The simulation presented in this chapter implements multiple perspectives in the form of goal evaluation functions. These functions are indeed constructed based directly on the nature of the economic properties of this specific organization. In other words: this simulation confirms that the specific criteria of how to measure 'distance' *can* be found using domain analysis. This shows that the answer to the main question is valid, and that it offers a solution at most one step away from actual criteria: the step of anlyzing the domain.

Now, we would like to know what we can learn from these results about different types of organizations. Remember from chapter 1, that we have natural, simulated and MAS organizations. The experiment results show us how a simulation corresponding to a natural organization can be set up using a MAS. This was done with a small but realistic economic setting with an economic cycle. Inside the simulation we have learned that reorganization based on multiple perspectives can be a success, though this success depends on the reorganization strategy. Consequently, usage of multiple perspectives strengthens the link between natural organizations and simulated organizations. Because the simulation uses a MAS, it is to be expected that the link between natural organizations and MAS organizations would also strengthen, if multiple perspectives are included in both types of organizations.

As a final note it can be said that the experiments were a success, confirming that the answers to the research questions posed some usable insights. Even so, looking at the two multiview strategies that failed in their aim, it must be said that developing actual criteria (specifications of how to measure 'distance') is not an process to be taken lightly. When building simulated and MAS organizations, careful thought will have to be given to finding a metric for measuring the distance between states of affairs.

5. Conclusion

In this chapter I will summarize the research questions and the search for answers. After that I will pose a final conclusion. I will finish this chapter and also this thesis with some thoughts for future work.

5.1 Results

5.1.1 Research Questions

The research questions have stimulated a fruitful search for answers in various disciplines of science. Let's start with recapitulating these questions from chapter 1:

Research Questions:

1. What are the criteria for evaluating the state of affairs in an organization?

   A. What is the level of domain-dependence of these criteria? How do we get these criteria?

   B. How can these criteria facilitate formal models that allow the specification of dynamic reorganization of agent societies?

The search for answers directed us to Organization Theory and informatics, in respectively chapters 2 and 3. Insights from these disciplines combined allowed me to construct the following answers to the questions:

Research Questions, answers:

1. A specification of how we can measure the 'distance' between the current properties of the organization and the desired properties of the organization.

   A. The specific criteria (specifications of how to measure 'distance' between different values of a state of affairs' property) are completely dependent on the domain of the organization. These criteria can be found by analyzing the domain of the organization.

   B. Specifications of ways to measure 'distance' between different values of organizational properties can be included in (formal) models used as a basis for MABS, because they offer us different perspectives. The organization can then be dynamically reorganized by using multiple perspectives on the organization.

These answers are not the whole story. There's two more things left to discuss. First, the answers must be useful. To test this, a simulation was built to do experiments with. In the next section (5.1.2) I discuss the experiment results.

Second, we must try and learn something about MAS from the answers. This is the subject of the final conclusions section (5.3).

5.1.2 Simulation Experiments

The simulation described in chapter 4 posed a small economic setting with a store as the organization. For this imaginary natural organization, a simulated version was built. This was done in order to combine the answers to the research questions with reorganization, which is very important (as mentioned in chapter 1) in dynamic, open environments.

The answers from chapters 2 and 3 were followed by the concept of perspectives: different ways of looking at an organization when considering a reorganization. These perspectives were the subject of the experiments. In the experiments, six reorganization strategies were tested, all using different triggers. Three 'uniview' strategies that didn't use the added value of extra perspectives, and three 'multiview' strategies that tried to use all perspectives.

From the experiments it can be learned that the criteria (specifications of how to measure 'distance') could technically be implemented. This was done in the form of different evaluations for multiple different organizational goals. The strategies were there to test whether these implementations were useful for a manager (in this case: a software agent). In the results, it turned out that using the perspectives was quite tough. Two multiview strategies were reduced to the uniview strategies looking in only one perspective. However, a final multiview strategy forced the agent to look at both perspectives, resulting in average scores for both goals.

In the end, it can be concluded that the multiple perspectives using different criteria for evaluating the state of affairs in an organization *can* be a useful addition when trying to reorganize. This finally means that the answers to the research questions *are* useful in researching reorganization of organizations.

5.3 Conclusions

The most important part comes right here at the end of this thesis, and is about what we can *learn* from the answers to the research questions. How does this knowledge help us move towards MAS reorganization theory?

