Universidade Técnica de Lisboa
Instituto Superior Técnico

Teaming Up Humans and Synthetic Characters

Rui Filipe Fernandes Prada

Licenciado em Engenharia Informática e de Computadores

Dissertação para a obtenção do Grau de Doutor em Engenharia Informática e de Computadores

Orientador: Doutora Ana Maria Severino Almeida e Paiva

Júri

Presidente: Reitor da Universidade Técnica de Lisboa

Vogais: Doutor Thomas Rist, FHA - Fachbereich Informatik, Augsburg, Germany;
Doutor Eugénio de Oliveira, Faculdade de Engenharia da Universidade do Porto;
Doutor José Manuel Tribolet, Instituto Superior Técnico da Universidade Técnica de Lisboa;
Doutor João Madeiras Pereira, Instituto Superior Técnico da Universidade Técnica de Lisboa;
Doutora Ana Maria Paiva, Instituto Superior Técnico da Universidade Técnica de Lisboa;
Doutora Helena Sofia Pinto, Instituto Superior Técnico da Universidade Técnica de Lisboa

Dezembro 2005

Resumo

Com o aparecimento e desenvolvimento de personagens sintéticas em ambientes virtuais, estes podem agora ser habitados simultaneamente por este tipo de personagens e utilizadores, todos interagindo, colaborando ou competindo entre si. No entanto, esta interacção nem sempre é a melhor e só será positiva se as personagens foram capazes de demonstrar um comportamento credível e coerente. Por isso, em cenários onde utilizadores e personagens sintéticas interagem como um grupo na resolução de tarefas é muito importante que as interacções sigam uma dinâmica de grupo credível.

Focando este problema, desenvolvemos um modelo que suporta a criação de uma dinâmica grupo credível num grupo de personagens sintéticas. Este modelo foi inspirado em teorias de dinâmica de grupo desenvolvidas em ciências humanas sócio-psicológicas. A dinâmica é conseguida através de uma caracterização dos diferentes tipos de interacções que podem ocorrer no grupo e lida com as interacções sócio-emocionais assim como as interacções relacionadas com a tarefa.

Este modelo foi implementado no comportamento das personagens sintéticas que colaboram com o utilizador na resolução de tarefas num ambiente virtual. As tarefas colaborativas foram desenvolvidas na forma de um jogo de computador que envolve o utilizador com um grupo de quatro personagens sintéticas constituído com o propósito de resolver os desafios propostos. O jogo foi utilizado num estudo que foi conduzido para determinar o efeito do modelo na interacção do utilizador com o grupo e os resultados mostraram que o modelo teve um efeito positivo na interacção dos utilizadores, em especial na sua confiança e identificação com o grupo.

Palavras Chave: personagens sintéticas, dinâmica de grupo, colaboração, interface pessoa máquina, sistemas multi-agente, ambientes virtuais

Abstract

In recent years, virtual environments have evolved from single user and single agent, to multi-user and multi-agent, thus, engaging simultaneously users and synthetic characters all interacting, collaborating or competing with each other. This new scenario created new challenges for the users' interaction with the environment, in particular, for their interaction with the autonomous synthetic characters. To engage in successful and believable interactions the synthetic characters must be able to show a coherent set of behaviours responsive to the user's actions. For example, in scenarios where users and synthetic characters engage the resolution of collaborative tasks, thus, interacting as a group, it is very important that their interactions follow a believable group dynamics.

Therefore, focusing on this problem, we have developed a model that supports the creation of a believable group dynamics on a group of synthetic characters, inspired by theories of group dynamics developed in human socio-psychological sciences. The dynamics is driven by a characterization of the different types of interactions that may occur in the group and deals with both the socio-emotional interactions and the task related interactions.

We have implemented the model into the behaviour of synthetic characters that collaborate, in a group, with the user in the resolution of tasks within a virtual environment. The collaborative task was developed in the form of a computer game that engages the user in a group with four synthetic characters constituted to solve together puzzle like challenges. This game was used in an experiment conducted to assess the effects of the model on the user interaction with the group. The results showed that the model had a positive effect on the users' interaction experience, specially in their trust and identification with the group.

Keywords: synthetic characters, group dynamics, collaboration, human-computer interaction, multi-agent systems, virtual environments

Acknowledgements

I would like to thank all my friends and family for their unconditional support. First of all, I would like to express my warmest thanks to Joaninha for being always there in the most difficult moments of the development of this work. These same words are extended to my parents and sister for all the help they offered.

Special thanks go to my supervisor, Ana Paiva, and the people from the GAIPS research group. In particular, to Carlos Martinho, Marco Vala, João Dias, Guilherme Raimundo, David Pontes, Daniel Sobral, Isabel Machado and Patricia Lima.

Nuno Otero and Ana Passos deserve a special mention for their help and comments concerning the social and psychological view of the group dynamics.

I would also like to thank my special friends, David Prada and Catarina Gouveia, for all their support and concern.

Finally, I would like to mention that this work has been developed with the support of a scholarship (PRAXIS XXI/BD/21658/99) granted by the *Fundação para a Ciência e a Tecnologia*.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Motivation

In the last few decades, the use of computers has become so widespread that it is now part of everyone's daily life. Today, people interact with computers most of the time in their daily routines: in their work, at home, while shopping, in the car, even on the beach through their mobile devices. This new reality has brought to the surface the importance of the study of human-computer interactions, making such study one of the crucial issues in the development of new computational systems. This dissertation is, therefore, about users, computers and their inevitable interactions.

Moreover, computers are primarily built as tools that solve problems for humans that interact with them. Thus, computational systems should be designed in a way that makes their use efficient for the task. But not only. They should also be designed in a way that makes their use natural and similar to the way we humans interact. Elements such as natural language interpretation, speech synthesis and recognition, graphical 3D interfaces, are now a constant need in our computer applications.

Users want and expect the interaction with computers to evolve, to become more natural in a way that alleviate the burden and difficulties of dealing with technology. Therefore, a very effective and efficient computational system may not have any success if, for example, users are *afraid* of using it. Indeed, human factors such as emotions may be involved in human-computer interactions. People's emotions influence their interactions with computers and reciprocally, the interaction influences the emotional state of the user [127]. On the other hand, humans are social animals. They live in society, get organized around groups, and even treat objects and machines as social actors.

Reeves and Nass [140] have shown that users interact with computers by applying the same social rules that they use in human to human interactions. More than that, their extensive work shows that people's interactions with computers are essentially social. They have found for example, that people prefer polite computers and, at the same time, they are polite to computers; that users like to be flattered by computers; that users assign personalities to computers and prefer to interact with computers that match their own personality; that people use social conventions on gender when interacting with computers if they can make the distinction between male and female (e.g. a computer that plays a male voice versus a female voice). These behaviours were not only observed in novice users but also among expert computer users.

These results suggest that when designing user-computer interactions one should also take into account some social factors in the relation between applications and humans.

The fact that people actually treat computers the same way as they treat a human has opened a new focus for research on new ways of human-computer interaction: to build interfaces that behave socially and meet the users' social expectations. Following this line of research, the use of *synthetic characters* emerges as a way of bringing "life" to the computational objects and to, thus, provide some social dimension to the user interface. *Synthetic characters* are life-like entities, that inhabit a synthetic world, usually having an anthropomorphic appearance. The user's interaction with these characters is intended to be natural and as close as possible to the interactions users can have in the real world.

Synthetic characters have been used in many different domains. In commercial applications, for instance, there are characters that perform as salesperson: for example REA [19] the Real-State Agent helps the user to find her/his perfect house. Other characters are used in educational and training scenarios: for example STEVE [144] was used to train teams of people to handle emergency situations that may arise inside a military ship; Cosmo[97] was developed to help students understand the mechanisms of network routing; and in Teatrix [104][134] the characters make part of a cast of actors that together with the students create a fairy tale story and, consequently, help the children develop their literacy skills. Additionally, in entertainment, *synthetic characters* are used, for example, in computer games, in cinema, and in some exhibit installations. Examples of this are the (void*)[22], which provides a very entertaining experience to users while they influence the way two characters dance, and S3A

[107] that gives the users a chance to interact with two synthetic dolphins.

1.2 The Problem

With the emergence of *synthetic characters*, collaborative virtual environments are now populated with this type of characters and users at the same time, all interacting, collaborating or competing with each other. Examples of this can be found in some computer games, more specifically Role Playing Games [1], such as "The Temple of Elemental Evil"[173] and "Star Wars: Knights of the Old Republic"[21], and in virtual internet communities such as Active Worlds [1] [2].

In these environments, users may collaborate with several *synthetic characters* and form groups that engage the resolution of tasks. However, this interaction with synthetic characters is not always the best, and it will only be positive if the characters are able to show a coherent and believable behaviour. Believability in *synthetic characters* highly depends on the richness of the characters' actions and interactions, on their expressions, and, more importantly, on how well they lead the user to the suspension of disbelief as stated by Bates in 1994 [17]. Moreover, in the case where characters perform in groups, *synthetic characters* must not only be believable as individuals, but also convey a believable group dynamics.

Given that humans are good at social relationships, if the rules for using technology are in essence socially inspired, then no additional skills should be required from the user to interact with such applications. Aspects of human social interactions may include, among others issues, speech, recognition of each others' emotional state, collaboration and cooperation.

However, most of the research conducted on believability of *synthetic characters* is centred on the interactions between a user and a single character (see for example the case of Rea [19], Sam [27], Greta [124]) or on the interactions of the *synthetic characters* themselves [172] [148] without considering the user within the group and without a common collaborative task. Thus, there is a clear need for research on scenarios that engage the user in collaboration with

[1] A role-playing game (RPG) is a game in which each participant assumes the role of a character (such as a brave medieval knight or a futuristic spaceship captain) that can interact within the game's imaginary world and its characters. Characters usually form groups and act together in the search for a solution to the world's quests.

[2] Active Worlds supports a virtual community over the internet in a virtual 3D environment in which users register as citizens and may organize and participate in many socials events as parties and conferences, and can even enroll in a university. This community has many autonomous characters that enrich the environment.

a group of *synthetic characters*.

The work presented in this document addresses the problem of creating and promoting the existence of groups of humans and *synthetic characters*, while making the emergent interactions and collaboration natural and believable. To achieve such group believability it was assumed that the interactions and the group dynamics should resemble the collaboration and interactions that emerge in human groups. Given this, the main hypothesis that will be proven with this work is:

> *If the interactions in a group of* synthetic characters *follow similar dynamics as the interactions in human groups do, the synthetic group will become more believable, which then improves the interaction experience of a user that actively participates in the group.*

1.3 Why Model Group Behaviour in *Synthetic Characters*

The results of this research can be used in many different application areas. The first, and perhaps the most relevant one, is entertainment. For example, there are many multi-player computer games that may have several players and *synthetic characters* engaged as a team. One particular case are the Role Playing Games where players, together with some *synthetic characters*, form a group of heroes that undertake the quests of a fantasy world. In this type of games the social interactions are an important part of the game, specially those that take place between the members of the group. And, therefore, the role of the *synthetic characters* is usually very restricted. Additionally, players frequently have some control over the characters, which reduces their autonomy. For example, in the "Star Wars: Knights of the Old Republic"[21] game, the player starts the adventure with one character, but as the game evolves other characters join the player's quest and s/he will end up controlling simultaneously an entire party of several characters.

This fact increases the distance between the player and her/his character and decreases the role playing of the game. For this reason, and in order to achieve a better level of role playing, Role Playing Games are often played by several users each one controlling a single character. In these scenarios the *synthetic characters* are limited to the role of servants or companions that follow their masters and do not actively participate in the group. However, if *synthetic characters* could interact and collaborate in a natural way within a group of human players, they could take a more active participation, thus, play more central roles in the group. Furthermore, in the absence of other human players, these *synthetic characters* may bring the same levels of role playing to the game and make it as enjoyable as if there were only humans involved.

Another prominent area of application is cinema, where the use of *synthetic characters* to simulate populations of actors, specially in crowds, is now a widespread practice. One well known example is the "The Lord of the Rings" trilogy [115]. These films included numerous fighting scenes involving armies of thousands of warriors, the major part of these being played by synthetic actors. Furthermore, it is an undeniable fact that the use of *synthetic characters* made possible the realization of scenes that were very difficult to attain before the use of these techniques. However, with the current technology, the synthetic actors will always be relegated to the minor roles that do not involve high social engagement with the main characters. Thus, enhancing the social abilities of such *synthetic characters*, as it is the purpose of this work, will make it possible to, in the future, allow these characters to play more central roles in a film. Imagine for example, an autonomous *synthetic character* playing the role of one of the members of the fellowship of the ring when they decided the ring's fate in the Elrond's council scene at Rivendell.

Moreover, in education and training these results can be used in applications that train team work, such as the application where STEVE [144] was used. The team training can be enhanced by including some social training to endow the learners with the ability to manage the group social relationships as well as the action cooperation procedures. For this, it is important that the synthetic participants behave in a believable way towards the group and its members.

The same ideas can be applied to children's education. Researchers have found that learning in group may foster the learners' knowledge building ability [163]. For example, Aebli [3], supported on Piaget's theory of cognitive development [126], stated that learning how to behave in group is fundamental to early children's development, since working and discussing with others requires that children take different points of view and see the other individuals' perspectives. This effort helps children achieve a more flexible and logical reasoning moving their thought from egocentric to operational [3]. This process of children's development can be

[3]Piaget divides the child's cognitive development into four different stages: the *Sensor-Motor (Egocentric)* stage (age 0-2) when a child's behaviour is primarily motor the child perceives and manipulates but does

supported by a computer software that simulates believable group interactions.

1.4 Research Goals

In the development of this research several goals were pursued. First of all virtual environments with *synthetic characters* can be seen as multi-agent systems [38]. Thus, our first goal was to build a framework that supports the creation of multi-agent systems with the presence of human participants.

The second goal concerns the creation of a model of the dynamics of groups with synthetic and human participants. The idea is to apply human-human group dynamics and relational theories to model the interactions and dynamics that emerge in groups that are created in a virtual environment cohabited by synthetic and human participants.

The third goal was to evaluate the effects of the model developed for the second goal on users that interact in virtual environments with groups containing *synthetic characters* as members, and by this means validate our hypothesis.

1.5 Methodology

The work described in this dissertation was conducted in the following steps:

- First, knowing beforehand that virtual environments with *synthetic characters* are usually implemented with the use of multi-agent systems, we have reviewed the field of multi-agent systems paying special attention to the agents' mechanisms of interaction, namely in collaborative scenarios.

- Then, we have reviewed the notion of virtual environment and its interaction issues, specially when it concerns the interaction with *synthetic characters*. In this process, we have explored the notion of social intelligent agents, which are *synthetic characters* that make use of social skills to interact with humans in a more natural and believable way.

not reason; the *Pre-operational* stage (age 2-7) when symbolic thought develops (uses language and mental representations) and children are seen to remember, imagine and pretend; the *Concrete Operational* stage (7-11) when children begin to learn how to handle the basics of logical thought but still rely on concrete objects, they can perform mental operations with concrete materials but not with abstract ideas; the *Formal Operational* stage (11 plus) when abstract problems can be solved and the ability to formulate hypotheses is reached.

- Additionally, we have reviewed several systems that use *synthetic characters* and have some of the characteristics that we believe to be important for the group interactions: (1) characters that are able to establish social relations with the user, (2) characters that are able to establish social relations with each other, (3) characters that collaborate with the user and (4) characters that behave in group.

- At this point, we have developed a framework that supports the creation of multi-agent virtual environments and promotes clean integration between the visual and the behavioural components of the system.

- Furthermore, having in mind that our *synthetic characters* should follow a believable group dynamics, as stated in our hypothesis, we have reviewed some theories of group dynamics developed in human social psychological sciences. We gave emphasis to those theories that are centred on the perspective of the individual and that defined groups as systems of interacting members, because they seemed easier to apply to computational systems.

- Then, inspired by the reviewed theories of group dynamics, we have designed a model for the group dynamics of *synthetic characters*. The model defines a characterization of each individual member of a group and the social relations that are built with the other members during the group process. The dynamics is driven by a system of categorization of the group interactions which influence and, at the same time, are influenced by the individual characterization and the established social relations.

- In order to evaluate the model for synthetic group dynamics, we have designed a computer game that engages one user in a group with four *synthetic characters*. This group is successively challenged, as the game evolves, with puzzle like collaborative tasks. The *synthetic characters* were implemented as autonomous agents that make use of the synthetic group dynamics model to drive their behaviour, and the whole system was integrated using the proposed multi-agent framework.

- Finally, we have conducted a study with some students from our university to evaluate the effects of the synthetic group dynamics model in the user's interaction with the group of *synthetic characters*. In the study, subjects were divided in three different groups that interacted respectively with three different versions of the game, one not using the model and the other two using the model but with different group conditions.

1.6 Contributions

Our work produced some scientific contributions that we believe to be relevant for the research fields of multi-agent systems and *synthetic characters*. The first one, is related to the multi-agent framework that we describe in chapter 6. This framework was developed having in mind the integration issues of the two different type components that are usually present in the development of virtual-environments with *synthetic characters*: the behavioural components, that implement the intelligent behaviour of the simulated entities, and the presentational components, that present the simulated world to the user.

These issues are usually not explored in the frameworks that people typically use to build their systems. Which are either, very dependent of the visualization system (e.g. the 3D graphics engine) or very detached from the presentation perspective disregarding the user's interaction with the system. Our proposal merges the two perspectives without creating strong dependences between the two. This allows, for example, to support the integration of several different views on a simulated world in order to support the collaboration of people with different needs, e.g. between sighted and unsighted people.

The ideas of this framework were successfully used in different contexts, for example, in the VICTEC project [122] [168], and are being incorporated in a more general framework that is being developed in our research group [167].

Furthermore, our main contribution is the SGD Model that we developed to support the dynamics of the interactions in groups with synthetic and human participants, which is described in chapter 5. This model incorporates some ideas regarding the socio-emotional and social influence structure of the group that are usually not found on similar research. In addition, we have successfully used the model and confirmed that it has some positive results on the users' interaction in groups with synthetic members. We also believe that this model can be easily extended and that it can be easily applied to other contexts where several synthetic agents engage in a collaborative task.

Moreover, our ideas and results have been successfully published in some relevant conferences, such as the *International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, the *International Working Conference on Intelligent Virtual Agents (IVA)* and the *International Conference on Advances in Computer Entertainment Technology (ACE)*. We have published a total of five papers ([129] [130] [133] [131] [132]) and we would like to give a special mention to the one published in the AAMAS'2005 conference [131], because it was selected as a candidate for the best paper award.

Finally, we think that this work brings attention to some issues of user interaction with *synthetic characters* that have not been explored by the research community before, in particular, the interaction in collaborative groups. We believe that we made clear the importance of the believability of group interactions in *synthetic characters* and established some bases for future work on the field.

1.7 Outline

This document is divided into nine different chapters including this one.

The second chapter describes the notions of agent and multi-agent systems and their common definitions and issues.

On the third chapter we describe the notions of virtual environments and *synthetic characters* and discuss their interaction with humans. In this chapter, several systems that make use of *synthetic characters* are also described, in particular systems with characters that present some social skills.

The fourth chapter describes some theories of group dynamics developed in human social psychological sciences that sustain the design of our model for the dynamics of groups of *synthetic characters*.

The fifth chapter describes the Synthetic Group Dynamics Model (SGD Model) that we propose to attain the believability of the group interactions of *synthetic characters*.

The sixth chapter describes a framework that was developed in the context of this work to support the creation of multi-agent virtual environments.

The seventh chapter presents the group task, in the form of a collaborative computer game, that was developed to be used in the study conducted to evaluate the effects of the model in the interaction of users with a group of four *synthetic characters*. This chapter also presents some details of the development of the game, in particular, the integration of the SGD Model on the synthetic members of the group.

The description of the evaluation study and discussion of its results are presented in chapter eight.

Finally, in chapter nine we present some conclusions and future work.

Chapter 2

Multi-agent Systems

2.1 Introduction

Collaborative virtual environments are often populated by synthetic characters and users at the same time, all of them interacting, collaborating or competing with each other. These environments can be seen as composed by a ”world” (virtual world) and a set of entities that perceive and act in such world (the users and the characters). This notion is clearly related to the notion of a multi-agent system [178] where each individual character is seen as an autonomous agent acting and perceiving the world.

In addition, many concepts of multi-agent systems draw inspiration from theories of human social intelligence, which is also one of the main focus on building synthetic characters [38]. However, although this association exists, the community of multi-agent systems has for the past few years, focused its research primarily on the formal and engineering aspects of multi-agent systems, thus addressing problems such as their formalization, optimization, multi-agent planning, agent's communication, cooperation and negotiation, without regarding to issues of believability and interaction with users. Nevertheless, multi-agent systems constitute a strong technology support for the construction of collaborative virtual worlds.

Therefore, to establish a common ground for the research here presented and to clarify some terms that define the starting point for the construction of believable groups, this section presents an overview on the field of multi-agent systems, discussing its common definitions and problems and how it impacts on the research here presented.