As mentioned in the introduction, the steps are as follows. First we examine human organizations and their theories. From this we try to extract some general principles. Those principles can then be tested in a simulation, and finally be made useful in the context of MAS. In this thesis, I followed these steps to learn something about reorganization. I have found that criteria used in reorganization decisions come down to some kind of measurement upon the organizational properties. To test this, I implemented a form of this measurement in a simulation, showing that these general principles on evaluation criteria can be somehow applied. This finally brought us with these general principles, which might be applied to real MAS in future work.

5.4 Future research

We have learned some things in this thesis. We have not learned everything there is to learn though. As always, new research questions have to be posed to learn more (though I doubt we will *ever* learn *everything*). Specifically, from this thesis several new research questions follow.

First, more knowledge from OT can probably be used to learn about simulated and MAS organizations. It would be good to start with an overview of which parts of OT can be stretched to cover these other organizations as well. Then all this stretching has to be researched, one piece at a time.

Second, we must learn more about the links between the different types of organizations. How can simulated organizations provide insights in natural- and MAS organizations? For the first category, a lot of work has already been done. Many simulations already research different aspects of natural organizations. Especially reorganization strategies for human organizations can be improved with further research on using simulations as a basis for decision support tools. Multiple perspectives are only a small part of the vast body of improvements needed in simulations to realistically simulate real situations.

Finally, and perhaps most interestingly, MAS organizations can be improved with research such as this thesis. Compared to human organizations, we have little theory about MAS organizations and specifically how to reorganize those organizations. We can learn much more about MAS and their interaction with human organizations through further research. This research will have to address several issues, including the very important issue of creating efficient reorganization strategies. Directly following this thesis, it would be very useful to try and implement organization evaluation criteria as measurements upon properties of the organization in a real MAS.

As a final note I *don't* think we can learn everything, but in the end I *do* think this thesis –and future research proposed here– allows us to *have fun being useful*.

Appendix A: Simulation Parameters

This appendix describes in detail what the parameters of the simulation are for. The first column is the name of the parameter that is used in this thesis. The name used in the Repast environment is in the second column. Finally in the last column the effect of the parameter on the simulation is described.

Parameter	Name	Effect
$BANK_{MIN}$	BankValueMinimum	This is the lower boundry of the bank value parameter. If the bank value drops below this value the simulation is stopped, because the store went bankrupt. It corresponds with the maximum amount of money the store can borrow at the bank.
$BANK_{T}$	BankValueStart	This is the amount of money the store starts with. This value is copied to a bank variable at the start of a simulation run. The value of this parameter is useless during a run, as the internal variable is used to keep track of the actual bank value.
COST(CONST)	ConstantCost	These are constant costs the store will always have.
C_1	CostTrend	This is the amount of money by which the constant cost is increased per timestep.
D_2	DemandExternalAmplitude	This is the amplitude of the external demand function. The difference between the highest and lowest point in the external demand function equals twice the value of this parameter.
D_1	DemandExternalBase	The external demand has this value as a base demand.
D_3	DemandExternalFrequency	This value determines how far apart two external demand peaks are apart. Normally, two peaks are about 3.2 timesteps apart, and this is multiplied by the inverse of this constant D_3.
D_4	DemandExternalTrend	This is the amount of money by which the constant demand is increased per timestep.
D_{JOB}	DemandHumanWithJob	A human with a job (i.e. a HUMAN or $HUMAN_{OTHER}$) spends this amount of money at the store, generating as much demand.
$D_{NO\text{-}JOB}$	DemandHumanWithoutJob	A human without a job (i.e. a human that's not a H or H_{OTHER}) spends this amount of money at the store, generating