2.2 The Definition of Agent

Multi-agent systems are systems composed by a multitude of interacting entities, here called *agents*. Thus, to describe multi-agent systems, we surely need to define the term *agent* beforehand. One problem with that is that there is not a clear definition for it, not even a close agreement of the attributes that agency should bear. This difficulty comes from the fact that an agent definition depends largely on the domain where it applies. Most common definitions however, agree on the fact that agents are always placed inside an environment and agree on some properties that agents should have: autonomy, sensing and acting capabilities. Figure 2.1 resumes the agent interaction with the environment.

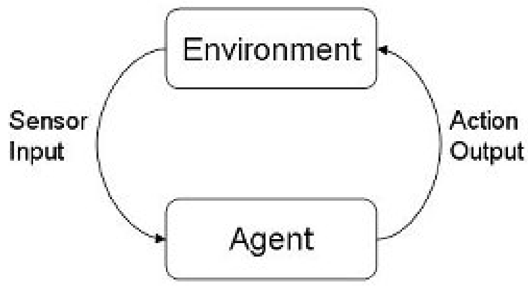

Figure 2.1: An agent in its environment. The agent receives stimuli from the environment through the sensory input and acts using the action output process.

The most common definitions of agent include the Russell and Norvig [147], Pattie Maes [105], Hayes-Roth [69] and Wooldridge and Jennings [179] [178] definitions which are presented below:

> *"An agent is anything that can be viewed as perceiving its environment through sensors and acting upon that environment through effectors"* Russel and Norvig [147]
>
> *"Autonomous agents are computational systems that inhabit some complex dynamic environment, sense and act autonomously in this environment, and by doing so realize a set of goals or tasks for which they are designed."* Pattie Maes [105]
>
> *"Intelligent agents continuously perform three functions: perception of dynamic conditions in the environment; action to affect conditions in the environment; and*

> *reasoning to interpret perceptions, solve problems, draw inferences, and determine actions."* Hayes-Roth [69]
>
> *"An agent is a hardware or (more usually) a software based computer system that enjoys the following properties: autonomy, social ability, reactivity, and pro-activeness."* Wooldridge and Jennings [179] [178]

Or, more recently, the definition found in [178], where

> *"Agents are complex systems with two important capabilities. First, they are at least to some extent capable of autonomous action of deciding for themselves what they need to do in order to satisfy their design objectives. Second, they are capable of interacting with other agents not simply by exchanging data, but by engaging in analogues of the kind of social activity that we engage in everyday of our lives: cooperation, coordination, negotiation and the like."* Wooldridge [178]

The first three definitions are centred on the notion of environment and enhance the need of the agent to interact with it through its sensory input and action output processes. In addition, Wooldridge's definitions enhance the social dimension of the agents, together with their autonomy.

Moreover, the diverse studies conducted on agents have proposed several properties to characterize them. This set of properties change according to the domain where the agents are used, which makes it difficult to determine a set that applies to every agent and every domain. As a summary, the definition of the most common properties of agents is presented here:

- **Autonomy:** agents operate without the direct intervention of humans or others, having control over their actions and internal state;
- **Social ability:** agents interact with other agents (and possibly humans) via some kind of agent communication language;
- **Reactivity:** agents perceive their environment and are able to respond to its changes in a reasonable time interval;
- **Pro-activeness:** agents do not simply act in response to their environment; they are able to exhibit goal-oriented behaviour by taking the initiative;

- **Adaptability:** agents adapt their behaviour according to the experience that they gradually gain while interacting with the environment. Agents learn from their previous experience.

- **Mobility:** agents can move in the environment from one place to another.

- **Emotional behaviour:** agents can express emotional behaviour. Emotions are synthesized and can influence the agent thinking and drive its course of actions. Emotions may also play a very important role when the agent is interacting with humans.

- **Personality:** agents may exhibit personality traits that influence their behaviour. The agent behaviour must always be consistent with the agent personality.

- **Rationality:** agents act in order to achieve their goals and should not take any action that goes against any of them. All the time, agents act in a way that, according to their knowledge and abilities, is the best to achieve their goals.

- **Persistence:** agents are persistent while pursuing their goals. If a plan to achieve their goals fails, they should elaborate a new one and continue acting towards the satisfaction of their goals.

- **Temporal continuity:** agents are continuously running processes in their life cycle (either running active in the foreground or sleeping/passive in the execution background) and not ”one-shot” computations that map a single input to a single output and then terminate.

- **Veracity:** agents will not intentionally communicate false information.

- **Benevolence:** agents do not engage in conflicting goals and will always try to do what they are asked for.

- **Delegation:** agents, or users, can delegate actions to other agents. The delegated agents should perform the actions taking into account the interests of the agent that requested the actions and should report the results of the actions to it.

2.3 The Characterization of the Environment

As we can see from the definitions of agent presented, an agent can not be clearly established without the characterization of the environment where it will act. The agent will always be related with the environment and its performance highly depends on the environment specific characteristics. Russell and Norvig [147] propose a characterization of environments as follows:

- **Accessible versus inaccessible:** an accessible environment is one in which the agent can obtain complete, accurate, up-to-date information about the environment's state.

- **Deterministic versus non-deterministic:** a deterministic environment is one in which any action has a single guaranteed effect, there is no uncertainty about the environment's state that results from the performance of an action. This is not certain in a non-deterministic environment.

- **Static versus dynamic:** a static environment is one that remains unchanged except by the result of the agent's actions. Conversely, in a dynamic environment, this state can change as a result of other processes beyond the agent's control.

- **Discrete versus continuous:** an environment is discrete if there is a limited number of distinct actions and perceptions defined. Chess is discrete; there are fixed numbers of possible moves each turn. Taxi driving is continuous - the speed and location of the taxi and other vehicles sweep through a range of continuous values.

- **Episodic versus non-episodic:** in an episodic environment, the agent's experience is divided into "episodes". Each episode consists of the agent's perception and, then, action.

From this description we can easily conclude that the scenario where the agent will have more problems in its performance is when the environment is simultaneously inaccessible, non-deterministic, dynamic, continuous and non-episodic.

In the case of virtual environments with several agents (agents that can also represent the user), we can say that accessibility may be possible (given that the whole environment is simulated), determinism of the agent's actions can also be achieved (although the presence of humans and other agents constrain it), discrete and episodic environments are also possible, but static environments are highly improbable.

However, in the concrete scenario that is addressed in this dissertation, which involves a collaborative environment with both autonomous agents and users that engage in rich social interactions, the agents' actions will certainly not be deterministic, since it will not be possible for the agents to predict with accuracy the social reactions that their actions might provoke on the others, specially on users. Additionally, the presence of human agents in the environment makes the agents' experience to be hardly episodic.

2.4 Agent Architectures

The definition of agent presented above does not compromise the agent to any specific internal architecture. However, to achieve certain properties, agents may have to embed certain functional elements.

To describe the components of the agents, researchers often classified the type of control architecture that they have. According to Wooldridge and Jennings [179] there are three major classes of agents based on their architecture:

- **Reactive Agents:** A reactive agent does not include any kind of central symbolic world model, and does not use complex reasoning on its decisions. The decision is made on the present without considering any past experiences. Brooks' subsumption architecture [24] is the most referenced reactive architecture.

- **Deliberative Agents:** A deliberative agent is an agent that contains an explicit model of the world, and uses reasoning based on the world model to make its decisions. This kind of agents are sometimes divided into two sub categories: deductive agents, those that use pure logic on its reasoning (e.g. theorem-proving algorithms), and practical reasoning agents that make its reasoning based on actions (e.g. action-planning algorithms). The best known and more referenced deliberative architecture is the Belief/Desire/Intention (BDI) architecture [45] [138] [137] developed at Stanford Research Institute.

- **Hybrid Agents:** Hybrid agents are those that are neither completely reactive nor completely deliberative. They use both approaches on their internal architecture, often incorporating the two aspects through a layered architecture. One of the best known hybrid architectures is the Georgeff and Lansky Procedural Reasoning System (PRS) [59].

But, perhaps, the type of architecture that had the major impact in the construction of intelligent agents is the Belief/Desire/Intention (BDI), since it is one of the most referenced and used agent architectures.

The BDI architecture was a result of the work of the Rational Agency project at Stanford Research Institute and is based on the theory of human practical reasoning developed by Michael Bratman [23]. It is a deliberative agent architecture in which the agent state is described by a set of three mental states: its *beliefs*, its *desires* and its *intentions*. The *beliefs* represent the knowledge that the agent builds about the world at each instant. The *desires* are the needs and wants of the agent and derive the agent general goals. The *intentions* represent commitments about tasks and actions to take in order to achieve the agent's goals. They represent the agent's current focus.

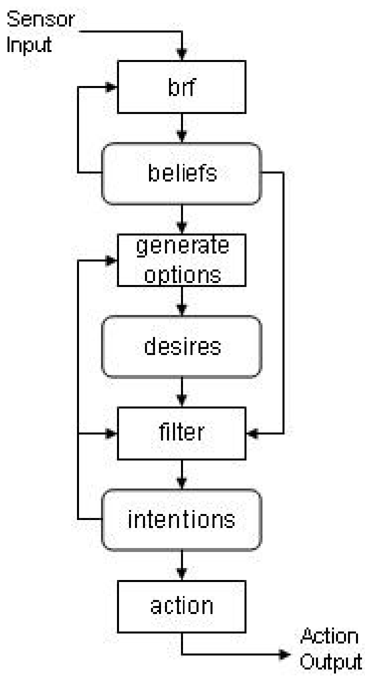

Figure 2.2: Diagram of a generic belief-desire-intention architecture.

In figure 2.2 one can see the abstract BDI architecture that was originally proposed by Rao and Georgeff [138]. The architecture includes seven components that were described as follows:

1. A set of current *beliefs*, representing information that the agent has built about the environment.

2. A belief revision function (brf), which takes a perceptual input and the agent's current *beliefs* and determines a new set of *beliefs*.

3. An option generation function (generate options), which generates the options available to the agent, its *desires*, on the basis of its current *beliefs* and its current *intentions*.

4. A set of *desires*, that represents the agent's objectives and their associated priorities and payoffs.

5. A filter function, which represents the agent's deliberation process, and which generates the agent's *intentions* according to its current *beliefs*, *desires* and *intentions*.

6. A set of current *intentions*, representing the agent's focus, the goals that the agent is committed to achieve.

7. An action selection function (action), which chooses an action to perform based on the agent's current *intentions*.

2.5 Defining Multi-Agent System

A multi-agent system combines several agents which can interact with each other in an environment. Furthermore, in addition to sharing a common environment the agents are often linked through organizational relationships. Jennings [88] proposed a typical structure of a multi-agent system which is shown in figure 2.3.

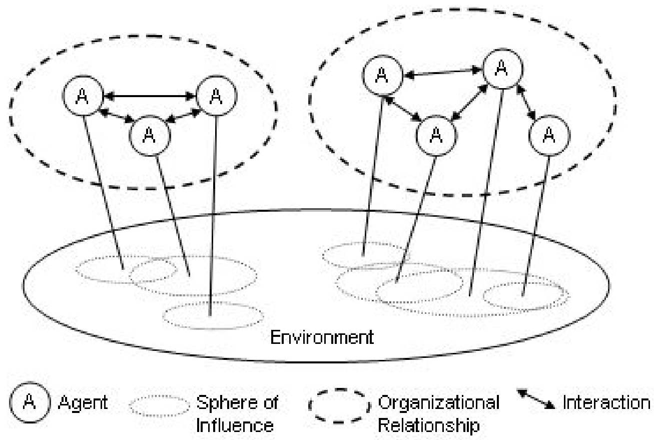

Figure 2.3: Typical structure of a multi-agent system.

Additionally, Jennings presented the notion of agent's influence sphere, which defines the portion of the environment that the agent can influence or act upon. Influence spheres of different agents can, and often do, intersect each other. For that reason agents may have dependency relations, which are categorized according to Jennings [88] as:

- **Independent:** there is no dependence between the agents.
- **Unilateral:** one agent depends on the other, but not vice versa.
- **Mutual:** both agents depend on each other with respect to the same goal.
- **Reciprocal dependent:** both agents depend on each other but not necessarily for the same goal.

The central characteristic of a multi-agent system is that its agents have the opportunity to interact with each other. However, the type of interaction differs according to the attitude of the participants in the interaction. Usually the agents take one of these three attitudes:

- **Competitive:** the success of one agent depends on the failure of others.
- **Cooperative:** the participants engage in the same common goals.
- **Self-Interested:** the participants have private independent goals but can cooperate in order to achieve them.

When several agents engage in an interaction event, such event is called an encounter. In these encounters, agents' interests may conflict which means that a proper resolution must be found. In the case of agents that are not cooperative this resolution may become very difficult or even impossible to achieve.

2.6 Resolving Agent Encounters

One of the main concerns in multi-agent systems is how to manage the interactions between the several agents. These interactions should be not only efficient and effective as possible, as far as the individual agents is concerned but also for the overall multi-agent system. Thus, what may be good for one individual agent, may be bad for the whole society.

One well known example from the game theory is the Prisoner's Dilemma [9]. This example describes a scenario, where two suspects accused of the same crime are both held by the police, but are inquired separately. They are both guilty and previously agreed that would deny the crime at any circumstance. However, the police told each of them that if one of the suspects confesses the crime while the other denies it he could go free and the other gets a ten year sentence. Additionally, they know that the sentence is of six years if they both confess the crime but they only take the minimum sentence of one year if they both deny it. The best choice concerning both prisoners well-being is clearly to cooperate and deny the crime, since each would get a one year sentence. But, if they concern only with their own well-being they may decide to betray the other by confessing the crime in the expectation of getting free. However, by following this behaviour probably both get a six year sentence.

Furthermore, even if the several agents share the same objectives and act in a cooperative way to reach a resolution of an encounter, it might not be an easy task. Wooldridge [178] identified two main reasons that make the effective resolution of multi-agent encounters a difficult task:

1. First of all, in multi-agent systems agents are often build by different entities and have different internal objectives, therefore agents do not usually share a common goal.
2. Secondly, being autonomous by definition agents make their decisions and take their actions dynamically on the environment. Thus, several agents may influence and change at the same time the same portion of the environment.

To overcome these difficulties, researchers have identified the need for coordination. Which, according to Wooldridge [178] implies the definition of a mechanism to manage the interdependencies between the agents and their activities. Jennings [87] in his work presented three reasons to support the importance of coordination in multi-agent systems:

1. **Existence of dependencies between the agents:** interdependency occurs when the actions that agents perform to reach their individual goals are interrelated. This may happen when one agent's decisions have impact on other agents' decisions (e.g. if one agent's actions change the conditions that make the actions of other agents possible). In these cases agents must coordinate their actions in order to make it possible for each one to reach their own goals.

2. **Existence of global restrictions:** if the system is constrained by global restrictions (concerning resources, time, cost and others) the individual agent's actions may not be enough to satisfy the restrictions.

3. **No agent has resources, knowledge or ability to perform a complete task or solve a complete problem:** most tasks and problems need very distinct knowledge and abilities. Thus, a multi-agent system often has agents specialized in different domains or agents that have access to different knowledge.

In addition, Nwana et al. [120] discussed two more reasons related to the gains in efficiency that coordination can bring to multi-agent systems:

1. **Efficiency:** through information exchange and division of tasks coordination can boost performance and efficiency of a multi-agent system. For example even if two agents are able to perform their current tasks they can exchange them if others can perform the tasks more efficiently. In another example an agent can share its resources and information to improve the performance of a task being performed by another agent.

2. **To prevent anarchy and chaos:** multi-agent systems are normally not centralized and due to the complexity and dimension of such systems anarchy can easily emerge. Agents have a local non global view of the environment and do not have a global authority over the system. Agents can interact with every other agent that they see and, if not coordinated, these interactions may lead to confusion and disorder turning the agent society into an anarchy.

In the next sections we will describe some of the best known strategies and protocols for addressing the problems of coordination. We will describe auctioning, negotiation, argumentation, and discuss some coordination strategies that can be applied to multi-agent systems when the agents are cooperative.

2.6.1 Auctioning

Auctions are, due to their simplicity, one of the strategies used when the coordination problem is concerned with allocation of resources or the competition for something exclusive, for example some piece of information that is not common knowledge.

In an auction agents can have two different roles. They can be either the auctioneer, that is, the agent that calls the auction, or be one of the bidders. The idea behind the auction is that the auctioneer has something to offer and calls for an auction to get the best deal that he can among all the interested agents, the bidders.

The notion of "worth" is the central notion that drives the interaction in the auction. The offer proposed has a value, which might be common and fixed among agents or change from agent to agent. The goal of the auctioneer is to maximize the value while the bidders want to minimize it.

Auctions have been widely studied, and therefore several different protocols for resolving the auctions have been proposed. These protocols differ in the way they deal with three different characteristics of the auction:

1. **Number of rounds:** the auction can be only one round, when the bidders make their bid simultaneously, or last for several rounds. In the second case the bidding can be ascending or descending as the value of the bids increases or decreases.

2. **Bidding:** the bidding can be open-cry, which means that all participants have information about the current bids, or sealed, when the agents don't have information about the others' bids.

3. **Winner determination:** auctions differ in the way they determine the winner and the value he pays for offer. The two best-known possibilities are the first-price and second-price auctions. In the first-price, the winner is the one that makes the highest bid and the value of the offer is the value of that offer, in the second-price the winner is also the highest bidder but the value of the offer is the second highest bid.

There are four types of auctions among the most used: the English auction, which is a first-price open-cry ascending auction; the Dutch, which is an open-cry descending auction; the first-price sealed bid auction, which is a one round auction; and the Vickrey auction, which is a second-price one round sealed auction.

If the participants on the auction are not fair, auctions may have several problems. For example bidders can collude and unite against the auctioneer by dropping the offers to a value lower than the real worth of the auctioned good. On the other hand the auctioneer can lie in relation to the bids received in the case of sealed bid auctions. For example in a second-price auction the auctioneer can fake the second best bid and give it a higher value and, therefore, receive a higher price for the good than he should.

In an auction the offer that the auctioneer first proposes cannot be changed until the end of the process, nor can the bidders have any participation in the design of the proposal. This may not be desirable. To give all participants a chance to change the proposal being discussed, other strategies need to be used, like, for example, negotiation.

2.6.2 Negotiation

In a negotiation, one agent presents a first proposal that is iteratively changed by all participants until a common agreement is reached. The agreement might not be possible if the participants are not willing to make some concessions. Rosenschein and Zoltkin [146] make a differentiation between two different domains in which negotiation occurs. They are:

- **Task-oriented domains:** when the proposals discussed are based on tasks to be performed by the agents.
- **Worth-oriented domains:** when the proposals are based on the worth of the environment state.

Additionally, Rosenschein and Zoltkin proposed a protocol for the negotiation between two agents, which they called the monotonic concession protocol. The protocol consists in a sequence of rounds where each of the participants make, a concession, in turns, on their part of the proposal until it is accepted by the other.

The rules of the protocol are as follows:

1. Negotiation proceeds in a series of rounds.
2. On the first round, both agents simultaneously propose a deal from the negotiation set.
3. An agreement is reached if the one of the agents realizes that the proposal made by the other is at least as good as the proposal it presented.
4. If no agreement is reached, then the negotiation proceeds to another round. The agents cannot make a proposal less preferable (e.g. with less worth) than they previously did.
5. If neither agents makes a concession in the same round then the negotiation ends with no agreement.

Negotiation strategies have clear advantages comparing to auctions, specially when there is a strong need to have an adaptable proposal. But as recognized by some researchers they present particular limitations. Jennings in 2001 [46] emphasized these two:

- **Positions cannot be justified:** when making a proposal the agents cannot justify why they made it (e.g. when a car salesman proposes a price he justifies that price with the good engine the car has and the extras). This is especially important if the agent was delegated to do the job and should justify the agreement to the agent who made the delegation.

- **Positions cannot be changed:** during the negotiation process the function that evaluates the proposals is considered to be static, but it is useful that it may change because, as the negotiation proceeds, new relevant facts might become available.

2.6.3 Argumentation

Argumentation strategies emerged, as an extension to the negotiation approach, to solve the limitations that were identified in the negotiation strategies [166] [48]. In argumentation based negotiation agents can justify their proposals and have a dynamic evaluation function to evaluate the proposals on the table. Gilbert [60] suggested that argumentation could occur in four different modes:

1. **Logical mode:** the logical mode is deductive in nature and resembles mathematical proof (e.g. if you accept A, and A implies B, then we must accept B)

2. **Emotional mode:** the emotional mode occurs when the argumentation appeals for feelings, attitudes, and the like.

3. **Visceral mode:** the visceral mode implies a physical indicator to stress the argument.

4. **Kisceral mode:** the kisceral mode involves arguments that imply intuition or mystical and religious arguments.

The common strategies to resolve argumentation based negotiation make use of dialogues. In these sense dialogues are sequences of arguments. The discussion is started when one agent places the first argument that represents the evidences and facts that it wants the others to believe. Then another agent answers by placing a counter argument that threatens, or

Type	Initial Situation	Main goal	Participants aim
Persuasion	Conflict of opinions	Resolve the issue	Persuade the other
Negotiation	Conflict of interests	Make a deal	Get the best for oneself
Inquiry	General ignorance	Growth of knowledge	Find a proof
Deliberation	Need for action	Reach a decision	Influence outcome
Information seeking	Personal ignorance	Spread knowledge	Gain or pass on personal knowledge
Eristics	Conflict antagonism	Reaching an accommodation	Strike the other party

Table 2.1: Types of dialogues proposed by Walton and Krabe.

invalidates, the first one. The first agent answers the new argument with other argument that threatens it and so on. If an argument stands without any threat the agent that called it wins the argumentation. Walton and Krabe in 1995 [175] suggested a typology for dialogues that is resumed in the table 2.1.

2.6.4 Coordination Strategies for Cooperative Agents

The coordination strategies and protocols described above can be applied to multi-agent systems whether the agents involved are cooperative or not. However, in situations where all the agents in the system are cooperative, some strategies can be used that take advantage of this. Some of these examples are the *joint intentions* strategy proposed by Levesque and Cohen [98] later extended by Jennings [85] [86], agent *mutual modelling* suggested by Genesereth that was based on his work on cooperation without communication [58] and strategies based on *norms and social laws* as Shoham and Tennenholtz suggested [156] [155].

- **Joint Intentions:** when an agent assumes a joint intention with other agents that form the team, it is assuming some responsibility concerning the other members of the team. This responsibility will constrain the agent's behaviour in a way that it cannot drop the common intentions and pursue his own goals independently. For example, if two agents decide they need to move an object together, because neither one could do it on their own, they cannot, while moving the object, change their minds and simply drop it. Jennings[85] defined this responsibility in terms of commitments and conventions. Commitments are viewed as pledges to undertake a specific course of action, while conventions provide a means of managing commitments in changing situations.

- **Mutual Modelling:** the idea behind this approach is to allow agents to build models of the others (theirs beliefs, intentions, and the like) and try to predict their course of actions based on those models. Using these predictions the agent can then coordinate its activities. In such coordination scenarios the use of explicit communication is not required.

- **Norms and Social Laws:** norms and social laws define patterns of behaviour that agents should follow. These norms don't enforce the behaviour but not following them will not be well accepted by the others. For example, while waiting to take a BUS people are expected to form a line and enter the BUS in the same order, there is not a rule that states this but if you don't respect the line and enter before, the others will, at least, look at you in a furious way.

2.7 Agent Communication

Communication is essential in multi-agent systems, since agents in a society or a group need to communicate in order to engage in dialogue when resolving their encounters.

In general, most of the agent communication languages (ACLs) and communication theories are based on the Theory of Speech Acts developed by John Austin [8] and later extended by Searle [150].

John Austin stated that certain classes of speech utterances have characteristics of actions in the sense that they can change the world's state in the same way as physical actions do. As an effect, those utterances can change, or influence, the mental state of others. Those utterances are referred to as *speech acts.* Austin distinguishes three different dimensions of the speech acts:

1. **Locutionary:** the physical action of the act itself. For example, when saying "step back" we make some lip movements and possibly some gestures.

2. **Illocutionary:** the meaning and contextual function of the act. For example, by telling someone to step back, you are warning them of a falling object.

3. **Perlocutionary:** the results of the act upon the listener. For example, alerting the listener to the falling object, in the sense that the listener acknowledges of the imminent danger.

One of Searle's [150] contributions to Austin's speech act theory was a classification of the speech acts according to the intention of the speaker using them. He proposed five different categories for the speech acts:

1. **Representatives:** using a representative act the speaker informs the hearer about some propositions.

2. **Directives:** using a directive act the speaker requests the hearer to take a course of action.

3. **Commissives:** commits the speaker to a course of action.

4. **Expressives:** this act allows the speaker to express a mental state, for example, emotions.

5. **Declarations:** this act changes an institutional state of affairs, for example, declaring war.

The Theory of Speech Acts is the base for the two most used agent communication languages: the KQML (Knowledge Query and Manipulation Language)[49] developed by the Knowledge Sharing Effort Group[47] and the FIPA-ACL[51] developed by the Foundation for Intelligent Physical Agents (FIPA) [52].

These languages categorize the communication messages in several *performatives* each containing a set of parameters that specify for example: the message sender, the receiver, its content, the language of content, the ontology used on the content, the reply instructions, and so on. You can see an example of such messages in figure 2.4 and the complete set of FIPA-ACL *performatives* on table 2.2. The example shows two messages, the first message is sent by the *client* agent to ask the *stock-server* agent for the price of a mouse. The second message is the answer sent by the *stock-server* agent.

1. (ask-one :content (PRICE mouse ?price) :receiver stock-server :reply-with price-question :language KIF :ontology computer-hardware)
2. (tell :sender stock-server :receiver you :in-reply-to price-question :content (= (PRICE mouse) 7))

Figure 2.4: Examples of KQML messages.

Passing Information
confirm, disconfirm, inform, inform-if, inform-ref
Requesting Information
cancel, query-if, query-ref, subscribe
Negotiation
accept-proposal, cfp (call for proposals), propose, reject-proposal
Performing Actions
agree, cancel, propagate, proxy, refuse, request, request-when, request-whenever
Error Handling
failure, not-understand

Table 2.2: FIPA-ACL *performatives*.

2.8 Concluding Remarks

In this section we have introduced the notions of agent and multi-agent system, and described some of their properties. In addition, we have discussed typical interaction scenarios for the agents in such systems and referred to the importance of communication in these interaction processes.

These concepts established the base for the development of the framework to support the creation of virtual environments inhabited by several synthetic characters that we describe in chapter 6.

Furthermore, from this description, that resumes the main issues related to the development of multi-agent systems, we can see that the main focus of the research community is on the efficiency of the resolution of the agent encounters. Therefore, although most of the achievements are based on human studies, questions concerning the believability of the interactions are rarely addressed. For example, in the set of *performatives* proposed by the FIPA-ACL (see table 2.2) we cannot find any *performative* related to the socio-emotional nature of human communication.

Chapter 3

Virtual Environments and Synthetic Characters

3.1 Introduction

The main focus of this research is on groups of synthetic characters that act in virtual environments and their interaction with the user. Thus, before presenting further details on the work conducted in order to develop and test our hypothesis we discuss the notions of virtual environment and synthetic character and the issues of their interaction with the users, focusing primarily on work that is relevant for our argument.

We start our discussion by defining virtual environment and some notions of the user interaction with such environments. Then we proceed to the discussion of synthetic characters and their interaction with the user. In the end we present some examples of systems that make use of synthetic characters that, in one way or another, are related to our work.

3.2 Defining Virtual Environment

To find a definition of *virtual environment* we have necessarily to refer to the term *virtual reality*, because these two notions are so close related. Some authors, such as Kalansky [91], even say that both can be used in the same way. Virtual reality first appeared with the tele-operator systems, which were systems that allowed users to operate a mechanism from a remote location. The main idea behind the tele-operator systems was to create a alternative reality that the user could manipulate as if it was the real thing.

Furthermore, a *virtual reality* system is defined by Vince [174] as a *"system that builds a reality that is believable and yet does not physically exist"*. By believable is meant that there is no way that a user could distinguish it from the real world. In fact, the ultimate goal of virtual reality was to build systems that could recreate the real world as it is. But, it soon became evident that this task was not as easy as people first expected.

The term *virtual environment* appeared later in the decade of 1990 as a wider notion of *virtual reality*. The new notion defined *virtual environments* as systems that create simulated realities that are acceptable, but not necessarily as the real world. This new notion includes systems such as, for example, Multi-User Dungeons (MUDS) [1] that were excluded from the first definition of virtual reality. In addition, Lawrence J. Hettinger [73] stated that the use of virtual environments is *"an approach of the design of human machine interfaces that seeks to provide users with a sense of immersion in a computer generated synthetic world"*.

As the use of virtual environments prospered and the domains where they can be used diverse new terms have emerged to specify certain, more specific, classes of virtual environments, for example:

- **Intelligent Virtual Environments (IVE):** are virtual environments that express some kind of intelligent behaviour[2].
- **Collaborative Virtual Environments (CVE):** are virtual environments that foster the collaboration between the users that interact in the environment.
- **Virtual Learning Environments (VLE):** virtual environments that support learning and teaching activities.

3.3 Interacting with Virtual Environments

One of the most important characteristics of a virtual environment is the level of immersion, or presence, that it induces in the user. Coomans and Timmermanns [35] defined this notion as *"the feeling of being deeply engaged where participants enter a make-believe world as if it*

[1] Multi-User Dungeons (MUDS) are virtual environments that simulate a world where several users interact, usually to play fantasy games and adventures. Normally, the display of the simulation is presented in text and the user interacts with the world through written commands. MUDS are often referred as "text-based virtual reality adventures".

[2] Some literature also refers to IVE as Immersive Virtual Environments.

is real". In addition, Smith et al. [158] argued that the level of immersion is a result of the strength of the mental environment versus the physical environment, the sense of presence depends on the distance between the mental and physical environment. If they are closely related, the person has the sense of "being here", while if they are distant the person has the sense of "being there". If a person has a complete sense of "being there" then s/he is fully immersed in the environment.

In 1965 Ivan Sutherland [165] presented the idea of the ultimate display, a system that could present information to all the senses of the users at a resolution equal or greater than they could discern. Users would, therefore, be unable to distinguish between such a system and reality. He described the first ideas about getting the user's sensors attention and how it is important to achieve a good immersive environment.

A system with the characteristics that Sutherland described is far from existing, but some devices that are a first approach to the ultimate display have been developed. For example, Head Mounted Devices (HMD) or Cave Automatic Virtual Environment (CAVE) systems [36].

In 1992 Sheridan [154] classified the factors that influence the quality of the immersion that a virtual environment may induce in two different dimensions:

1. **The vividness of the environment:** these factors are related to the appearance of objects in the environment, thus, how real they look.

2. **The level of interactivity:** these factors are related to the user's interaction with the environment. For example, how the navigation and the interactions with objects (global selection, local selection, and manipulation) are accomplished.

Furthermore, Bates [16] argued that the presence of inhabitants in virtual environments highly increases the user's sense of presence and immersion. This effect becomes better as the interactions between the user and the other inhabitants become richer. In other words, rich interactions with other inhabitants in a virtual environment provide the user with a good sense of presence and immersion. The other inhabitants in the environment may be other users as well as computer generated characters, which are often referred as *synthetic characters*.

Thus, creating a handful of *synthetic characters* that have the ability to engage in interesting interactions with users became a desirable goal to any virtual environment developer.

3.4 Synthetic Characters

A *synthetic character* is a computer generated character that lives in a virtual environment. It may be an active participant in the environment as it can interact with other things and change the environment's state. *Synthetic characters* can be seen as the life-like "embodied" representation of the concept of *agent* presented in chapter 2. They are autonomous and proactive agents that can have social and emotional behaviour.

Synthetic characters were introduced in virtual environments to enrich the user interaction and engage the user in higher levels of immersion. For that reason one of the key features of a synthetic character is its level of believability. Bates [17] defined a believable synthetic character as a character that provides the illusion of life, and thus leads to the audience's suspension of disbelief. A believable synthetic character is, therefore, one that makes the user "forget" that it is being generated by the computer. The more the user forgets about the computer behind the character the more believable the character is. As with virtual environments, the believability of a synthetic character can be addressed in two different dimensions:

1. **Visual believability:** a character should look believable. Its physical appearance should be such that the user perceives is as real, in particularly, the character's interaction with the environment should be presented to the user in a credible way (e.g. the character should grasp an object as expected according to the object size, weight and structure). Thalmann [169] and Badler [11] have some solid work on visual believability of synthetic characters. Their work is centered on human-like characters but the same ideas can be applied when creating other type of synthetic characters (e.g. animals).

2. **Behavioural believability:** a character should behave in a believable way. This means that the decisions it makes should be coherent with the user's expectations (e.g. the character should act rationally when a rational decision is expected). Believable behaviour includes many features such as social and emotional behaviour.

Both visual and behavioural believability are very important in the creation of a synthetic character. Neither one should be neglected and their levels should be balanced. For example, the user will hardly accept a very realistic character that behaves in a very dumb way, or on the other hand a very smart character that performs its actions in a very unnatural way. Additionally, visual and behavioural believability should be consistent with each other. For instance, a synthetic character should not show visual hints of behaviours that it cannot perform (e.g. a character should not have wings if it cannot fly, or should not wield a weapon if it cannot use it).

Furthermore, a synthetic character's appearance and behaviour should also be coherent with the context where the character is defined. The context may include cultural references, social norms or a background story among other things. For example, the situation described above of a character with wings that cannot fly might be acceptable if the context clearly justifies it.

Furthermore, in 1994 Bates [17] stated that emotions play a central role in the believability of a synthetic character. His arguments were grounded on the expertise that Disney animators gathered as character designers [170]. They said that, to give an emotional dimension to characters, three key issues must be considered:

1. *The emotional state of the character must be clearly defined:* At each instant the character must have a well defined emotional state.

2. *The thought process reveals the feeling:* Characters' actions must reflect the emotional state. Thus, emotions should be considered in the character decision making process.

3. *Convey the emotion to the viewers:* The character's emotions should be perceptible by the users that interact with it, if necessary, by using non realistic exaggerated expressions.

3.5 Interacting with Synthetic Characters

In virtual environments populated with synthetic characters the user's interaction will focus mainly on the interaction with the characters, thus, it is crucial to this interaction to be rich and believable. Furthermore, the user interaction with a synthetic character can be expressed in terms of the level of control that the user has over the character. This control may go from full control of the character's actions to no control at all:

1. **Full control:** the character can represent the user in the environment as an avatar; the user controls every action that the character takes in the environment.

2. **Partial control:** the user can give explicit orders to the character, but the character has the choice to decide the proper way to execute those orders. Sengers [151] presented this notion in the concept of semi-autonomous avatars.

3. **Internal influence:** the user can change some variables that influence the character's decision making process process (e.g. the emotional state). In this way the user has some control over the character but cannot precisely predict its behaviour.

4. **No explicit control (External influence):** the character is completely autonomous, the user can only influence its behaviour indirectly through the environment, for example, by placing objects that capture the character's attention, or by communicating with the character.

The control that a user has over a synthetic character constrains its autonomy, the more control the user has, the less autonomous the character is. In addition, this autonomy is never fully achieved if the environment is shared among several characters and even if the character "lives" alone in the environment its behaviour will be dependent of the environment's state. In 1994, Castelfranchi [32] stressed the notion of agent social autonomy, the autonomy in relation to the influence of other agents, and argued that agents should have at least a certain degree of social autonomy. This minimum autonomy should give, at least, control to the character over its goals and beliefs. Additionally, he stated that complete social autonomy is not useful at all because it will only happen if each character lives in a completely different world. On the other hand, a certain social dependence increases the richness of interactions in multi-agent systems.

Different types of control suggest the use of different kinds of interaction. When the user has some explicit control over a character - it being total, partial, or just some influence - the use of an external physical representation of the character is one prominent approach. Systems like Swamped [90] and the SenToy in FantasyA [121, 123] have successfully explored the use of such tactile control interfaces. In Swamped, the user uses a rubber chicken to control a virtual chicken in a virtual barn with the goal to protect its eggs from the threats of a raccoon. In FantasyA, the user uses the SenToy, a human like puppet that senses emotions, to influence the emotional state of characters that engage in a magic duel. Moreover, when characters are fully autonomous (e.g. the user has no explicit control over them) the user can interact with them through a character that s/he can control, through the use of objects in the environment or through direct communication with the character, for example by engaging in a dialogue.

Synthetic characters that have specific skills to engage users in dialogue's are referred to as Embodied Conversational Agents [31]. By definition, Embodied Conversational Agents (ECAs) are anthropomorphic interface agents that are able to engage a user in a real-time, multi-modal dialogue, using speech, gestures, gaze, posture, intonation and other verbal and nonverbal channels to emulate the experience of human face-to-face interaction. Several systems make use of Embodied Conversational Agents to interact with the user, for example: STEVE [144], REA [28], Greta [124] and the PPP Persona [5].

In 2001 Tomlinson and Blumberg [171] presented some guidelines to build successful interactive virtual environments with synthetic characters. They divided their considerations in six different points:

1. **Interaction Time Scales:** in the first seconds of interaction the character should respond to the user showing that it is aware of the user's presence; in the next seconds the character should show that it is understanding the user's interaction; in the first minutes of interaction the character should present an attractive behaviour to capture the user's interest; finally, the character should establish some long term relations with the user in order to give her/him motivation to return and interact again.

2. **Reciprocal Interaction:** characters should be able to interact with everything that they find in the environment: objects, other characters and the user.

3. **Dynamic Range:** the character's behaviour should have a wide range of possible interactions. A character that always presents the same behaviour is hardly convincing and becomes easily boring.

4. **Life Cycle:** characters should show some evolution and growth with the experience, they can learn from the user's interaction and/or change its appearance accordingly (e.g. getting old).

5. **Allusions to Existing Media:** the interaction should explore the prior knowledge that users can possibly gather from other media. Some cultural aspects should be used in order to help the identification of events in the environment (e.g. using cinematographic clichés).

6. **Impedance Match:** the character as a whole. There should be a balance between the body and the mind; the body should not raise expectations about behaviours that the character cannot perform (e.g. the character should not look smarter than he really is).

The balance between the characters is also important; one character should not become more interesting than the others.

3.6 Social Intelligent Agents

We have already discussed the fact that believability is a key issue concerning the user interaction's with a synthetic character, and that giving an emotional dimension to the character increases its believability [17]. Additionally, social dependence among synthetic characters may increase the richness of their interactions [32], thus, giving some social dimension to synthetic characters may foster their believability.

Reeves and Nass [140], in their studies, concluded that computers are social actors, in the sense that people apply essentially social interactions when interacting with computers. They also suggested that those social interactions follow the same rules as social interactions between people. For example, people prefer polite computers and are polite to computers, people like to be flattered by computers, people assign personalities to computers and prefer to interact with computers that match their own personality.

These results support the idea that a social dimension in a synthetic character is very important to convey its believability. Dautenhahn [38] defined social dimension or social ability of synthetic characters as *the ability that enables and drives synthetic characters to engage in behaviours that would require social skills when performed by humans.* Moreover, social ability and social behaviours are essential when interacting with other characters and/or the users, in particular when collaborating with others. Software agents, or synthetic characters, that have social skills and social behaviour are often referred to as Social Intelligent Agents [38].

Today, developing social skills on synthetic characters is one of the main goals of researchers in the field. The research has focused on several different aspects of social intelligence such as, the ability to engage in dialogues or the ability to build persistent relationships with the users. We have identified some of these aspects to be relevant concerning the goal of our research, which is to build believable working groups formed by synthetic characters and human participants. Thus, in order to have synthetic characters and users engaged in a believable group, it is very important that synthetic characters have the ability to:

1. Create social relationships with the users.

2. Create social relationships with the other characters.

3. Collaborate with the users concerning a task.

4. Behave in a group.

We have conducted a review on the current research on synthetic characters that is somehow related to each of these four items. This review is presented in the next section.

3.7 Case Studies of Synthetic Characters

3.7.1 REA: The Real Estate Agent

In 2001, Bickmore and Cassell [19] presented the notion of *Relational Agents* as a new paradigm for the user's interaction with synthetic characters. The idea was to enhance the interaction by encouraging the user, and the character, to build social relationships that support the interaction. In their test case (REA) they were using the ability to build relationships as a mean to increase the user's trust on the character and its suggestions.

REA (Real Estate Agent) [28] is a real-time, multi-modal, life-sized Embodied Conversational Agent, with a fully articulated graphical body, that can sense the user passively through cameras and audio input. REA is capable of speech with intonation, facial display and hand gestures. The system consists of a large screen where REA and its environment are displayed, a set of two cameras mounted on top of the screen used to track the user head and hand positions, and a microphone used to capture the user's speech input. REA simultaneously processes the organization of conversation and its contents using the FMTB (Functions, Modalities, Timing and Behaviours) model [29]. REA plays the role of a real estate salesperson that interacts with users (potential buyers) to determine their needs, shows them around virtual properties, and attempts to sell them a house (see Figure 3.1).

In the real states sales, it is necessary to build a strong relationship with the user, because of the significant commitment that is at stake and the long duration of the interaction that may take several months. To achieve such a strong relation and increase the user's trust, REA uses small talk strategies. People often engage in small talk to avoid awkward or confusing situations, for example, while waiting for her/his turn at the doctor people engage in small talk to avoid long periods of silence. Additionally, small talk may increase the familiarity and the solidarity between the interlocutors. Familiarity and solidarity are two factors that positively affect the process of building trust, therefore, using small talk as a strategy to increase the

Figure 3.1: REA interacting with a user.

user's familiarity and solidarity may foster the possibility of building trust relationships with the user.

REA's discourse manager incorporates two types of dialogue, small talk and task talk, and has the capability to interleave between them while keeping a record of both. During task talk, REA asks questions about the users' buying preferences, such as the number of bedrooms they need. During small talk REA can talk about the weather, events and objects in her environment, or tell stories about herself or real state. REA chooses a discourse action that minimizes the threat to the user and maximizes her/his trust. The selection of an action at a given time is based on the following factors:

- **Closeness:** REA continually assesses her "interpersonal" closeness to the user, which is a composite representation of the familiarity and solidarity. Each conversational topic has some level of closeness as a pre-requisite. If REA finds that she is not close enough to the user to present a topic, she will engage in discourse actions to raise closeness value (e.g. small talk).

- **Topic:** REA keeps track of the current and past conversational topics. Conversational moves that maintain the topic coherent are preferred. If REA wants to change to a different topic she has the ability to gradually change it, for example: talk about the weather, to talk about the weather in Boston, to talk about the real state in Boston.