		as much demand.
P(EXT)	ExternalDemandProportion	This indicates the proportion of demand coming from the external demand function (the function is multiplied by this value, before adding it to the total demand in a timestep).
D_5	ExternalDemandRandomProportion	This is the margin for the random variations in the external demand function. For example, if this has a value 0.1, each timestep the external demand will get a random variation from –10% to +10%.
	GraphUpdateFrequency	This determines how many timesteps must be calculated before plotting a new value in the graphs. Set this high (50 or so) to have the simulation run quicker.
	HeadStrategy	The strategy for the head agent can be selected through this parameter.
CAP(H)	HumanCapacity	The amount of demand a human employee can provide for. This value is the capacity for one human employee.
COST(H)	HumanCost	This is what one human employee costs for the store per timestep.
$SIZE_T(H_{OTHER})$	HumanOtherSize	This is the amount of humans that has work somewhere outside the store.
$size_t(h)$	HumanSizeStart	This is the starting amount of employees.
$SIZE_T(H_{\Sigma})$	HumanTotalSize	This is the amount of townsfolk. The amount of people with no job equals this value, minus the size of H and H_{OTHER}.
P(INT)	InternalDemandProportion	This value indicates what the proportion of the internal demand function is in the total demand, analogous to the P(EXT).
cap(m)	MachineCapacity	The amount of demand a machine employee can provide for. This value is the capacity for one human employee.
COST(M)	MachineCost	This is what one machine employee costs for the store per timestep.
$P_{MAX}(M)$	MachineProportion	This is the maximum proportion of machine employees that can be working at the store. The size of H divided by the size of M can never be above this proportion.
$SIZE_T(M)$	MachineSizeStart	This is the starting amount of machine employees.
$SIZE_T(M_{\Sigma})$	MachineTotalSize	This is the total amount of machines available to the store.
P(PRODUCTS)	ProductCostProportion	This value indicates how much the products coming with a certain capacity cost (i.e. the 'ingredients'). The cost of the products is calculated as the total capacity multiplied by this value.
	StretchFactor	This determines how much the profit goal evaluation gets stretched in the middle; this makes values around 50 easier to analyse.

Appendix B: Simulation Output

This appendix describes the four graphs that are plotted when doing a single run with the simulation. Note though, that the precise graphics might vary according to local settings.

Appendix B.1: Goal Evaluation

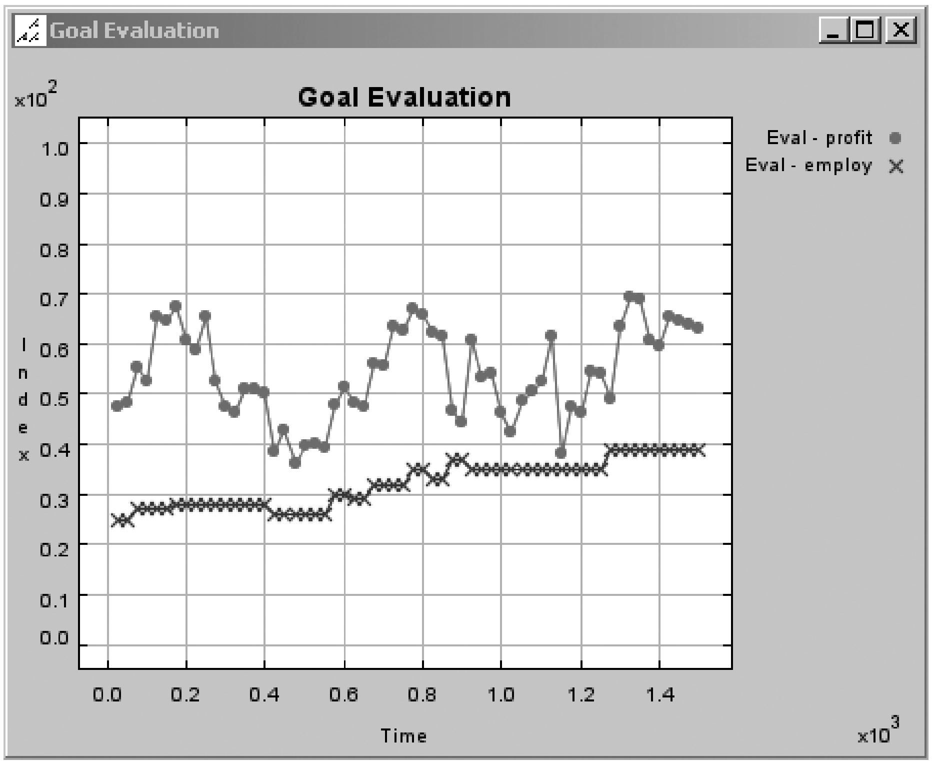

This graph plots the evaluation value of two goals over time. The X-axis shows timesteps, internally represented as "ticks". The Y-axis shows the [0…100] scale. The values for the first line "Eval – profit" are determined using the profit and total demand from the graph in C.3. Therefore it roughly follows the form of these other parameters. Similarly, the "Eval – employ" line follows the line of human employees in the graph from C.4.