- **Relevance:** REA maintains a list of the topics that she thinks the user is more familiar with, and always chooses a topic from that list. The list initially has topics that any person should be aware of, for example the weather. This list of topics is updated during the dialogue with inferences that REA makes with the users' answers.

- **Task Goals:** REA has a list of prioritized goals (e.g. to find the user's housing needs). She preferably chooses discourse moves that help in the satisfaction of these goals.

- **Logical preconditions:** Conversational moves have logical preconditions, which must be satisfied before REA can select such a move. For example, REA cannot ask about the user's current studies before she actually knows that she/he is a student.

This mechanism to build the user's trust has proved to have some effect on users especially in those with a disposition to be extroverts.

REA is an example of a synthetic character that successfully builds and maintains social relationships with users, however she lacks the ability to build such relationships with other characters[3] and her interactions are always one to one. What was relevant for the work presented in this dissertation is the idea to establish social relations with users and the fact that these relations may improve the user's interaction with the synthetic character.

3.7.2 Sam: the Castlemate

Sam [27] is a 3D animated virtual child that acts as a playmate to children and creates stories with them. The character is implemented as an embodied conversational agent that has the ability to tell stories as well as to listen to story bits made by interacting users. The user interacts with Sam through a castle that is shared between the physical and the virtual reality and a plastic figurine. The agent is projected on a screen with the physical half of the toy castle in from of it, the other half of the castle is represented as an image on the screen that gives the idea that the castle continues into the screen (see figure 3.2). The plastic figurine can exist either in the physical world or on the screen, so that the agent and the user can pass it back and forth between their worlds. This exchange is achieved with the help of a magic tower. Whenever Sam or the user places the figure on the tower and closes its door, the figurine will

[3]This was not one of the goals of the system, REA does not need to interact with other characters and they are not defined in REA's environment.

automatically move from one world to the other. The background of Sam is the real-time video of the child's environment, so Sam actually exists in the child's playing space.

Figure 3.2: A child interacting with Sam.

Sam detects a child's presence through a microphone and two pressure-sensitive mats placed in front of the castle. Once the child sits in front of the castle, the agent's gaze is guided by where the child sits on the mats. When the child is playing with the toys and narrating, the system uses audio threshold detection to determine when to give feedback. Sam's stories and other utterances were recorded from a real child.

The system is able to detect when the figurine is on the magic tower and when the door of the tower is closed, so that the child will never see the physical and virtual instantiations of the figurine simultaneously (when the door is opened and Sam has the figurine, it disappears instantly and Sam expresses surprise).

The interaction begins when Sam detects a child nearby. He tries to attract the child's attention by introducing himself and saying sentences like: "Does anyone want to play with me?". Once he gets the child's attention he'll say "Lets make up a story. I'm going to start.". Sam then tells the beginning of a story, moving the figurine around the castle, occasionally looking at the user to draw her/him into the story. Then Sam says: "I'll put the toy in the magic tower so you can tell a story", and places the figurine inside the back of a tower.

When the user opens the door at the front of the tower, s/he finds the toy that Sam had been playing with. S/he can pick it up and begin to tell a story, and while s/he does so, Sam watches her/him and the toy, nodding, smiling, and prompting "what happens next?" when the user hesitates. The user can, at anytime, place the toy back into the tower to give it to Sam. Sam will pick up the toy and build the story a bit further. Sam will give the toy back to the child and the interaction goes on and on.

Children (and even adults) that tried the system interacted with Sam for a considerable amount of time and seemed very pleased with their experience.

The Sam system presented a very interesting way of engaging users in a collaborative activity through the use of shared realities. However, the interaction and the collaboration is only in a one-to-one basis, one user with one character. Sam is also not able to explicitly build social relationships with the users it interacts with. Anyway, this work shares with our work the idea of having users and synthetic characters collaborating in a common task.

3.7.3 Boids: Simulated Flocking Creatures

In 1987 Reynolds [141] developed a computer model to support the simulation of coordinated animal motion such as bird flocks, animal herds and fish schools. He named these generic simulated flocking creatures as *boids*.

The model is based on three dimensional computational geometry to define the *boids*' individual maneuvers, position and velocity. The fundamental aspect of this work is that the group behaviour emerges from the interactions of the individual actions. Thus, the flocking behaviour emerges from the combination of three simple steering behaviours, which described how each individual *boid* should move taking in consideration the positions and velocities of its nearby flockmates:

1. **Separation and Collision Avoidance:** steer to avoid crowding and collisions with nearby flockmates (see figure 3.3).

2. **Alignment and Velocity Matching:** steer towards the average heading of nearby flockmates and attempt to match their velocity (see figure 3.4).

3. **Cohesion and Flock Centering:** steer to move towards the average position of nearby flockmates in order to stay close to them (see figure 3.5).

Collision avoidance and dynamic velocity matching are complementary. Together they ensure that the members of a simulated flock are free to move within the crowded area of

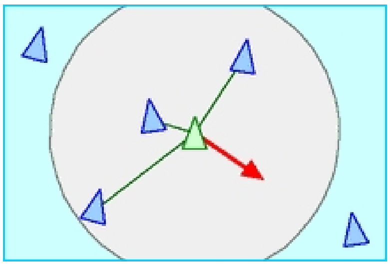

Figure 3.3: *Boids*' separation behaviour.

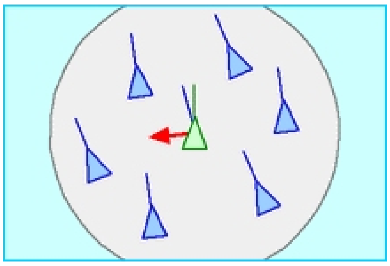

Figure 3.4: *Boids*' alignment behaviour.

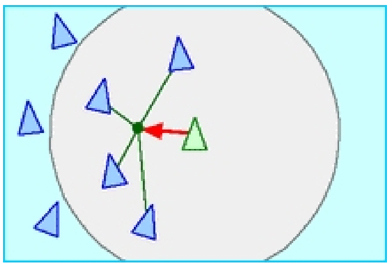

Figure 3.5: Boids' cohesion behaviour.

the flock's interior without bumping against another. Collision avoidance is the urge to steer away from an imminent impact. It is based on the relative position of the flockmates and ignores their velocity. Conversely, velocity matching is based only on velocity and ignores position. Even so, it is a predictive version of collision avoidance: if the *boid* makes a good job of matching velocity with its neighbours, it is unlikely that it will collide with any of them any time soon.

Flock centring makes a *boid* want to be near the centre of the flock. Because each *boid* has a localized perception of the world, the "centre of the flock" actually means the centre of the nearby flockmates. Thus, flock centring causes the *boid* to move closer to the centroid of the nearby boids. This behaviour has the particularity to sometimes split the flock apart, for the reason that as long as an individual *boid* can stay close to its nearby neighbours, it does not care if the rest of the flock turns away. This fact is coherent with the behaviour in real flocks.

The *boids* model was first implemented in the behaviour of the autonomous synthetic characters that represented birds and fishes in the "Stanley and Stella in: Breaking the Ice" movie (see figure 3.6). The results were very impressive and, since then, the *boids* model has been used in many other applications of behavioural animation. For example, the 1992 Tim Burton's film "Batman Returns" contained computer simulated bat swarms and penguin flocks which were created with modified versions of the original *boids* model.

Figure 3.6: A snapshot from Stanley and Stella in: Breaking the Ice.

Boids is an example of a system where the interaction of several synthetic characters produce an emergent group dynamics that results in a believable life-like group behaviour. However, the *boids* creatures lack the ability to build social relations, which we believe to be essential for the interaction with a user, as we expect to demonstrate with our work.

3.7.4 AlphaWolf

Another quite relevant work in the area of synthetic characters if the AlphaWolf [172] system, which is an interactive multi-agent system, developed by the Synthetic Characters Group at MIT Media Lab, that is based on the behaviour of packs of grey wolves. The system simulates a pack of six wolves, three adults fully autonomous and three pups that are directed by the user. Three participants may interact with the system at the same time, howling, growling, whining or barking into a microphone in order to tell their pup to howl, dominate, submit or play. In the course of actions, wolves build and maintain social relationships and are able to perceive the identity of its pack mates, recognizing them as distinct individuals. These relationships are maintained through a mechanism of "emotional memory" based on the somatic marker hypothesis suggested by Damasio [37]. The relationships are explicitly displayed as icons on the user interface, but can also be perceived in the interaction of two wolves. For example, if two wolves have a domination-submission relationship when they meet the dominator will stand steadily while the submissive one will sit in front of it (see Figure 3.7).

Figure 3.7: Domination-submission relationship. The black pup is dominating the white pup.

The mechanism by which the wolves form social relationships involves models of emotion, perception, learning an expression:

1. **Emotion:** Each wolf has a single emotional variable called dominance. A wolf's dominance value changes according to its interactions. For example if the wolf is bitten, its dominance value decreases while, if the wolf is the object of another wolf's submission, its dominance value increases.

2. **Perception:** Each wolf has a unique tag that is known to the entire pack. Through this tag each wolf can univocally identify other members of the pack. Each wolf has, at each instant, an object of attention; it can only perceive events that involve its object of attention. An interaction begins when a wolf A and a wolf B are reciprocally objects of attention of each other.

3. **Learning:** The first time a wolf interacts with another member of the pack it stores an emotional memory of the encounter. The emotional memory consists of three bits of information: (1) the tag of the wolf that was the target of the interaction, (2) the emotional state resulting from the interaction, and (3) a confidence value. Each time the two wolves meet again the wolf's emotional state will change towards the emotional state recorded in the emotional memory based on the confidence of the memory. After each new interaction the memory is revised.

4. **Expression:** The wolf's emotional state affects what it does and how it does it. Dominance affects the way a wolf performs and acts. In high dominance states the wolf will hold its tail and ears erect, while in low dominance states the wolf might walk with its tail between its legs. Additionally, if the wolf is performing autonomously, the dominance value affects its decision mechanism.

From the observation of people interacting with the system, it was concluded that the system highly engaged the users. In the beginning, people felt uneasy about howling, growling, whining, barking in front of others, but after a few moments of interaction they got completely immersed and unconcerned about being watched by others. The wolves seemed plausible to participants and met their expectations, and the relationships they built were mostly understood. People talked about the wolves relationships and it seemed to be one of the reasons for their enthusiasm.

In AlphaWolf the several synthetic characters have the ability to build social relationships among each other, and indirectly with the user through the character that s/he controls. But the user and the synthetic characters do not engage in the resolution of a collaborative task and do not have a strong notion of group.

3.7.5 The SIMS

The SIMS [43] [44] is a computer life-simulation game that allows the player to create characters - *sims* - and control their lives (see figure 3.8). Players can control everything from choosing a career to deciding what to eat for dinner and when to go to the bathroom.

Figure 3.8: A snapshot of "The Sims 2".

The *sims* have different purposes in their lives, they may aspire to a life of popularity, fortune, family, romance or knowledge; these aspirations convey the *sims*' daily fears and desires. Furthermore, each *sim* has a set of eight drives that gradually drop and that it needs to maintain at high levels (see figure 3.9). Those are the *hunger* that represents the need for food, the *comfort* that represents the need to have comfortable fittings, the *bladder* that represent the need to satisfy its primary physiology needs (e.g. use the toilet), the *energy* that represents the need to sleep and rest, the *fun* that represents the need to have fun, the *social* that represents the need to socialize with other *sims*, the *hygiene* that represents its need to be clean and tidy and the *environment* that represents the need to be in enjoyable, quiet, and clean environments.

Figure 3.9: The *Sims*' needs.

In addition, the relative importance of each of the *sim*'s needs depends on the *sim*'s personality. For example, a very neat character will be more susceptible to low hygiene and to messy environments. The *sims*' personality is defined in five different dimensions (see figure 3.10). The first dimension defines the *sim*'s tendency to be sloppy or neat, the second defines its tendency to be shy or outgoing, the third defines the tendency to be lazy or active, the fourth defines the tendency to be serious or playful, and the fifth defines the tendency to be grouchy or nice.

Figure 3.10: The Sims' personality traits.

The synthetic characters in the SIMS are semi-autonomous and will pursue their goals and desires, and try to fulfil their needs without the user's intervention. However, they will usually need the user's guidance to have a successful life. Thus, the role of the user in the game is to help a family of *sims* avoid their fears and make their desires concrete. This can be accomplished by: making their career options, buying them new things, making them fit and learned, and engaging them in relationships with the other *sims*.

The *sims*' relationships represent their mutual social attraction and are built according to their social interactions. In the game, the *sims* can talk, tell jokes, play, flirt, kiss, hug, irritate or fight with each other; and each one of these leads to a change in the social relationship of the interaction's participants. If the interaction is positive then the social attraction will increase, otherwise it will decrease. Social interactions are positive if they are of a positive nature like telling jokes, kissing or hugging, but only if they have a positive response by the receiver (e.g. if a kiss is refused by the other *sim*, the social attraction will decrease instead of increasing as it normally would if the kiss was accepted). Conversely, negative social interactions like irritating and fighting always lead to negative changes in the social relationships.

Furthermore, the set of interactions that are available changes dynamically according to the social relationship state. Which means, for example, that a *sim* can only try to kiss or hug another if their social relationship is positive, not neutral nor negative. In addition, a *sim* can only fight with another if their relationship is negative, not neutral nor positive.

The SIMS is a quite successful computer game that uses its characters' social interactions and relationships as one of the most important means to advance in the game. The synthetic characters in this game are able to build persistent relationships with each other and indirectly with the user if s/he strongly identifies her/himself with one of the characters that s/he controls. However, the characters do not collaborate in the resolution of a common task or build a strong notion of group.

3.7.6 Avatar Arena

The Avatar Arena was developed by Schmitt and Rist [148] as a test-bed for the simulation of multi-character negotiation dialogues. It is used as a demonstrator for the model of group negotiations that these authors have developed. Their model, unlike other approaches to agent-agent negotiation, focuses on the affective issues and dynamics of the social relationships in the negotiation, rather than the protocols and optimization of the process.

In this test scenario, users delegate the task of scheduling their meeting appointments to a virtual agent, their designated avatar. These avatars are later sent to a virtual arena where they talk and negotiate their owners' meetings' times and dates. After that, the negotiation results and process can be displayed to the users in the form of a simulation using embodied conversational characters (see figure 3.11).

According to the authors, the avatars need certain skills in order to be able to perform the task: (1) they need some understanding of the domain; (2) to be aware of the users' personal preferences; (3) to be able to show personality and affective behaviour; (4) to build social relationships with the other avatars; and (5) they need some conversation and negotiation skills.

To handle the domain's knowledge, the authors proposed a simple model for the charac-

Figure 3.11: Snapshot of the Avatar Arena taken during the display of a generated negotiation dialogue. Avatar Peedy, who has a negative attitude towards Genie, is going to make a negative comment on Genies' justification for his unwillingness to accept a certain meeting date proposal.

terization of meeting dates, which included five attributes:

1. time;
2. type of activity (e.g. business meeting, seminar, conference, medical treatment, family celebration);
3. type of person(s) to be met (e.g. family, friends, business partners);
4. type of temporal fixation (fixed or movable);
5. and location of the meeting.

Furthermore, during the negotiation process the avatar needs information about the already scheduled meetings, that can be initially entered by the user. It also needs to know the importance that the user gives to each of these meetings, which will certainly vary from meeting to meeting and from user to user. This information is essential for the decision whether to reschedule or cancel an already scheduled appointment.

The importance of each meeting is not individually set by the user. Rather, it is assessed by the identification of several general user's interests (see figure 3.12) that can be related to the meetings dates, such as: career, hobby, wellness, social contacts, culture, art and education.

For example, if the user expresses a high interest in the dimension "career" then meetings that involve activities that are related to work, such as participating in working meetings and business events, will have high importance values. In addition, a private activity may become more important if the activity is carried out together with workmates or the boss.

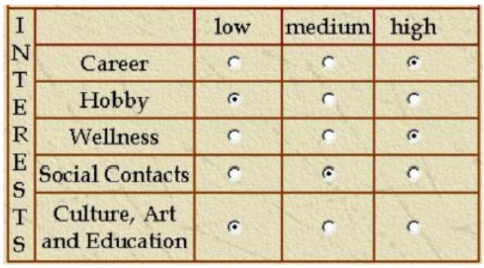

Figure 3.12: The user's individual interest profile.

Furthermore, since one of the main goals of the Avatar Arena is to show the dialogues resultant from the avatars' negotiations to the users, the believability of the conversations is a very important issue. Therefore, in order to achieve this believability, the avatars have different personalities that will be reflected in their conversational behaviour and negotiation style. To model the personality, the authors used three of the dimensions of the Five-Factor Model [109]: the Extraversion, Agreeableness and Neuroticism.

In addition, to increase its believability even further, an avatar's behaviour relies on the concept of a cognitive configuration that can be either balanced or unbalanced. This concept is based on socio-physiological theories of cognitive consistency, such as Heider's Balance Theory [71] and Festinger's theory of cognitive dissonance [50].

These cognitive configurations establish relationships between an avatar that is perceiving the world with another avatar and a third object (see figure 3.13). Furthermore, these configurations are based on the relations of social attraction (or distance) that express an avatar's attitude towards other entities. E.g. Peedy likes Genie, Peedy gives high importance to his career.

Thus, before starting a negotiating dialogue, all participating avatars have a certain social distance to each other that is initially set by the users. Afterwards, they make assumptions about the other avatars' interests values so that the corresponding cognitive configurations are balanced. For example, suppose that the career is an important value for the parrot Peedy who likes Genie, consequently this positive attitude towards Genie makes Peedy believe that the career is also an important value for Genie.

Furthermore, when an avatar discovers a mistake in its assumption about another avatar's interests, it may attain an unbalanced cognitive configuration and eventually change the social distance to the other avatar so that a balance is achieved again.

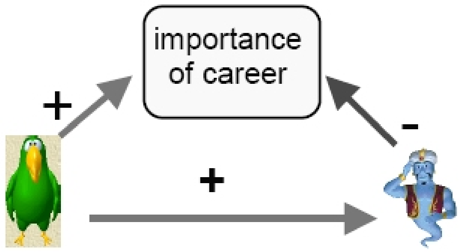

Figure 3.13: Example of an unbalanced cognitive configuration.

For instance, following our last example where Peedy values his career and additionally likes Genie: if, when talking to Genie about career opportunities, it turns out that Genie has no interest at all in issues that relate to career development, then Peedy's cognitive configuration of the situation becomes unbalanced (see figure 3.13). And, therefore, his positive social relation with Genie cannot be maintained.

In conclusion, Avatar Arena engages several synthetic characters in a group activity that involves negotiation of meeting appointments. Furthermore, each character has an individual personality and considers its social relations with the others when engaging in the negotiation dialogues. These relations are initially set and may change through the interaction. However, the user is not directly engaged in the group activity and can only be a spectator of the whole process.

3.7.7 STEVE in Teams

STEVE (Soar Training Expert for Virtual Environments) [145] is a synthetic character, developed by Rickel and Johnson, that has the ability to provide tutoring sessions to students. He has a variety of pedagogical capabilities one would expect of an intelligent tutoring system. For example, he can point out students' errors, and he can answer questions such as "What should I do next?" and "Why?". Furthermore, because he has an animated body, and cohabits the virtual world with his students, he can demonstrate actions, use gaze and gestures to direct the student's attention, and guide the student around the virtual world (see figure 3.14).

This makes STEVE particularly valuable for teaching tasks that require interaction with the physical world.

STEVE was initially used in one-to-one tutoring systems [142] [143] and was later extended to be used in a team training scenario [144] to interact with the students either as an instructor or a teammate.

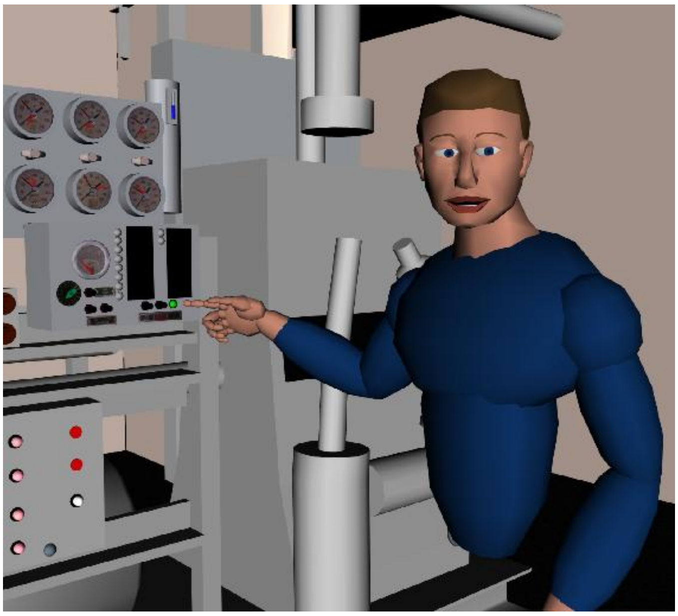

Figure 3.14: STEVE pointing out a power light to a student.

The team training scenario was designed to support the training exercises conducted at a naval training facility. The goal is to train the emergency procedures to handle possible malfunctions that may arise in a ship, for example, a loss of fuel or oil pressure in one of the gas turbine engines that propels the ship. Normally, due to the severity of such casualties, the several members of the crew must work together to successfully handle the problem.

The system can be used by several individuals who are immersed in a virtual environment that simulates the ship, its equipment and the scenario conditions. Each user gets a 3D view of the virtual world through a head-mounted display and interacts with the world via data gloves. In addition, each user is represented by an avatar in the environment, and is accompanied by an instructor (human or a STEVE agent) that coaches them in their role in the task. Humans and agents communicate through spoken dialogue. An agent speaks to a person by sending a message to the person's text-to-speech software, which broadcasts the utterance through the person's headphones.

Teams can be composed of any number of human and agent members. In addition, humans and agents can take the role of a student or a coach as desirable.

Each STEVE agent consists of three main modules: perception, cognition, and motor control. The perception module monitors messages from the system components, identifies relevant events, and maintains a snapshot of the state of the world. This may include the student's current view of the world. The cognition module, implemented in Soar [96], interprets the input it receives from the perception module, chooses appropriate goals, constructs and executes plans to achieve those goals, and sends out motor commands to the motor control module, which accepts the following types of commands: move to an object, point at an object, manipulate an object (about ten types of manipulation are supported), look at someone or something, change the facial expression, nod or shake the head, and speak. The motor control module decomposes these motor commands into a sequence of lower-level messages that are sent to the other software components to accomplish the desired effects.

Furthermore, most of STEVE's abilities to collaborate with students on tasks, either as a teammate or tutor, rely on his understanding of those tasks. Thus, as the scenario unfolds, STEVE must always know which steps are required, how they contribute to the task's goals, and who is responsible for their execution. Tasks are represented in STEVE's knowledge base as a standard plan representation [147] which consists in the description of the task at four different levels (see figure 3.15 for an example):

1. **Task Steps:** each task consists of a set of steps, which are either a primitive action (e.g. press a button) or a composite action (e.g. itself a task). Composite actions give tasks a hierarchical structure.

2. **Ordering constraints:** there may be ordering constraints among the steps, which define a partial order over the steps.

3. **Casual Links:** the role of the steps in the task is represented by a set of causal links [108]; each specifying that one step in the plan achieves a goal that is a precondition for another step in the plan or for the termination of the task.

4. **Roles:** a mapping of task steps to team roles.

When someone (e.g. a human or agent instructor) requests a team task to be performed,

Task transfer-thrust-control-ccs

Steps press-pacc-ccs, press-scu-ccs

Causal links press-pacc-ccs *achieves* ccs-blinking *for* press-scu-ccs;
press-scu-ccs *achieves* thrust-at-ccs *for* end-task

Ordering press-pacc-ccs *before* press-scu-ccs

Roles pacc: press-pacc-ccs; scu: press-scu-ccs

Figure 3.15: An example team task description.

each STEVE agent involved in the task (as a team member or instructor) uses his task knowledge to construct a complete task model. The request specifies the name of the task to be performed and assigns a person or an agent to each role in that task. Thus, starting with the task description for the specified task, each agent recursively expands any composite step with its task description, until the agent has a fully-decomposed, hierarchical task model. Role assignments in the request are propagated down to subtasks until the task model specifies which team member is responsible for each step. Furthermore, during the execution of the task, STEVE agents continually monitor the state of the virtual world, identify which goals in the task model have already been satisfied, and use a partial-order planning algorithm to construct a plan for completing the task.

Moreover, depending on his role in the task, each STEVE agent will use his current plan to guide his behaviour in the team. Thus, if an agent is only serving as a team member, it simply performs its role in the task, waiting, when appropriate, for the actions of its teammates, and communicating with them when necessary. In contrast, an agent serving as an instructor for a human student can demonstrate the student's role in the task, explaining each action it takes, or it can monitor the student as she/he performs the task, answering questions when the student needs help.

The team coordination is achieved through two modalities: first, team members can coordinate their actions by simply observing the actions of their teammates and second, team members may inform their teammates when a goal has been achieved, when they are starting an activity, or when they detect an abnormal condition, by explicitly communicating through speech. In addition, the team speech communication is modelled as explicit speech acts in the task descriptions. Thus, uttering a sentence to a teammate becomes part of the procedure and should be learned by the students.

STEVE agents have been tested on a variety of naval operating procedures and they are able to engage in a collaborative resolution of complicated tasks[4] with human team members. However, all the interactions between the team members are connected to the task and there is not a possibility for deeper social engagement. For example, STEVE agents do not have the ability to model and build social relationships with the other elements of the team.

3.8 Concluding Remarks

In this chapter we have discussed the notion of virtual environment and argued that the use of synthetic characters in such environments may improve the user's interaction experience. Furthermore, we have identified the importance of the believability of such characters when they are built to interact with humans. It was affirmed that giving social skills to a synthetic character is one of the best ways to convey such believability.

Concerning our goal to build a believable group of synthetic characters that engage in the resolution of a collaborative task with humans participants, we have identified that the characters must have the ability to build social relationships with all the group members, including the user, and let these relationships influence their behaviour, in addition to the ability to engage in group interactions and collaborative task resolution.

Finally, we have reviewed some work conducted on synthetic characters that explore the social behaviour of the characters, and relate it with our goal's requisites. We conclude that although several systems have handled some of the requisites, these could not be found all together in the same system. In other words, we did not find a system that made use of synthetic characters, where those characters could successfully interact in a group, collaborate with a user in a common task, and establish dynamic social relations with the several group members (see table 3.1). To build a system with all the mentioned characteristics is the fundamental goal of this work.

[4]The most complex scenario tested involved five team members and a task with several subtask levels, that all together consisted of about three dozen actions.

System	**Soc Rel Users**	**Soc Rel Characters**	**Collaborate**	**Group Behaviour**
REA	yes	no	yes	no
Sam	no	no	yes	no
Boids	no	no	no	yes
AlphaWolf	yes	yes	no	yes
SIMS	yes	yes	no	no
Avatar Arena	no	yes	yes	yes
STEVE	no	no	yes	yes

Table 3.1: The comparison of the several reviewed systems in relation to the requisites identified for our work: (1) the ability to create social relation with the user, (2) the ability to create social relation with the other characters, (3) the ability to collaborate with the user on a common task, and (4) the ability to behave in a group.

Chapter 4

Studies on Group Dynamics

4.1 Introduction

In this chapter we will review some studies and theories of group dynamics that strongly influenced the development of this thesis. Group dynamics is a research field that emerged from the social psychology sciences that studied the nature of groups of people, including the rules that guide the group development, the interaction between the group members and the interaction between several groups.

Different perspectives of group and its process have been taken in these sciences, however, we have concentrated our study on those that define group as a system of members that engage in interaction processes. We made this decision because we considered that these were more easily adapted to a computational system.

We start our review by introducing the concept of group and group process. Then, we describe the different factors that influence the group process, having a special focus on studies of social power and social attraction, since they constitute one of the main drives in the model that we propose for the group dynamics of synthetic characters (see chapter 5).

4.2 The Definition of Group

If we want to study group dynamics it is important that we agree on a definition of group first. However, having a clear general definition of group is not a trivial task as the distinction between a collection of people and a group of people is often not very clear, since groups may emerge in many different social contexts and the interactions of its members may widely vary, for example, from close face to face to more distant and formal. Nevertheless, several different definitions have been proposed [26, 68, 111]. These definitions may contrast concerning the perspective they use to look at the group, which is closely related to the domain where the definition is applied to. However, in general, these definitions are based on the notions of *interaction*, *interdependency* and *mutual perception (or identification)*:

- **Interaction:** one of the differences between a collection of people and a group of people is related to the level of interactions between their members. A group has more interactions and those interactions are more relevant and interrelated. The relevance and interrelation of interactions highly depends on the context of the group.

- **Interdependency:** the members of a group have some interdependency, which means that one member's behaviour affects the other members. This fact is often related to the goal that originated the group in the first place. For instance, all members of the group are interested in the resolution of the same problem and depend on each other to achieve a solution.

- **Mutual Perception (Identification):** all members of a group have the perception of the group. They can identify the group, its members and recognize that they belong to the group themselves.

Furthermore, in 1984 McGrath [111] proposed a definition of group inspired by the mathematical notion of fuzzy sets that defines a group in terms of degree of *groupness*. McGrath identified four different factors that influence the degree of *groupness* of a group of individuals:

- **The number of members that the group contains:** McGrath does not define a boundary for maximum and minimum elements of a group, which was always a difficult problem to attain. But the number of elements is one of the factors that determine the degree of *groupness*. The higher the number of elements of a collection of people, the less probable it is to be considered a group, therefore, the lower is its degree of *groupness*.

- **The interactions between its members:** the occurrence of interactions between the individuals in the group is a determinant factor for the degree of *groupness*. The higher the number of interactions that occur between the elements of the group, the higher is its degree of *groupness*.

- **The history of interactions of the group:** if the same collection of people has had previous interactions it is more likely to be a group and, therefore, it has a higher degree of *groupness*.
- **The probability of future interactions:** if a group of people has a strong probability of having future interactions then the group has a higher degree of *groupness*.

This is a very flexible definition since it does not present an absolute and strict definition of group. Therefore, it easily includes all different types of groups even those that are unlikely to be considered a group according to other definitions, for example, groups that do not have a previous history, like laboratory ad-hoc groups that are created just for research scenarios, or a jury that is formed just to present its judgement on a trial.

Additionally, McGrath applied the same fuzzy logic approach when defining the notion of member of a group. According to McGrath, members of a group are seen as members of a set and not as logical parts of the group. Their belonging is defined as a degree and they can belong to many different groups at the same time having different degrees of *belongingness* for each group.

4.3 The Group Process

The term group dynamics was first introduced by Kurt Lewin [99, 100] in the studies he conducted between 1939 and 1946. From these studies, he formulated the field theory that is known to be the first group dynamics theory.

After the publication of Lewin's studies many other researchers focused their work on the study of group processes and dynamics. Consequently several different theories on group dynamics emerged. Cartwright and Zander [26] in 1968 resumed the theories that contributed most to the group dynamics development and divided them into several categories according to their orientation:

- **Field theory:** the theory developed by Lewin [99], based on the notion of space of life.
- **Interaction theories:** based on the notion of a group as a system of interacting members, developed by Bales [12], Homans [76] and Whyte [176].
- **System theories:** consider the group as a system with observable inputs and outputs. Newcomb [116] was the first to refer to the group as a system on his work on groups as

"communication systems". Furthermore, Miller [113] and Stogdill [162] developed other approaches to groups as systems inspired by biology.

- **Sociometric theory:** developed by Moreno [114], studies the interpersonal choices of the group members.
- **Psychoanalytic theories:** first proposed by Freud [55]. They study the processes of motivation and defence of an individual within a group. This theory has been developed by other authors such as Bion [20], Stock and Thelen [161].
- **Cognitive theories:** studies the knowledge that an individual builds about its social world, and how this knowledge influences her/his behaviour. Festinger [50] and Heider [71] have made some contributions to this theory.
- **Statistics theories:** they study statistical processes for finding the group dynamics concepts, developed by Cattel [33] and Hemphill [72] among others.
- **Formal models:** some authors have attempted to construct formal models based on logic and mathematics in order to deal rigorously with some aspects of groups. Examples of these models can be found on the work developed by French [53].

Our approach to model group dynamics of synthetic characters mainly follows theories classified as system and interaction theories. Nevertheless we also include several ideas from other different perspectives, such as, sociometry and cognitive theories.

In the perspective of system and interaction theories, the group process is composed of a set of interaction processes that occur during an interval of time $t = [t1, t2]$ (see figure 4.1). The interaction processes are affected by preceding factors that are present at time *t1* and result in some consequences at time *t2*. McGrath and others [110, 67, 111] have studied the group process using this perspective and have defined and classified preceding factors, interaction processes and their consequences:

- **Preceding factors:** are variables that influence the interaction processes; they can be categorized in three different dimensions:

    1. *Individual level:* these factors are related to the individual characteristics of each member of the group. They include the individual capabilities and skills, the attitude and motivation and other demographic and biographic traits.

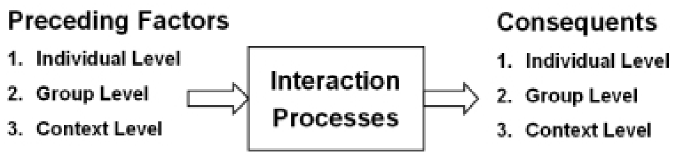

Figure 4.1: The group as a system of interacting members with observable inputs and outputs.

2. *Group level:* these factors relate to the group topology. They include the dimension of the group, its composition, its hierarchy structure, the distribution of influence and power, the role of each member and the organization of the communication along with other sociometric factors.

3. *Context level:* these factors are associated with the nature of the task and the environment where it must be performed. Additional context factors, like inter-group relations and social and cultural factors, are also considered.

- **Interaction processes:** are the interactions and exchanges that occur between the members of the group. These processes can be categorized according to their frequency, form and content.

- **Consequences:** the interaction processes have consequences that change the situation that was initially verified when the processes first began. These changes occur at three different levels:

1. *Individual level:* results that affect the individual members. For example, the individual social condition can change and be facilitated. Furthermore, individuals' attitude and motivation on the group may also change.

2. *Group level:* results that change the group topology. The group structure may change or even emerge if the group was not previously structured. The influence and power distribution can be modified resulting, for example, in changes in the leadership.

3. *Context level:* results in the task level. The task efficiency and execution can change.

Type	**Categories**
Social-Emotional	1. Seem Friendly
	2. Release Tension
Positive	3. Agree
Instrumental	4. Give Suggestion
	5. Give Opinion
Active	6. Give Information
Instrumental	7. Ask for Information
	8. Ask for Opinion
Passive	9. Ask for Suggestion
Social-Emotional	10. Disagree
	11. Show Tension
Negative	12. Seem Unfriendly

Table 4.1: Bales IPA system of categories.

4.4 Interaction Processes

According to Jesuino [89] most of the results that are available on group interaction processes are due to Bales' studies and his associates' [12]. Bales developed a method for group analysis based on the observation of the interactions that occur between the members of a group. He called this method the Interaction Process Analysis (IPA).

Bales concluded that a group, in its process, faces two main problems: the *instrumental problems* which are related to the task, and the *social-emotional problems* which are related to the social and emotional relations of the members. Furthermore, Bales divided the interactions concerning the task problems into *active* and *passive*. *Active* interactions are those where the members give information while *passive* interactions are those where the members receive information. The interactions concerning the *social-emotional problems* can be *positive* or *negative* as they raise positive or negative social-emotional responses on the members.

The IPA system proposes twelve categories for the possible interactions of group members throughout the group process. Table 4.1 resumes the IPA categories and its classification according to the type of problems that they relate to (instrumental or social-emotional).

disagreement Bales defined each of these categories as follows:

1. **Seem friendly:** acts that show solidarity, raise the self esteem of others, raise the status of others, and show general affability. For example, a member may cheer up another when she/he makes an negative contribution for the group.

2. **Release tension:** acts that show tension release and general satisfaction. For example, telling a joke in the beginning of a meeting.

3. **Agree:** acts that show passive acceptance, understanding, concurrence, and compliance. For example, show agreement with a proposal made by another member.

4. **Give suggestions:** acts that provide direction without depriving others of their autonomy. For example, suggest the date for the group's next meeting.

5. **Give opinion:** acts that evaluate. These often begin with ”I think,” ”I believe,” ”It seems to me,”.

6. **Give information:** acts that provide objective information. For example, telling the group the room's availability for scheduling the next meeting.

7. **Ask for information:** acts that ask for objective information. For example, ask for the availability of the meeting room.

8. **Ask for opinion:** acts that ask for evaluation and analysis. For example, ask one member of the group what she/he thinks about having the next meeting on the next week.

9. **Ask for suggestions:** acts that ask for direction and ways of acting. For example, asking the group for possible dates to schedule the next meeting.

10. **Disagree:** acts that show passive rejection or lack of concurrence. For example, show disagreement with a proposal made by another member.

11. **Show tension:** acts that indicate a plea for help, withdrawal, or lack of comfort with the discussion. For example, complain about the current discussion.

12. **Seem unfriendly:** acts that deflate the status and self esteem of others. For example, show hostility towards a member when she/he fails one of the group tasks.

When applying the IPA system to group observation, Bales found that each group has three different problems to solve concerning the task resolution: the *orientation of the task*, the *evaluation of the information* and the *control of the decisions*. Furthermore, he observed that each of these problems cause different levels of interaction throughout the group interaction time. At the beginning, most of the interactions are related to the *orientation of the task* - categories 6 and 7, then the interactions concentrate on *evaluation of the information* - categories 5 and 8, and in the last moments the most common interactions are connected with the *control of the decisions* - categories 4 and 9. Figure 4.2 shows a simple diagram with the distribution of the interactions related to each of these problems over time.

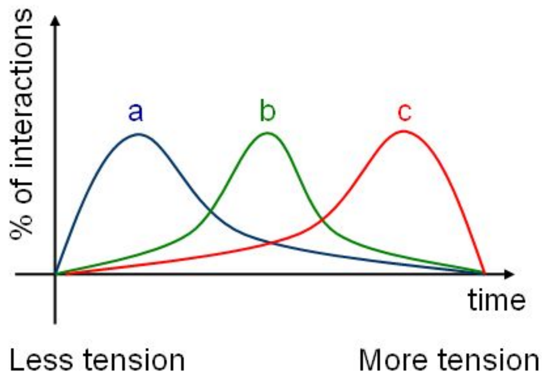

Figure 4.2: The distribution of the IPA system interactions over time. a) orientation; b) evaluation; c) control.

Additionally, Bales observed that, in the beginning of the group process, the tension between the members is usually very low and, therefore, the group makes less use of the social-emotional interactions. As the interactions evolve and the decisions need be made, the tension in the group gets much higher and the social-emotional interactions become more frequent.

Other researchers have studied the IPA system and have drawn some more conclusions regarding the interactions of the members within a group. Hare [68] and McGrath [111] resumed the most important results:

- There are always members with more participation on the group than others.
- Usually, there is a member that takes the initiative of starting most of the interactions. This member is referred to as the *top-initiator*.

- The *top-initiator* addresses the group as a whole, while the other members refer to specific members in their interactions.

- If we order the members of the group by their percentage of interactions in the group, the first one - the *top initiator* - will be involved in almost fifty percent of the group interactions, the second one is involved in one quarter of the interactions and the next members are successively involved in half of the interactions of the previous one. E.g one eighth, one sixteenth, etc.

- The level of interaction of each member highly depends on some factors, like their position in the communication network, their status and role in the group and their competence and motivation.

4.5 Individual Factors

In 1981, Shaw [153] reviewed several studies on the influence that the individual characteristics have on the group process and resumed some of their conclusions. These studies looked at biological, social and psychological aspects of people.

One of the factors studied was the chronological age of people. It was found that older people usually have more participation on the group process and that there is also a tendency for the leader of the group to be the oldest member.

The intelligence and competence of each member also influences the participation on the group. More intelligent and competent people have more interactions on the group and have more influence on the decisions of the group. Additionally, they take the role of the leader more often.

Personality and motivation are other factors that influence the interactions of a member in the group. Highly motivated members interact much more in the group than non-motivated members. Furthermore, the personality affects the way that a person interacts in the group. Its was found that people with higher positive social orientation, in other words that easier relate to others, contribute more to the group.

4.6 Group Factors

Another type of factor that influences the group interactions are those related to the group characterization. This characterization is usually viewed in terms of group *dimension*, *composition* and *structure*.

Group Dimension According to Bales [14] the dimension of the group is a factor that contributes to the disparity in the participation levels of the different members of a group. The more elements a group has, the more probable it is that the group interactions are monopolized by just a few members. For example, in small groups most of the members interact and contribute to the solution, while in large groups only some of the members will interact and contribute actively to the solution.

In addition to the influence that the dimension of the group induces on the way the interactions flow in the group, it also affects the results of the interaction process. The dimension of the group is referred to by Bass and Norton [15] as one of the factors that contribute to the emergence of a leader in an unstructured group. Groups with more elements favour the emergence of a leader. The group dimension can also be one of the factors that determines the task's execution. For example, some tasks might require at least a certain number of elements and, therefore, if the group is too small the task can not be completed.

Group Composition The second relevant variable of the group characterization is the composition of the group. Shaw [153], in his revision, identified two significant dimensions of the group composition: homogeneity/heterogeneity of competencies and homogeneity/heterogeneity of socio-emotional and personality traits.

Groups that are more homogeneous in terms of socio-emotional and personality spend less time on socio-emotional interactions needed to maintain the group's cohesion and, therefore, concentrate more on the interactions relevant to the task's resolution. On the other hand, groups that are incompatible in terms of social and emotional characterization have more problems to achieve the resolution of the task and spend much more time on socio-emotional problems. In addition, members in compatible groups are, generally, more satisfied with the group interactions.

Concerning the homogeneity of competencies, the results show that more heterogeneous groups perform better. They come up with a solution more easily and in general reach a better and more creative solution.