Appendix B.2: BANK variables

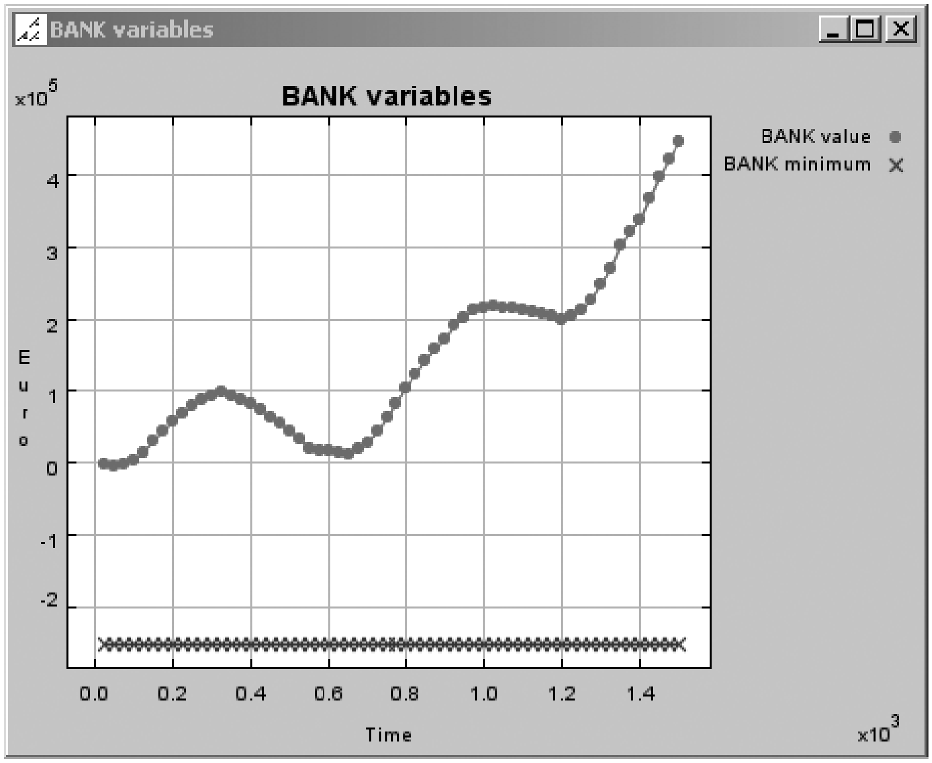

This is probably the simplest of all four graphs. The bottom line indicates the minimum value of the other line. If the "BANK value" line ever drops to the other line, the store has gone bankrupt. Any other variance in these line is determined directly by profit: if you gain money the line in this graph goes up, if you loose money it will drop.

Appendix B.3: Variable Functions

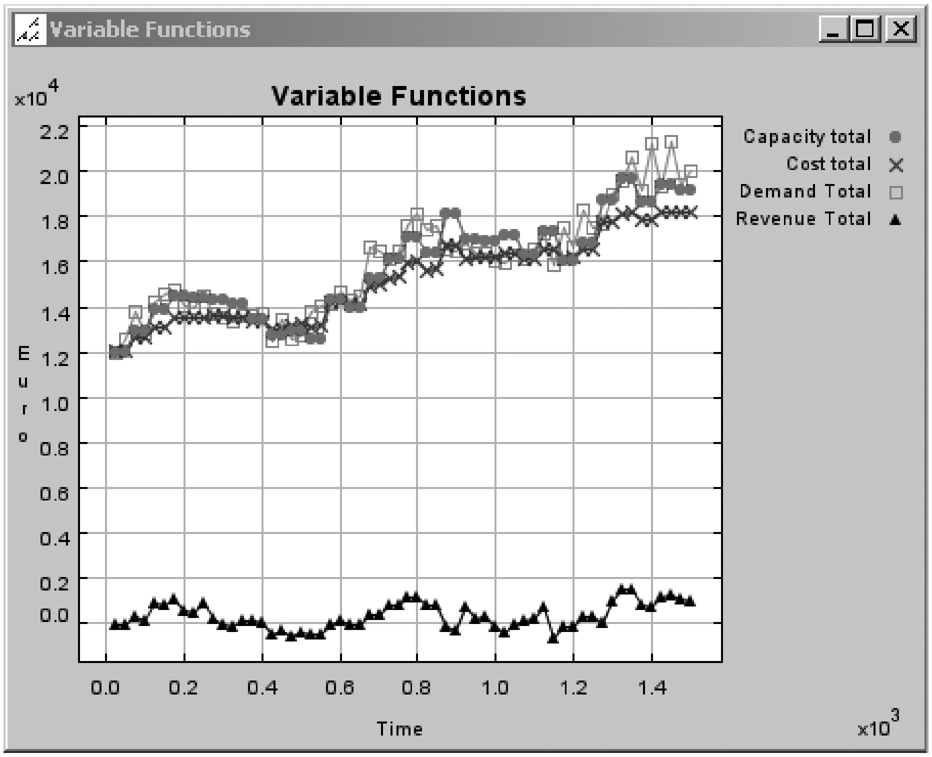

This is the most 'chaotic' graph. At the top, 3 lines all follow roughly the same path. The total demand minus the total costs forms the proift line at the bottom. As an exception to this, if the capacity line is below the demand (as is the case in the last plotted values), the profit line is formed by the capacity minus the cost.