Group Structure According to Collins and Raven [34] the group structure can be characterized in terms of the interpersonal relations of the group members. Furthermore, this structure is usually divided into several different dimensions. According to Jesuíno[89], the most common are:

- the *structure of communication*, which reflects the communication network that connects the group members,
- the *structure of social power*, which reflects the social influence that members may exert on each other,
- and the *structure of interpersonal attraction (or sociometric structure)*, which reflects the social attraction relations established between the group members.

Moreover, associated with the group structure is the concept of social role. Each member in the group may be identified with a social role, which depends on the group structure, and entails her/his *position*, *status* and *social norms* in the group:

- **Position:** Shaw [153] has defined group position as *"the aggregate characterization of the different parts of a group that can be associated with a person, which defines the person's place in the group"*. The concept of position resumes the differentiation of a person in a group concerning all its dimensions. Therefore, for example, defining the position of a member in the group means characterizing the member's social attraction and social power in relation to the other members of the group.
- **Status:** the social role may reflect the member's function in the task and, additionally, have a status, attributed by the other members, that reflects the role's relative prestige in the group. For example, a member that is identified as the head of a group (e.g. the chief of a department) has some prestige that will grant her/him with a good position in the group beforehand.
- **Social Norms:** together with a given social role, some social norms, that define the expected and acceptable behaviour of the members that are associated with that role, are usually found.

Many different social roles may be identified in a group and they highly depend on the group characterization and the context of the task. Nevertheless, there are some invariant roles that can be identified in certain conditions. For example, Bales and Slater [13] observed that people in a group are able to determine two kinds of experts for two different roles:

- the *experimental expert role*, which is attributed to the member who is responsible for most of the ideas and is more concentrated on the task;
- and the *socio-emotional expert role*, which is attributed to the member that is more popular and is more concentrated on the socio-emotional problems and the satisfaction of the group.

Additionally, they observed that, although these two roles may be taken by the same person in the group, this is often not the case.

4.6.1 Social Power

The social relations that are established between the members of a group constitute the main structuring mechanism of the group process. Thus, in order to understand the dynamics of the group interactions, it is imperative to understand the dynamics of these social relations. Furthermore, these social relations are usually divided in two different sets: the relations of social power and the relations of interpersonal attraction. This section (4.6.1) describes the notion of social power and discusses the characterization of different types of social power and their dynamics, while the next section (4.6.2) will describe the notions related to the interpersonal attraction relations.

French and Raven [54] defined the notion of *social power* as the influence exerted by a social agent on a person, where the social agent can be another person, a social role, a norm, a group or part of a group. This influence is defined in terms of the psychological change on the person's perceptions, emotions and behaviours that can be attributed to the action of the social agent.

However, social power only defines a potential influence on the person. On one hand, the person may mobilize some efforts to resist that influence, and on the other, the social agent may decide to use less than his full power in the influence process. Therefore, the social influence exerted on a person results from the strength of power that the social agent induces reduced by the strength of the resistance mobilized by the person.

Furthermore, French and Raven [54] [139] have proposed a characterization of power according to its social source and the relation between the person and the social agent:

1. **Reward Power:** based on the perceived ability to mediate rewards, its strength increases with the magnitude of the rewards, but it also depends on the probability of the administration of the reward. The rewards can be seen as the administration of positive valences or the removal of negative ones.

2. **Coercive Power:** based on the perceived ability to mediate punishments. It is similar to the reward power, and its strength depends on the magnitude of the punishments and the probability of their effective application.

3. **Legitimate Power:** sometimes referred to as organizational authority, it is based on the perception that someone has the right to prescribe given behaviours. This kind of power stands on the notion of legitimacy that involves the common acceptance of the person and the social agent, that the latter has the right to exert influence over the first. This may be due to established social roles like, for example, in parent/son or master/servant relations, but can also be related to established commitments, for example, a promise to help. In addition, the bases for legitimate power are often related to social norms, for instance, in some cultures, the aged are granted the right to determine other people's behaviours.

4. **Referent Power:** based on perceived associations between the person and the social agent, for example, due to an affective relationship. This type of power is based on the desire of the person to become closely attached, or identified, with the social agent or the desire to maintain such attachment if it already exists. Furthermore, the strength of the influence exerted by these means varies according to the strength of such desires. Moreover, this influence is often exerted passively by the social agent and is reflected on self restrictions adopted by the person regarding her/his behaviour. For example, one might adopt a certain attitude to conform to and be accepted by her/his friends or not to offend a beloved one.

5. **Expert Power:** based on the perceived distinctive knowledge, expertness, abilities or

skills attributed to the social agent. Its strength varies with the perception of the extent of that knowledge or abilities, which is usually compared to the person's own expertness.

6. **Information Power:** based on the perceived control of the information needed in order to reach an important goal. This type of power is similar to the expert power, and its strength depends on the amount of control on the information that is attributed to the social agent and how important the information is for the person.

These different types of social power are interrelated and are often combined in the process of social influence. For example, a speaker of renown expertise on a given field may influence the opinion of a listener using the expert power. However, if the listener dislikes the speaker, this influence will be reduced by a negative referent power.

In addition, the process of influence depends on the perception of the person that is being influenced. For example, if a member is attached to a group and he conforms to its norms only because he fears to be ridiculed or expelled from the group, this would be called coercive power. However, if he conforms in order to obtain praise, it is the case of reward power. Furthermore, if the member conforms to avoid discomfort or gain satisfaction independently from the group's responses, then it is the case of referent power. On the other hand, if the member conforms with the majority's opinion based on the respect for the collective wisdom of the group then the influence is due to expert power.

The perception of power also has effects on the group process. For example, according to Hurwitz et al. [81] and Lippitt et al. [101], the members of a group with higher social power are usually more appreciated by the rest of the group and their interactions are more likely to drive the group's behaviour. In addition, they are more attracted to the group and are more satisfied with the group's interactions. They also tend to underestimate the efforts of the members that have less influence in the group, and to use their power on them as a way of self assertion [92] [93].

4.6.2 Interpersonal Attraction

The first studies on the structure of interpersonal attraction were conducted by Moreno[114]. He developed the Sociometry[114] technique that measured the individual interpersonal preferences, through a questionnaire, to attain the structure of the interpersonal attractions in the group. This technique supports the process of identification of some of the group's constructs like, for example, coalitions and the popularity among members.

The social relations of interpersonal attraction define the affective attitude of each member of the group towards the other members. This attitude reflects the affective ties that each member establishes with the others, which can be either positive or negative. For example, the positive ties can be associated with the attitude of liking the other, while the negative ties can be associated with the attitude of disliking.

This interpersonal attraction is not necessarily reciprocal, thus, for example, if person A has a positive attraction for person B, this does not necessarily means that person B has a positive attraction for person A. Nevertheless, some results show that reciprocity in interpersonal attraction relations is often a reality and that it is an important factor for the development of strong social attraction ties [117] [41].

In addition, the relations of interpersonal attraction in a group tend to be balanced. Heider [70] justifies this fact with the need that people have to maintain balanced cognitive configurations. Furthermore, he developed a framework for studying the structural arrangements between social actors and their attitudes, which is referred to in literature as the Balance Theory [71].

The Balance Theory is centred on the concept of a POX triple where P is a person, O another social actor and X and object, which may be a third person, an idea, a rock group, or anything else that P and O both acknowledge. This triple represents a cognitive configuration built by P, which relates P's beliefs in O's attitude towards X and P's own attitude towards X and O. Figure 4.3 represents an example of such a configuration where each line represents one of the attraction attitudes involved: P towards X, P towards O and O towards X.

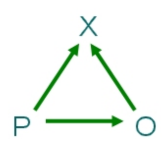

Figure 4.3: An example of a POX cognitive configuration, according to Heider's Balance Theory.

Those attitudes can be either positive (e.g. P likes O, O has a favourable attitude towards X, and so on) or negative (P dislikes O, P disapproves of O, and so on). Moreover, having three different elements with two value relations between them makes it possible to build eight different configurations, which represent the eight cognitive states of the POX triple studied in the Balance Theory. Heider suggested that some of these cognitive configurations are fraught with tensions that make them unstable, particularly if the attitudes are strong. He divided the eight configurations in two different sets: the stable configurations set, shown in figure 4.4, and the unstable configurations set, shown in figure 4.5.

Figure 4.4: Stable POX cognitive configurations.

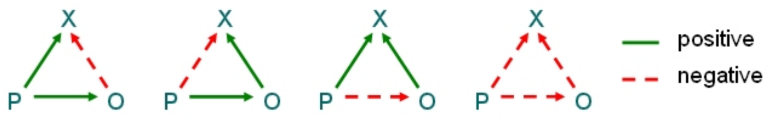

Figure 4.5: Unstable POX cognitive configurations.

The hypothesis in the Balance Theory states that people avoid unstable cognitive configurations and that, if they realize the existence of one, they mobilize their efforts to resolve it and change it to a stable state. For example, suppose that initially a person P has a positive attraction for an object X and another person O, and, additionally, believes that O also has a positive attraction for X (this is the first case in figure 4.4). The cognitive state is balanced so P is fine with the situation. But, later P discovers an unpleasant feature about X and develops a negative attitude towards it. The cognitive state becomes unbalanced (second case in figure 4.5), which creates a certain strain and tension on P. Thus, P will try to recover the balance. According to Heider's model P has two options: (1) change her/his attitude towards O and develop a dislike for O (fourth case in figure 4.4), or (2) reconsider her/his attitude towards X and recover the initial attitude of liking X (first case in figure 4.4). In addition, there is a third option in which P tries to influence O to change her/his attitude towards X (second case in figure 4.4), but this situation is more uncommon because it involves more of P's efforts, and in certain circumstances may be totally impossible. See figure 4.6 for a graphical representation of the example described above.

The Balance Theory has been used in some computational systems of complex social net-

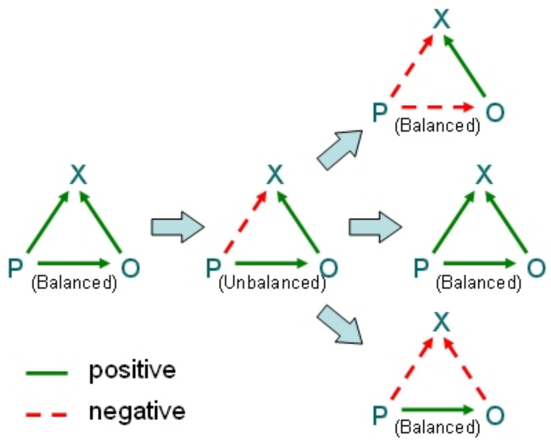

Figure 4.6: Example of the POX configuration dynamics.

works [80] which make use of signed graphs to represent the social structural arrangements [25] [118] [79]. It has additionally been used to convey a more believable social behaviour to social intelligent agents [148].

Interpersonal attraction is often related to the concept of group cohesion, which can be defined in terms of *"the number and strength of mutual positive attitudes among the members of a group"* [125] and the member's attraction for the group [103]. Thus, the higher and more frequent the positive interpersonal attraction relations in the group are, the higher is the group cohesion level.

Furthermore, the level of cohesion influence the group process. For example, members in more cohesive groups interact more often [10] [102] and are more satisfied with the group's interactions [64] [106].

4.7 Context Factors

The context characterization is another factor that influences the group process. This characterization usually includes factors concerning, for example, the physical media of communication, the social, cultural and economic constraints, the intergroup relations and the task. However, we have kept the context factors very simple in our model, therefore, we will not include all these in this description. We will only focus on the task related factors.

4.7.1 The Task

According to McGrath [111] the task is the most important of the context factors, since most of the situations that involve a group usually involve the resolution of one or more tasks. The task becomes even more important if the goal is to study the group's effectiveness and performance. Therefore, it is essential to develop task typologies and descriptions for a better definition and understanding of the critical role of tasks and their associated group processes.

In 1969, Hackman [66] presented a definition of task, which states that a task *"consists of a stimulus complex and a set of instructions which specify what is to be done vis a vis the stimuli. The instructions indicate what operations are to be performed by the subject(s) with respect to the stimuli and/or what goal is to be achieved."*. From this definition we can clearly distinguish three important components of the definition of a task:

1. **Task stimuli:** which define the information and events that the task generates and presents to the members of the group.

2. **Operational Instructions:** which define the procedures to complete the task. It relates the task stimuli with correspondent operations.

3. **Goal Instructions:** which define the task's goals and the requisites to achieve them.

In addition, Hackman points out that the task becomes *"what the group members subjectively define it to be"* [66]. Thus, the task may be redefined individually by the members of the group. The group may also agree with a common redefinition. However, the task, as intended by the person who requested it, is not necessarily perceived as such.

Furthermore, Hackman [65] proposed a typology of tasks that classified tasks into: *production tasks*, which require the production and presentation of ideas or images (e.g. idea generation tasks); *discussion tasks*, which require an evaluation of issues; and *problem-solving tasks*, which require a specification of a course of action to be followed to resolve a problem (e.g. action planning).

One of the most referenced task classifications was proposed in 1972 by Steiner [160]. His classification focused on the task outcome and the imposed constraints that governed the means of accomplishing the outcome. It is based on three different criteria:

1. **Task composition:** it defines if the task can be divided into subtasks. According to this classification, a task can be:

- *Unitary:* if the task cannot be divided into subtasks (e.g. pull a rope).
- *Divisible:* if it is possible to divide the task in several subtasks (e.g. the construction of a house).

2. **Task output:** it defines the kind of output desired by the accomplishment of the task, which can be either:

   - *Maximization:* centred on the quantity of output (e.g. lift the most weight as possible, generate lots of ideas).
   - *Optimization:* centred on the quality of the output (e.g. generate the best idea, reach a targeted performance).

3. **Members' Contributions:** it defines the type of relation between the individual task contributions of the several members. Thus, a task is:

   - *Disjunctive:* if only one of the individual contributions is chosen as the best result (e.g. problem solving, maths calculations).
   - *Conjunctive:* if all the contributions are needed for the result but, the result is limited to the least proficient member (e.g. mountain climbing, running in group).
   - *Additive:* if the result is the sum of all individual contributions (e.g. lift a weight, pull a rope).
   - *Discretionary:* if the result is the weighted sum of all individual contributions, and the weights are chosen by the group (e.g. add an individual contribution to the solution according to the member's role in the group, clean a room).
   - *Compensatory:* if the result is the average of the individual contributions (e.g. money market forecasts, estimate the temperature of a room).

In a more recent work, McGrath[111] integrated the work of several researchers into a framework to classify group tasks. This framework divides tasks in four main groups with two subdivisions, each, according to the behaviour and processes required to complete a task:

1. **Generate:** tasks related to the generation of options.

   - *Planning Tasks:* based on the notion of action-oriented planning, the goal is to generate a plan of action.

- *Creativity Tasks:* based on the creativity of the group, the goal is to generate new ideas.

2. **Choose:** task related to the choice from several options.

   - *Intellective Tasks:* based on the notion that the task has an implied correct answer, the goal is to find that answer.
   - *Decision-making Tasks:* based on the concept that the correct answer is the one preferred by the group, the goal is to reach an agreement to decide the preferred answer.

3. **Negotiate:** the task is related to the resolution of conflicts.

   - *Cognitive Conflict Tasks:* based on cognitive conflicts, the goal is to resolve conflicts of diversion of viewpoints.
   - *Mixed-Motive Tasks:* based on behavioural conflicts, the goal is to resolve conflicts of motive interest.

4. **Execute:** the task is related to the execution of actions.

   - *Contest/Battles:* based on the notion that the task involves several parties and that one can achieve a victory, the goal is to resolve conflicts of power and compete for victory.
   - *Performances:* based on the notion of excelling, the goal is to perform against objective standards.

Additionally, in order to make the relations between the several categories explicit in order to facilitate the establishment of comparisons between them, McGrath designed the Task Circumplex that can be seen in figure 4.7.

Group tasks have been studied and classified with the main purpose of identifying the factors that influence the group process, with emphasis on the group performance. One of the findings concerns the performance of groups on tasks that sum the individual contributions of its members (e.g. Steiner's additive tasks). It was found that, although superior to the individual performance, the group performance is usually much lower than the sum of the individual contributions [95]. This effect, referred as *social loafing*, entails the idea that the

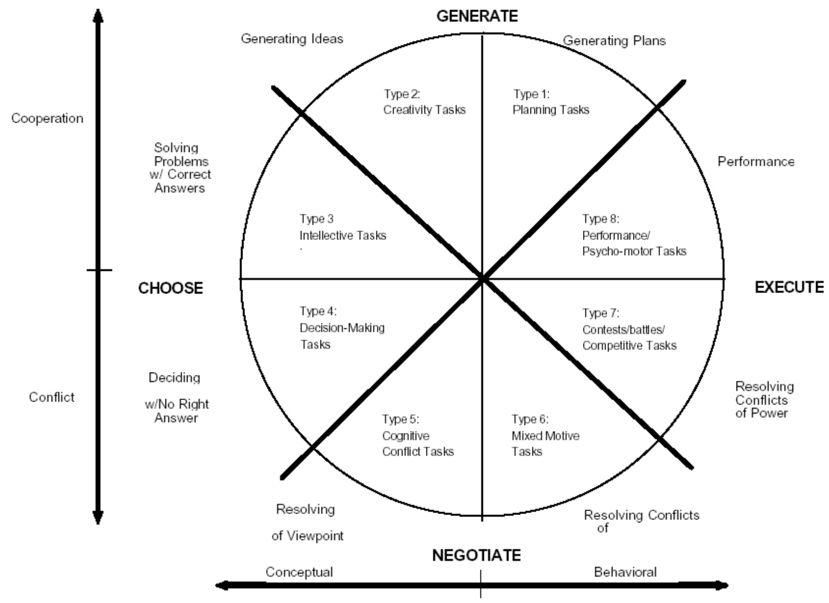

Figure 4.7: The Task Circumplex design by McGrath.

members may not be fully engaged in the group and that they contribute with less than they could. While Steiner [160] justifies this effect with losses in the process of coordination, Holt [75] and others confirm the idea of the low engagement of the members due to low motivation or low identification with the group.

Concerning disjunctive tasks the results show that the performance of the group is not necessarily superior to the individual performance [160]. Moreover, the performance among groups increases with the group size, since it increases the probability of having one member capable of solving the task in the group. Conversely, performance in conjunctive tasks is always inferior to the individual performance, since the performance is limited to the less capable member. Furthermore, the group members are usually more satisfied with the interactions on disjunctive tasks than they take from interactions in conjunctive tasks.

4.8 Concluding Remarks

This chapter presents some notions of group dynamics based on human social psychology studies that supported the creation of our model for group dynamics in groups of synthetic characters described in chapter 5. The group is described as a system of interacting members that are subject to certain factors, which influence their interactions. These factors are classified at different levels: the individual level, concerning the individual characteristics of the members, the group level, concerning the group characterization and structure, and the context level, concerning the group's social and cultural circumstances. We have emphasized the description of the interpersonal relations of power and social attraction, since they form the primary base for group structuring and we rely on them to achieve the structure of our synthetic groups. We decided to exclude the description of the relations of communication, because we have concentrated our studies on virtual environments where the structure of communication of the synthetic characters is very simple.

In conclusion, the studies reviewed in this chapter form a solid base to support the creation of a model to achieve the believability in the group interactions of autonomous synthetic characters that engage the user in a collaborative task.

Chapter 5

Modelling Group Dynamics

5.1 Introduction

To achieve the believability of the interactions of several synthetic characters when engaging in group for the resolution of a collaborative task we have developed a model to characterize and drive such interactions. This model, that we named the Synthetic Group Dynamics Model - SGD Model, was inspired by the theories of group dynamics developed on human social psychological studies, in particular developed by Cartwright and Zander [26], Bales [12] and McGrath [111]. It is based in the principle that a *character must be aware of the group and its members and should be able to build a proper social model of the group and reason with it.* Thus, it defines the type of knowledge that the individual characters should build about the group and its members and how this knowledge drives their interactions in the group. This chapter describes the SGD Model in detail.

5.2 Target Groups

It is not our intention to propose a model that applies to all possible types of groups, nor do we assume that creating such a model would be an easy task. For that reason, we want to explicitly define the types of groups that our proposal applies to, before starting the definition of the model.

First of all, we have focused our study on groups that involve a user with several synthetic characters that perform in a virtual environment. Therefore, these groups engage, at the same time, human and synthetic members.

In addition, the group must be committed to the resolution of collaborative tasks in the virtual environment. Thus, the group interactions should evolve in such a way that makes the resolution of those tasks possible.

We will also only consider groups with few members (small groups) and without a strong organizational structure. Which means that we are not going to deal with groups as crowds or complex organizations and societies of agents.

Furthermore, the synthetic characters in these groups are built as autonomous cognitive agents that need the ability to engage in conversations in order to discuss with the group, and to manipulate objects in the virtual environment (e.g. get, give, use and drop items) in order to perform the necessary actions that will solve the task.

The agents are expected to be socially autonomous as discussed by Castelfranchi [32], in the sense that they have autonomy on their goals and their beliefs. Nevertheless, this autonomy is only partial and relative since the characters' performance is influenced by the other characters in the group and by the environment. Note that the performance can be influenced but never controlled. The agent will always make a decision based on its own goals and beliefs.

Moreover, the user is represented by an agent (avatar) in the system that is not autonomous but completely controlled by the user.

5.3 The Group Process

The SGD Model considers the group as a system composed by several agents, which engage in interaction processes that drive the dynamics of the system (the group process). This system is characterized at four different levels:

1. **the individual level** that defines the individual characteristics of each group member, such as their personality;

2. **the group level** that defines the group and its underlying structure;

3. **the interactions level** that defines the different classes of interactions and their dynamics;

4. **the context level** that defines the environment and the tasks that the agents can perform.

These four levels represent the knowledge that agents should build in order to implement the SGD Model in their behaviour. Furthermore, in addition to this knowledge, the agents' behaviour in the group relies on three different processes:

1. **Classification of the Interactions:** the agent is aware of the actions in the group and classifies them into categories of interaction with specific semantics. For example, in this process the agent interprets if certain actions are helpful for the group of not. This process uses the information on the four levels of the agent's knowledge, specially on the interaction level, that defines the possible categories of interaction, and in the context level that defines how should the actions of the group be interpreted, for example, by means of social norms.

2. **Propagation of the Interaction Effects:** then, based on the identified category, the interaction produce some changes on the knowledge, in particular in the individual and group level. For example, the interaction may change the social relations established between the members that it engages.

3. **Influence of the Agent's Actions:** finally, the agent's perception of the group and its members influences the actions that if performs in the group. For example, if the agent is not motivated it will not try to solve the group's tasks.

5.4 The Individual Level

The individual level defines the knowledge that the agent should build concerning the individual characteristics of each of the group members. This knowledge defines the members' abilities and their personality:

1. **The agent's abilities:** these define the actions that each agent can perform in the environment associated with their levels of expertise (e.g. how good the agent is while performing each of these actions). The set of abilities is important to determine the agent level of expertise in the group, which is an important factor to define the agent's position in the group.

2. **The agent's personality:** we define the agent's personality using two of the dimensions proposed in the Five Factor Model [109]: Extraversion and Agreeableness. We only

consider these two dimensions because they are associated with the ideas of dominant initiative and socio-emotional orientation proposed by Bales [2] while the other dimensions are more related to the task resolution which is not our main focus.

(a) *Extraversion:* is related to the dominant initiative of the agent. It will influence the agent's frequency of interaction.

(b) *Agreeableness:* is related to the socio-emotional orientation of the agent. It defines the type of socio-emotional interactions that the agent will favour. More agreeable agents will favour positive socio-emotional interactions, while less agreeable agents will favour negative socio-emotional interactions.

Each individual can be represented, in the knowledge base, as a set of facts in the form of predicates, like those shown in the following equations (5.1, 5.2, 5.3):

$$Agent(name) \tag{5.1}$$

$$Skill(agentName, skillName, skillLevel) \tag{5.2}$$

$$Personality(agentName, extraversion, agreeableness) \tag{5.3}$$

The equation 5.1 acknowledges the existence of an agent in the system that is identified by a given *name*; while the equation 5.2 represents the knowledge concerning the agents' skills, where *agentName* identifies the agent that has the skill, *skillName* defines the symbolic name of the skill, and *skillLevel* defines the level of expertise of the agent when using the skill; finally, equation 5.3 represents the agent's personality, where *extraversion* and *agreeableness* respectively define the levels of extraversion and agreeableness of the personality of the agent identified as *agentName*.

5.5 The Group Level

The group level defines the knowledge that the agents should build concerning the group and its underlying structure, and additionally the agents' attitude towards the group.

First of all, the group is defined as a set of individuals that follows the definition presented in section 5.4. But, more than just a set, the group is a unique and identifiable entity with an inherent structure.

1. **The group identity:** identification is an important factor in the definition of a group. For that reason the group needs a unique name to allow it to be clearly distinct in the environment and enable the agents to recognize it and refer to it.

2. **The composition:** the composition is the set of individuals that are associated with the group. The composition may change over time as new members may be admitted or excluded.

3. **The structure:** the group structure is defined in different dimensions. According to Jesuino [89] the most common are the structure of communication, the structure of power and the structure of interpersonal attraction (sociometric structure [114]). As we are handling small groups the structure of communication should not be complex, since all characters may communicate directly with each other, thus, we decided not to include it in our model. The group structure is then defined in two dimensions: the *structure of power* that emerges from the members' social influence relations, and the *structure of interpersonal attraction* that emerges from the members' social attraction relations.

The group can be represented, in the knowledge base, as a predicate that follows the definition 5.4. Where, the *identity* defines the symbolic name of the group, and *members* defines a set with the name of the agents that belong to the group.

$$Group(identity, members) \tag{5.4}$$

Furthermore, since the group structure emerges from the social relations established between its members, the group characterization also depends on the definition of these social relations. Which, as said before, can be of two different types:

1. **Social attraction:** these relations define the interpersonal attraction of the members in terms of like (positive attraction) and dislike (negative attraction) attitudes. These relations are unidirectional and not necessarily reciprocal (e.g. if agent A has a positive attraction for agent B, this does not necessarily mean that agent B has a positive attraction for agent A).

2. **Social influence:** relations of influence define relations of power, they quantify the capacity of one agent to influence the behaviour of another. The influence is defined as

the difference between the power that one individual can exert on another and the power that the other is able to mobilize to resist [54].

These social relations can be defined, in the agent's knowledge base, as predicates that follow the two definitions 5.5 and 5.6.

$$SocialInfluence(source, target, value) \tag{5.5}$$

$$SocialAttraction(source, target, value) \tag{5.6}$$

Social relations are directed from one agent, the *source*, to another, the *target*, and are assessed by a *value* which can be positive, zero or negative. For example *SocialAttraction(A,B,50)* denotes that agent A has a positive social attraction for (e.g. likes) agent B with a value of 50.

Furthermore, the agent's social relations in conjunction with its level of expertise determine its position in the group. This position reflects the agent's relative significance in the group which defines how important its contributions are and how well accepted they are by the group. For example, actions performed by agents that have more social influence on the group members have stronger effects on the group process. Thus, the group position defines the agent's relative power in the group, which directly depends on (1) the overall social influence that the agent may exert on the others, (2) the attraction that the others feel for the agent and (3) the agent's relative expertise in the group. The group position may be computed using the formula 5.7, where *SkillLevel(A,G)* denotes the relative skill level of the agent in the group, and *SocAttraction(A,B*) and *SocInfluence(A,B)* denote, respectively, the value of the social attraction that agent A has for agent B and the social influence that agent A exerts on agent B.

$$\begin{aligned} &\forall G, A : Group(G) \wedge A \in Members(G), \\ &Position(A, G) = SkillLevel(A, G) + \\ &\sum_{m \in Members(G)}^{m} SocAttraction(m, A) \\ &+ \sum_{m \in Members(G)}^{m} SocInfluence(A, m) \end{aligned} \tag{5.7}$$

In addition to the relations that agents build with each other, agents also build a relation with the group itself. This relation captures the member's attitude towards the group and supports the notion of membership. Furthermore, this attitude reflects not only on the agent's motivation in the resolution of the group task and its engagement in the group's interactions, but also on its attraction for the group.

1. **Motivation:** The agent's motivation defines the level of engagement of the agent in the group's interactions.

2. **Group Attraction:** the group attraction assesses the level of attachment of the agent to the group. Agents with high attraction for the group are very tied to the group while agents with low attraction for the group are not very attached and can easily leave the group.

Agents may belong to several groups at the same time, and may have different values for their motivation, group attraction and group position for each one. This information can be represented, in the agent's knowledge base, as a predicate (one for each group) that relates these three concepts and the agent to the group, which follows the definition 5.8.

$$Membership(agent, group, motivation, attraction, position) \tag{5.8}$$

In the definition 5.8, *agent* and *group* represent the identifiers of the agent and the group, respectively. And the *motivation*, *attraction* and *position* define the agent's motivation, attraction, and position in this particular group.

Additionally, a reference to the group's current tasks and their goals, and the history of the more relevant events are found in the group level. For example, the group's history may store information about the group's creation and the admission or exclusion of members, the more recent group interactions, or information about the tasks that the group has already completed.

5.6 The Interaction Level

The interaction level describes the knowledge that the agent builds concerning the group's interactions and their dynamics. These dynamics reflects, on the changes that the group's interactions induce in the agent's perception of the group and, therefore, in the knowledge it builds about the group, and in the rules that drive the behaviour of the agent in the group.

The central notion in this knowledge and dynamics is the concept of interaction, which is related to the agents' execution of actions. In fact an interaction only occurs when agents execute actions that can be perceived and evaluated by others. An interaction consist of several actions, either performed simultaneously or in a certain execution pattern, therefore, more than one agent may be responsible for it. In addition, other agents may support the interaction but not be directly involved in its execution. For example, agents may agree with a certain interaction and explicitly show their support for its execution without performing a single action concerning the interaction other than the declaration of support.

Moreover, each interaction has a certain strength in the group that defines its relative importance in the group process. This strength is directly related to the position in the group of the agents that are responsible for the interaction (e.g. those that executed the action that originated the interaction) and to the position of the agents that supported the interaction (e.g. those who agreed with it). Additionally, each interaction may only affect certain members of the group, for example when a member of the group encourages another to perform a task, the effects of the encouragement interaction will only be reflected on the agent that was encouraged.

To summarize an interaction is defined by:

1. an *action* or pattern of *actions* that identify the type of interaction;
2. the set of *performers*, which define the agents that are engaged in the execution of the interaction;
3. the set of *supporters*, which defines the agents that support the interaction without being directly involved in its execution;
4. the set of *targets*, which defines the agents that are affected by the interaction;
5. and its *strength* in the group, which determines the relative importance of the interaction in the group.

During the group process, the several agents, that constitute the members of the group, interact by performing actions in the environment and affecting each other, for example, they may communicate to each other by using speech acts. Then, from their perception of these events, each agent may build a fact in their knowledge base that identifies the occurrence of an interaction. This fact can be defined as a predicate that follows the definition presented in 5.9.

$$Interaction(type, performers, targets, supporters, strength, time) \tag{5.9}$$

Where *type* defines the category of the interaction (see section 5.6.1 for more details), *performers* defines a set with the names of the agents that were responsible for the occurrence of the interaction, *targets* defines a set with the names of the agents that are influenced by the interaction, *supporters* defines a set with the names of the agents that support the interaction, *strength* defines the strength of the interaction in the group, and *time* defines the time of occurrence of the interaction.

5.6.1 The Classification of the Interactions

In order to model the dynamics of the group process we have divided the several possible group interactions into different categories. This categorization is then embedded on the knowledge that the agent has a priori, that will support the agent's process of perception and identification of the interactions when it asserts new interaction facts in its knowledge base.

Furthermore, although the interaction is closely related to the actions that the agents perform, its classification is more than just the classification of the actions themselves. It depends on the actions' results, on the context of the execution, and also on the agents' perception of the group. Thus, for example, the same action can be perceived as a positive interaction to the group by an agent but as a negative by another.

The classification that this model presents was based on the categories that Bales proposed on his IPA system [12]. Thus, it follows the same main distinction between socio-emotional and instrumental interactions, and divides the interactions into positive and negative (see figure 5.1).

On the socio-emotional level we use six categories similar to those presented by Bales. We consider three positive socio-emotional interactions (*agree, encourage and encourage group*) and three negative social emotional interactions that are opposed by symmetry (*disagree, discourage and discourage group*).

- Positive socio-emotional interactions

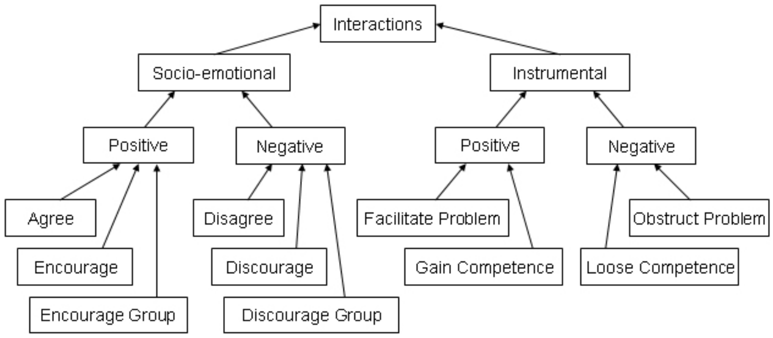

Figure 5.1: The categories of the interactions.

1. **Agree:** This class of interactions show the support and agreement of an agent towards one of the interactions of another agent consequently raising the importance of that interaction in the group.
2. **Encourage:** These interactions represent an agent efforts to encourage another agent, consequently facilitating its social condition (e.g. increasing its motivation).
3. **Encourage Group:** This class of interactions are similar to those in the *Encourage* category but apply to the group itself. These interactions encourage the group and facilitate the group's social structure (e.g. its cohesion).

- Negative socio-emotional interactions

1. **Disagree:** This class of interactions show disagreement of an agent towards one of the interactions of another agent, consequently decreasing the importance of that interaction in the group.
2. **Discourage:** These interactions represent an agent's hostility towards another agent and its efforts to discourage it.
3. **Discourage Group:** This class of interactions are similar to those in the *Discourage* category but apply to the group itself. These interactions discourage the group and raise the entropy of its social structure.

The categories proposed by Bales at the instrumental level focus mainly on speech acts. And, in addition, there is not a clear connection between the instrumental interactions an the task itself. However, in the context of virtual environments, the interactions that are not based on speech acts are very important because the agents may manipulate the objects defined in the environment. Also, the design of the interactions' influence on a problem solving group and its members is easier if the interactions' definition is based on the concept of "problem". Therefore, following these two principles we defined four instrumental interactions: two positive (*facilitate problem, gain competence*) and two negative (*obstruct problem, loose competence*), that do not have a direct correspondence in the IPA instrumental categories.

- Positive instrumental interactions

  1. **Facilitate Problem:** This class of interactions represents the interactions of an agent that solves one of the group's problems or facilitates its resolution.
  2. **Gain Competence:** These interactions make an agent more capable of solving a problem. This includes, for example, the learning of new capabilities or the acquisition of information and resources.

- Negative instrumental interactions

  1. **Obstruct Problem:** This class of interactions represents the interactions of an agent that complicates one of the group's problems or makes its resolution impossible.
  2. **Loose Competence:** These interactions make an agent less capable of solving one problem. For example, by forgetting information or loosing the control of resources.

5.6.2 The Interactions' Dynamics

The interactions create the dynamics in the group. Such dynamics are supported by the classification presented in the previous section 5.6.1 and is modelled by a set of rules that follow the ideas found in the social psychological theories of group dynamics, like for example, the theory of social power by French and Raven [54] and Heider's balance theory [71].

These rules define how the agent's and the group's state influence the occurrence of each kind of interaction and how the occurrence of each type of interaction influences the agent's and group's state.

5.6.2.1 The Occurrence of the Interactions

The interactions of a member in the group depend highly on the individual characterization of the member as well as her/his perception of the group's state. Thus, the member will interact in a completely different way according to different group situations, for example, in groups with different elements or with different emergent structures. In this section we will describe a set of rules that guide these different behaviours by describing the conditions that are more favourable for the occurrence of each type of interaction.

In general, the frequency of the interactions depends on the agent's *motivation*, *group position* and *personality* [153] [111] [2]. Thus, highly motivated agents engage in more interactions, as well as agents with high extraversion or a good position the group. In turn, agents that are not motivated, with a low position in the group, or with low levels of extraversion will engage in few interactions or not interact at all. These elements of the model are captured by the rule synthesized in the following equation (5.10):

$$\begin{aligned} \forall G, I, A : Group(G) \wedge Interaction(I) \wedge A \in Members(G), \\ Extraversion(A) \wedge GroupPosition(A, G) \wedge \\ Motivation(A, G) \vdash Starts(A, I, G) \end{aligned} \tag{5.10}$$

In equation 5.10, *Group(G)* is a predicate that identifies G as a group, *Interaction(I)* identifies I as a group interaction, *Members(G)* denotes the set with the members of the group G, *Extraversion(A)* denotes the level of extraversion of agent A, *GroupPosition(A,G)* denotes the group position of agent A in the group G, *Motivation(A,G)* denotes the motivation of agent A in group G, and finally, *Starts(A,I,G)* denotes that agent A starts the interaction I in the group G. In addition, the $\vdash$ denotes a relation of causality between the elements on the left and the right of the equation. In this case, it expresses that the agent's extraversion, group position, and motivation affect the probability of the agent to start an interaction in the group.

The agent's personality also defines some of the agent's tendencies for the social emotional interactions [2]. Thus, agents with high levels of *agreeableness* will engage more frequently in positive socio-emotional interactions while agents with low *agreeableness* will favour the negative socio-emotional interactions. This leads us to the rules expressed in the following equations (5.11 and 5.12):

$$\forall G, I, A : Group(G) \wedge SocEmotInteraction(I) \wedge A \in Members(G),$$
$$High(Agreeableness(A)) \vdash Starts(A, I, G) \wedge Positive(I) \quad (5.11)$$

$$\forall G, I, A : Group(G) \wedge SocEmotInteraction(I) \wedge A \in Members(G),$$
$$Low(Agreeableness(A)) \vdash Starts(A, I, G) \wedge Negative(I) \quad (5.12)$$

In equations 5.11 and 5.12, *SocEmotInteraction(I)* is a predicate that identifies I as a socio-emotional interaction, *Positive(I)*, identifies the interaction I as positive, *Negative(I)* identifies the interaction I as negative, *Agreeableness(A)* denotes the level of agreeableness of agent A, and *High(V)* and *Low(V)* are predicates that determine if its input value is high or low. In this case these predicates verify if the levels of agreeableness of the agent are high or low, respectively.

Furthermore, the agent's skills influence the occurrence of the instrumental interactions. Thus, more skilful agents will engage in more instrumental interactions than non skilful agents [111]. This fact is expressed in the following equation (5.13):

$$\forall G, I, A : Group(G) \wedge InstrInteraction(I) \wedge A \in Members(G),$$
$$Skilful(A, G) \vdash Starts(A, I, G) \quad (5.13)$$

In equation 5.13 the *InstrInteraction(I)* identifies I as an instrumental interaction and *Skilful(A,G)* identifies agent A as being skilful in group G.

Moreover, agents with higher *position* in the group are usually the targets of more positive socio-emotional interactions while the agents with lower *position* are the targets of more negative socio-emotional interactions [111][1]. In addition, when an agent is considering to engage in a socio-emotional interaction its social relations with the target are very important. Members with higher social influence on the agent and/or members to which the agent is positively socially attracted will often be targets of positive socio-emotional interactions, otherwise they will often be targets of negative socio-emotional interactions. The next two rules express these tendencies:

[1] Note that an agent has a high position in the group if it has high influence over the others and/or if the others have a high social attraction for it.

$$\begin{aligned}\forall G, I, A, B : Group(G) \wedge SocEmotInteraction(I) \wedge \{A, B\} \subset Members(G), \\ High(Position(B, G)) \wedge High(SocAttraction(A, B)) \wedge \\ High(SocInfluence(B, A)) \\ \vdash Starts(A, I, B, G) \wedge Positive(I)\end{aligned} \tag{5.14}$$

$$\begin{aligned}\forall G, I, A, B : Group(G) \wedge SocEmotInteraction(I) \wedge \{A, B\} \subset Members(G), \\ Low(Position(B, G)) \wedge Low(SocAttraction(A, B)) \wedge \\ Low(SocInfluence(B, A)) \\ \vdash Starts(A, I, B, G) \wedge Negative(I)\end{aligned} \tag{5.15}$$

In equations 5.14 and 5.14 the *SocInfluence(A,B)* denotes the value of the social influence that agent A exerts on agent B and *SocAttraction(A,B)* denotes the value of the social attraction that agent A has for agent B. Furthermore, *Starts(A,I,B,G)* denotes that agent A starts the interaction I in group G with agent B as the target of the interaction.

5.6.2.2 The Influence on the Group's State

When agents get the perception of the execution of an interaction they react to it according to the classification that they internally give to the interaction. These reactions are translated into changes in the perceived state of the group.

The *positive instrumental interactions* will increase their performers' *social influence* on the members of the group, by means of expert and information power [54], as well as its own *motivation*. Which means that any member that demonstrates expertise and solves one of the group's problems or obtains resources that are useful to its resolution, will gain influence over the others. In turn, members that obstruct the problem or loose competence will loose influence in the group and become less motivated[2]. These rules are summarised in equations 5.16 and 5.17.