Capacity is determined by how many agents are employed by the store, this can be seen in the graph from C.4. The manager tries to hire and fire agents to make sure this capacity line follows the demand line as closely as possible.

The "Cost total" line is determined by several factors. Important here is that the line follows the same path as capacity and demand for two reasons. First, the constant costs are increased a little bit each timestep. Second, when capacity is increased, agents were hired and employee costs therefore increased as well.

Appendix B.4: Agent Set Sizes

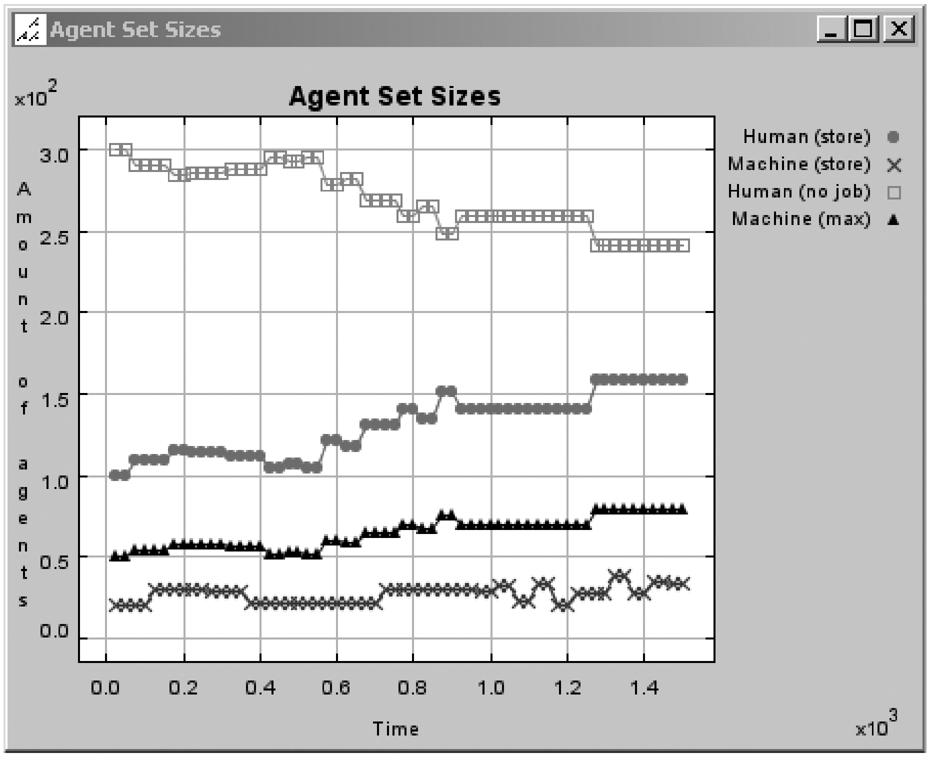

This graph shows how many agents are in different sets. The topmost line shows how many humans don't have a job. The second line from the top shows how many humans are employed by the store. If the first line drops, the second one rises with the same value, and vice versa.

The bottom line is the amount of machines. The line directly above that shows how many machines can be employed, which is determined by a proportion parameter. In this example graph the maximum number of machines is ½ of the amount of human employees.

Appendix C: List of Abbreviations

CAP	Capacity
CommonKADS	Common Knowledge Acquisition and Design Support
COT	Computational Organization Theory
D	Demand
DECIS	DElft Cooperation on Intelligent Systems
HOT	Human Organization Theory
MAS	MultiAgent System(s)
MABS	MultiAgent Based Simulation(s)
OperA	Organizations per Agents
OT	Organization Theory
R	Revenue
REPAST	Recursive Porus Agent Simulation Toolkit
TNO	Nederlandse Organisatie voor Toegepast Natuurwetenschappelijk Onderzoek (Dutch Organization for Applied Natural Research)

References

[1] Carley, K., Gasser, L. (1999), *Computational Organization Theory*, The MIT Press. From: "Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence", 299-330, editor G. Weiss, ISBN 0-262-23203-0.

[2] Dignum, F., Sonenberg, L. (2002), *A Dialogical Argument for the Usefulness of Logic in MAS*. From: "The Journal of Artificial Societies and Social Simulation, Volume 7, 4".

[3] Dignum, V. (2004), *A model for organizational interaction: based on agents, founded in logic*. SIKS Dissertation Series No. 2004-1, ISBN 90-393-3568-0.

[4] Dignum, V., Dignum, F., Furtado, V., Melo, A., Sonenberg, L. (2004), Towards a Simulation Tool for Evaluating Dynamic Reorganization of Agent Societies.

[5] Dignum, V., Dignum, F., Sonenberg, L. (2004), *Towards Dynamic Reorganization of Agent Societies*. From: "Workshop on Coordination in Emergent Agent Societies, ECAI 2004", 70-78, editor G. Vouros.

[6] Edmonds, B. (2002), *How Formal Logic Can Fail to be Useful for Modelling or Designing MAS*. From: "The Journal of Artificial Societies and Social Simulation, Volume 7, 4".

[7] Edmonds, B. (2002), *Comments on "A Dialogical Argument for the Usefulness of Logic in MAS"*. From: "The Journal of Artificial Societies and Social Simulation, Volume 7, 4".

[8] Etzioni, A. (1964), *Modern Organizations*, Englewood Cliffs, NJ: Prentice-Hall.

[9] Fraasen, B.C. van (1980), *Arguments Concerning Scientific Realism*. From: "The Scientific Image", 6-21, 23-25, 31-40, Clarendon Press.

[10] Furtado, V., Melo, A., Dignum, V., Dignum, F., Sonenberg, L. (2005), *Exploring congruence between organizational structure and task performance: a simulation approach*. From: "Proceedings of AAMAS05 worskhop, From Organizations to Organization Oriented Programming in MAS", 49-64, editors O. Boissier, V. Dignum, E. Matson, J. Sichman.

[11] Melo, A., Belchior, M., Furtado, V. (2005), *Analyzing Police Patrol Routes with the Simulation of the Physical Reorganization of Agents*. From: "Proceedings of

AAMAS05 workshop, Multi-Agent-Based Simulation", 93-106, editors L. Antunes, J. Sichman.

[12] Moss, S., Norling, E. (2005), *Having Fun Being Useful*. From: "Proceedings of AAMAS05 workshop, Multi-Agent-Based Simulation", 3-8, editors L. Antunes, J. Sichman.

[13] Musgrave, A. (1985), *Realism versus Constructive Empiricism*. From: "Images of Science", 197-221, University of Chicago Press.

[14] Padgham, L., Winikoff, M. (2002), *Prometheus: A Pragmatic Methodology for Engineering Intelligent Agents*. From: "Third International Workshop on Agent-Oriented Software Engineering".

[15] Perrow, C. (1970), *Organizational Analysis: A Sociological View,* Tavistock Publications Limited, ISBN 0-422-73480-2.

[16] Rollinson, D. (2002), *Organisational Behaviour and Analysis - An Integrated Approach, 2nd edition*, Pearson Education Unlimited, ISBN 0-273-65133-1.

[17] Schreiber, A.Th. et.al. (2000), *Knowledge Engineering and Management:The CommonKADS Methodology*, MIT Press, ISBN 0-262-19300-0.

[18] Wijngaards, N., Kempen, M., Smit, A., Nieuwenhuis, K. (2005), *Towards Sustained Team Effectiveness*. From: "Proceedings of AAMAS05 worskhop, From Organizations to Organization Oriented Programming in MAS", 77-89, editors O. Boissier, V. Dignum, E. Matson, J. Sichman.

[19] Wooldridge, M. (2002), *An Introduction to MultiAgent Systems*, John Wiley & Sons, ISBN 0-47149691-X.

[20] Zambonelli, F., Jennings, N.R., Wooldridge, M. (2003), *Developing Multiagent Systems: The GAIA Methodology*. From: "ACM Transactions on Software Engineering and Methodology, Vol. 12, No. 3", 317-370.

[21] Dutch TNO DECIS-lab website
http://www.decis.nl

[22] Repast Agent Toolkit website
http://repast.sourceforge.net

[23] Simulation Download Site
http://www.phil.uu.nl/~heijmans