[2]It can be argued that certain people with certain personality traits become more motivated when they fail to achieve a task, however, this is not the most common behaviour, and, therefore, we did not model it.

$$\begin{aligned}
\forall G, I, A, B : Group(G) \wedge InstrInteraction(I) \wedge \{A, B\} \subset Members(G), \\
Starts(A, I, B, G) \wedge Positive(I) \wedge \\
Motivation(A, G) = m_1 \wedge SocInfluence(A, B) = si_1 \\
\vdash Motivation(A, G) = m_2 : (m_2 > m_1) \wedge \\
SocInfluence(A, B) = si_2 : (si_2 > si_1)
\end{aligned} \tag{5.16}$$

$$\begin{aligned}
\forall G, I, A, B : Group(G) \wedge InstrInteraction(I) \wedge \{A, B\} \subset Members(G), \\
Starts(A, I, B, G) \wedge Negative(I) \wedge \\
Motivation(A, G) = m_1 \wedge SocInfluence(A, B) = si_1 \\
\vdash Motivation(A, G) = m_2 : (m_2 < m_1) \wedge \\
SocInfluence(A, B) = si_2 : (si_2 < si_1)
\end{aligned} \tag{5.17}$$

By their turn, *socio-emotional interactions* are associated with changes in the *social attraction* relations. An agent changes its attraction for another agent positively if it is the target of positive socio-emotional interactions by that agent and, negatively, otherwise. The *encourage* interaction has the additional effect to increasing the target's *motivation* in the group. The next equations (5.18, 5.19, 5.20 and 5.21) resume these rules:

$$\begin{aligned}
\forall G, I, A, B : Group(G) \wedge SocEmotInteraction(I) \wedge \{A, B\} \subset Members(G), \\
Starts(A, I, B, G) \wedge Positive(I) \wedge SocAttraction(B, A) = sa_1 \\
\vdash SocAttraction(B, A) = sa_2 : (sa_2 > sa_1)
\end{aligned} \tag{5.18}$$

$$\begin{aligned}
\forall G, I, A, B : Group(G) \wedge SocEmotInteraction(I) \wedge \{A, B\} \subset Members(G), \\
Starts(A, I, B, G) \wedge Negative(I) \wedge SocAttraction(B, A) = sa_1 \\
\vdash SocAttraction(B, A) = sa_2 : (sa_2 < sa_1)
\end{aligned} \tag{5.19}$$

$$\begin{aligned}
\forall G, I, A, B : Group(G) \wedge SocEmotInteraction(I) \wedge \{A, B\} \subset Members(G), \\
Starts(A, I, B, G) \wedge Encourage(I) \wedge Motivation(A, G) = m_1 \\
\vdash Motivation(A, G) = m_2 : (m_2 > m_1)
\end{aligned} \tag{5.20}$$

$$\forall G, I, A, B : Group(G) \wedge SocEmotInteraction(I) \wedge \{A, B\} \subset Members(G),$$
$$Starts(A, I, B, G) \wedge Discourage(I) \wedge Motivation(A, G) = m_1$$
$$\vdash Motivation(A, G) = m_2 : (m_2 < m_1) \quad (5.21)$$

Encourage(I) and *Discourage(I)* in the equations 5.20 and 5.21 identify I as an *encourage* or *discourage* interaction, respectively.

Agents also react to socio-emotional interactions when they are not explicitly the targets of the interaction. Following Heider's balance theory [71], if an agent observes a positive socio-emotional interaction in an agent that it feels positively attracted to, then its attraction for the performer of the interaction will increase. Similar reactions occur in the case of negative socio-emotional interactions. If in the latter example the agent performed a negative socio-emotional interaction, then the observer's attraction for the performer would decrease. These rules are shown in the following equations (5.22, 5.23, 5.24 and 5.25):

$$\forall G, I, A, B, C : Group(G) \wedge SocEmotInteraction(I) \wedge \{A, B, C\} \subset Members(G),$$
$$Starts(A, I, B, G) \wedge Positive(I) \wedge SocAttraction(C, A) = sa_1 \wedge$$
$$High(SocAttraction(C, B))$$
$$\vdash SocAttraction(C, A) = sa_2 : (sa_2 > sa_1) \quad (5.22)$$

$$\forall G, I, A, B, C : Group(G) \wedge SocEmotInteraction(I) \wedge \{A, B, C\} \subset Members(G),$$
$$Starts(A, I, B, G) \wedge Negative(I) \wedge SocAttraction(C, A) = sa_1 \wedge$$
$$High(SocAttraction(C, B))$$
$$\vdash SocAttraction(C, A) = sa_2 : (sa_2 < sa_1) \quad (5.23)$$

$$\forall G, I, A, B, C : Group(G) \wedge SocEmotInteraction(I) \wedge \{A, B, C\} \subset Members(G),$$
$$Starts(A, I, B, G) \wedge Positive(I) \wedge SocAttraction(C, A) = sa_1 \wedge$$
$$Low(SocAttraction(C, B))$$
$$\vdash SocAttraction(C, A) = sa_2 : (sa_2 < sa_1) \quad (5.24)$$

$$\forall G, I, A, B, C : Group(G) \wedge SocEmotInteraction(I) \wedge \{A, B, C\} \subset Members(G), \\ Starts(A, I, B, G) \wedge Negative(I) \wedge SocAttraction(C, A) = sa_1 \wedge \\ Low(SocAttraction(C, B)) \\ \vdash SocAttraction(C, A) = sa_2 : (sa_2 > sa_1) \quad (5.25)$$

The intensity of the interactions' effects described on the previous rules directly depends on the strength of the interaction in the group. For example, *encourage* interactions performed by members with a better position in the group will increment more the target's motivation. So, the interactions' strength depends on the agent's group position, thus, we can say that the group position is a key factor and the main driver for the dynamics of the group. Therefore, to perform well in the group, an agent should take care of its social relations with the other members in the group, since these social relations support its position in the group.

5.7 The Context Level

The context defines the knowledge that the agent builds about the environment where the agents live and the nature of the group's tasks. One of these important definitions is the task model, that allows the agent to interpret the group interactions in terms of their effects on the task, and, therefore, allows the agent to classify them in the instrumental categories.

Additionally, the context may also define some social norms, that will guide the agent in the interpretation of the social-emotional interactions. The social norms define the acceptable behaviours and the misconducted interactions. Thus, supported on these definitions an agent can check, for example, if an action directed to itself is a positive or negative socio-emotional interaction.

These two definitions are important for the group dynamics, however, our model does not deal with any type of tasks or social norms and, therefore, does not define any particular mechanism for the creation of social norms or the definition of group tasks. These are dependent of the context where the model is going to be applied and should be defined for each specific context.

Nevertheless, it is clear that the link with the context would benefit from an extension of the common agent communication languages (ACLs) with two performatives that represent the social emotional interactions, the *encourage* and *discourage* interactions. Thus, the ACL should be extended with a communicative act that represents the efforts to encourage another agent, and a communicative act that represents the efforts to discourage an agent. Furthermore, the performatives for the agree and disagree interactions are usually present in the common ACLs, however, their semantics is only instrumental and should also include the socio-emotional perspective.

5.8 Concluding Remarks

We have discussed that the believability of the interactions of a group of synthetic characters are crucial if the group interacts with users, specially if the group is committed to the resolution of a collaborative task. Therefore, we have developed a model to support the creation of a believable group dynamics on groups of synthetic characters that we described in this chapter.

This model describes the knowledge that each of the autonomous agents, that implement the synthetic characters, should build about the group and its members, and how this knowledge should drive the agents' behaviour, specially concerning the frequency of interaction and their socio-emotional engagement, which are two of the main factors in the dynamics of human groups. Concretely, the group dynamics are achieved through a categorization of the different types of group interactions, and the definition of rules that define the effects that each different class of interaction has on the group's state as well as the interaction's conditions of occurrence.

Furthermore, in addition to the group's state, the group's dynamics also depend on the individual characteristics of the group's members, such as their personality and skills, and the context that involves the group, such as the model of the task that the group needs to perform and the established social norms.

Chapter 6

Supporting the Creation of Multi-agent Virtual Environments

6.1 Introduction

The creation of a synthetic character always has two different facets. The first is the creation of a believable visual representation for the character in the virtual environment, and the second is the creation of mechanisms to implement the character's believable behaviour.

These two facets involve different problems and different perspectives concerning the development of the synthetic character, and, therefore, make use of different tools. For example, we can use the MASSIVE framework [63] [18] to support the development of distributed collaborative virtual environments, and the SOAR [96] and JAM [78] architectures to support the development of the intelligent behaviour of the synthetic characters. However, these tools usually concentrate on the problems of each of the perspectives and do not promote a clean integration of the two, which, consequently, decreases the reusability of the work. Therefore, there is the need for a tool that, not only supports a clean and fast development of multi-agent environments in their two facets, the visual and the behaviour believability, but also establishes a clean integration of the several components to allow them to be reused on different scenarios. For example, if we build a system with several synthetic characters collaborating to solve a problem in a virtual environment and later decided to completely change the visual implementation of the system, we would like to leave the components of the behaviour of the characters untouched and only be concerned with the changes in the visual system.

In this chapter we present a framework that establishes a clean interface between the visual and the behavioural components of multi-agent virtual environments. The first ideas for this framework came from the NIMIS [77] project, where we had the need to create synthetic characters in two different scenarios: (1) to build the synthetic characters for the Teatrix [128] [104] [134] application, and (2) to build a synthetic character that performed as the tutor for the NIMIS desktop and three of its applications.

The framework was used in the development of the application implemented for the evaluation of this thesis that is described in chapter 7.

Furthermore, the ideas of this framework were successfully used in other contexts, for example, in the VICTEC project [122] [168], and are being incorporated in a more general framework that is being developed in our research group [167].

6.2 The Framework

To establish a clean interface between the visual and the behavioural components, in order to interconnect them without creating strong dependencies, the proposed framework defines an abstract layer between them. Thus, conceptually, we divide the components of a system that implements a virtual environment populated with synthetic characters in three different layers[1] (see figure 6.1):

1. the *presentation layer* that contains the components that present the simulated world to the user (e.g. 3D visualization of the world);

2. the *abstract layer* that contains an abstract representation of the simulated world and the synthetic characters (e.g. an abstract model); and

3. the *behaviour layer* that contains the components that implement the intelligent behaviour of the synthetic characters. (e.g planning algorithms)

In this approach, the abstract layer plays the central role since it constitutes the interface between the two other layers and establishes the communication between them. It defines an abstract world, which consists in an abstract representation of the entities that are presented

[1] A similar division can be found in the concept that supports the Model-View-Controller software pattern [94].

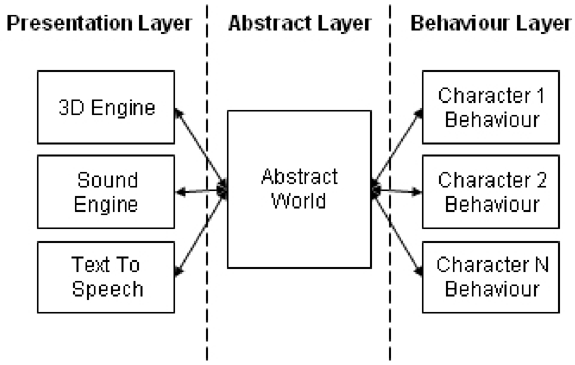

Figure 6.1: A three layer approach of a virtual environment with synthetic characters.

in the environment (e.g. the synthetic characters, and the objects that they manipulate) and their properties and relations.

Additionally, each of these entities may as well be represented in the other layers. For example, each synthetic character will be represented as an entity in the abstract world *(abstract layer)*, with a visual 3D representation in the 3D virtual environment *(presentation layer)* and with a module responsible for its behaviour *(behaviour layer)*, e.g. its mind.

Furthermore, components in the presentation and behaviour layers may only communicate indirectly with each other though the abstract world. For example, consider a component responsible for weather in the virtual environment; by definition it belongs to the behaviour layer, thus, it cannot make changes in the lightning conditions of the 3D environment, which is a presentation component, but it can indirectly request them by changing the state of the abstract model in the abstract layer (e.g. by changing a property that represents the environment's time of the day, from day to night).

This framework only deals with the definition of the abstract layer and with its interaction with the presentation and behaviour layer. Thus, it does not deal with any detail in the construction of the presentational and behavioural components.

6.2.1 The Abstract World

The abstract world defines an abstract model of all the entities that are relevant for the simulation of the virtual environment, which are typically those that the agents and the users may interact with. In other words, an entity is anything that may be identified by the agents. Thus, in order to define an entity one must always define its unique identifier so that there is always a possibility to unambiguously refer to it. Furthermore, the abstract model follows the Entity-Relationship [152] approach to model knowledge, which describes an entity in terms of its properties and relations with the other entities. An entity is defined by a fact that complies with the following definition:

$$Entity(id, properties) \tag{6.1}$$

Where *id* denotes the entity given id and *properties* denote the set of properties that characterize the entity. Furthermore, properties are pairs that map a name to a value, as shown in 6.2. For example, *Property(colour, green)* denotes a property named "colour" with the value "green".

$$Property(propertyName, propertyValue) \tag{6.2}$$

Relations are also identified with a name and relate several entities in facts defined as follows:

$$Relation(relationName, relationProperties, entities) \tag{6.3}$$

Where, the *relationName* denotes the name of the relation, *relationProperties* is the set of properties that characterize the relation and are defined with facts as shown in 6.2. *Entities* is a set of pairs that associate the id of an entity with the set of properties that characterize it in the relation. For instance, the relation between block A and block B that expresses that block A is under block B, as shown in the blocks world [157] example in figure 6.2, may be expressed in the following way (the keywords that identify the relation and the properties were shortened to 'R' and 'P', respectively, in order to facilitate the readability of the expression):

R(onTopOf, {}, {(blockA, [P(position, bottom)]), (blockB, [P(position, top)])})

This relation was named "onTopOf" and was defined without any property. It relates two entities identified as "blockA" and "blockB", both with a "position" property defined, that identifies their role in the relation (e.g. whether they are on top or below). The complete set of expressions that describe the state of the example of the blocks world, shown in figure 6.2, is presented on figure 6.3

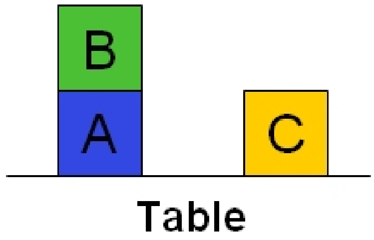

Figure 6.2: An example of the blocks world.

E(blockA, {P(colour, blue)})

E(blockB, {P(colour, green)})

E(blockC, {P(colour, orange)})

E(table, {})

R(onTopOf,{},{(blockA,[P(position,bottom)]),(blockB,[P(position,top)])})

R(onTopOf,{},{(table, [P(position, bottom)]),(blockA,[P(position,top)])})

R(onTopOf,{},{(table, [P(position, bottom)]),(blockC,[P(position,top)])})

Figure 6.3: A definition of the state of the blocks world example, using an Entity-Relationship approach.

Furthermore, agents are a special type of entities that have capabilities to sense and change the abstract world. Thus, in addition to the properties and relations, an agent has a set of sensors and a set of effectors that grant it the sensing and acting abilities. An agent is defined as follows:

$$Agent(id, properties, sensors, effectors) \tag{6.4}$$

Where *id* denotes the agent's unique id, *properties* denotes the set of properties that characterize the agent, *sensors* denotes the set of sensors that the agent may use in order to get information from the world, and *effectors* denotes the set of effectors that the agent can use in order to change the world.

Using the concepts presented here (agent, entity, property, relation, sensor and effector), we may model any particular aspect of a multi-agent virtual environment. However, not all of these should be necessarily included in the model. For example, there may be several entities that are needed in the visual representation of the environment, but which are irrelevant for the agent's interaction. Therefore, only the entities that the agents and the users can interact with, should be included in the abstract model. Likewise, only the relevant properties and relations should be modelled. For example, 3D objects that are only used as decoration in the virtual environment should not be included in the abstract model, since they are irrelevant for the simulation. Removing irrelevant entities from the abstract model simplifies the representation of the information that is handled by the behaviour components.

6.2.2 The Architecture

The contents of the abstract world can only be changed through the action of the agents by the means of their effectors. Therefore, agents are an essential part of the interface between the abstract world and the external components, which are defined in the other two layers (the behavioural and the presentational).

These components communicate with the abstract world by sending messages that request the execution of the agents' effectors and receive information from the abstract world through the messages that are generated by the agent's sensors. All this communication is managed by a central module that is attached to the abstract world: the Communication Manager (see figure 6.4). This module synchronizes the communication by coordinating the messages in two different queues: the *input queue* that stores the requests sent by the external components, and the *output queue* that stores the information generated by the agents' sensors that needs to be redirected to the external components.

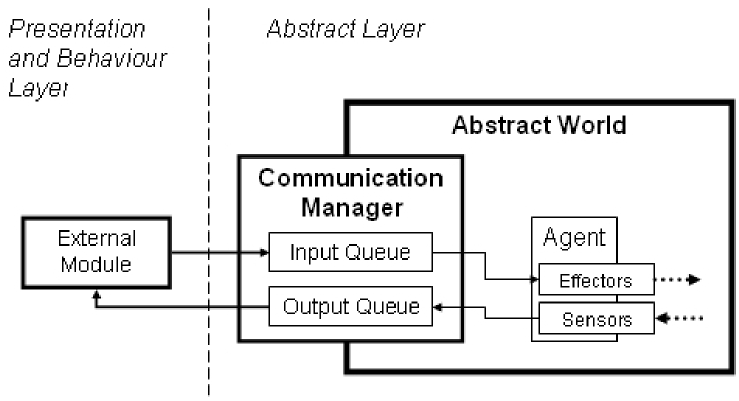

Figure 6.4: The framework architecture.

Furthermore, the type of requests that can be sent to the Communication Manager corresponds to the type of requests that can be asked from the agents concerning their effectors (see section 6.2.5). Thus, there are two different types of requests:

1. The request to *start the execution of an effector*, which asks an agent to start the execution of one of its effectors. This request must define the agent and the effector identifiers, as well as the arguments for the execution. For instance, consider that we have an agent in the blocks world with the goal to pile up the blocks and this agent has an effector named "move-block". In order to have the agent use its effector to move one block from one place to another, its behaviour module must send a request to the Communication Manager stating the agent's name, the "move-block" effector and, additionally, the id of the block that should be moved and the id of the entity (block or table) that is the destination of the movement.

2. The request to *stop the execution of an effector*, which asks one agent to stop the execution of one of its executing effectors. Unlike the *start execution* request, this request only needs to specify the agent and the effector id.

The abstract world, as a dynamic system, runs on successive cycles that are brought forth by a timer. This timer is responsible for starting the update cycle of the abstract world successively in equal time intervals. The update cycle runs on three different and sequential phases (see figure 6.5):

1. **Process incoming messages:** in this phase the messages found in the input queue of the Communication Manager are interpreted and transformed into the correspondent effector execution requests. Thus, agents are notified of the executions that they should start and of those that they should stop.

2. **Update Agents:** this phase updates the agents, what more concretely means that the execution state of the agents' effectors is updated. This may involve three different things (check section 6.2.5 for more details):

    (a) First, the effectors that were requested to start will be started with the given arguments.

    (b) Second, the effectors that were in execution and that were requested to stop are stopped.

(c) Finally, the execution of the effectors that are still in execution is stepped forward.

3. **Process outgoing messages:** in this phase the messages that are found in the output queue of the Communication Manager are sent to the correspondent addressees.

Figure 6.5: The abstract world update cycle.

6.2.3 Events

The activity in the abstract world generates events that can be used by the agents' sensors to build their perceptions (see section 6.2.4 for more details). These events are, in most cases, generated in four situations:

1. When the execution state of an effector is changed (e.g. the effector has started, or stopped; or has finished or failed).

2. When the properties of the entities are assigned to new values or when they are first created.

3. When the relations established between the entities change (e.g. a new entity is added or removed, a property of the relation has changed).

4. When new entities are added to the abstract world.

In addition to these situations, events may as well be generated within the execution of effectors. For example, one way to implement a communication effector that sends a given message from one agent to another (e.g. a "tell" effector) is through the use of events. The "tell" effector, in its execution, may generate an event containing the message that is then caught by the receiver's sensors.

Figure 6.6 presents a scheme that resumes the possible situations that generate events. These situations are always a result of either a new request from the external modules or the execution of an effector.

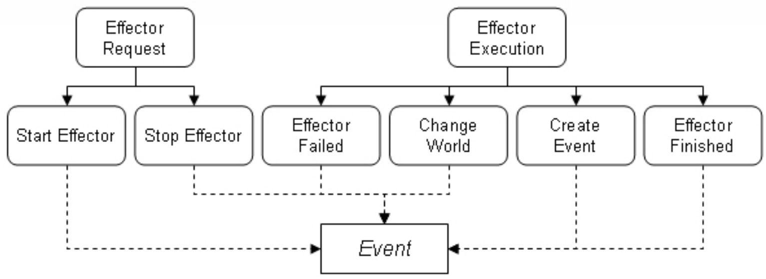

Figure 6.6: Situations that may generate events in the abstract world.

Furthermore, except for the events that may be generated in the effectors' execution, which are user-defined, the framework classifies the possible events in eleven categories:

- **Effector Started:** these events are raised when the execution of an effector has just started. The event contains information with the id of the agent that started the execution, the id of the effector and the execution arguments.

- **Effector Finished:** these events are raised when the execution of an effector has successfully ended. The event contains information with the id of the agent that was executing the effector, the id of the effector and the execution arguments.

- **Effector Failed:** these events are raised when the execution of an effector can no longer continue. The event contains information with the id of the agent that was executing the effector, the id of the effector and the execution arguments.

- **Effector Stopped:** these events are raised when the execution of an effector was deliberately stopped. The event contains information with the id of the agent that was executing the effector, the id of the effector and the execution arguments.

- **Entity Added:** these events are raised when a new entity is added to the abstract world. The event contains information with the id of the new entity.

- **Entity Removed:** these events are raised when an entity is removed from the abstract world. The event contains information with the id of the entity.

- **Property Changed:** these events are raised whenever the value of an entity's property is changed. The event contains information with the id of the entity that holds the property, the property name, and the old and new values.

- **Entity Added to Relation:** these events are raised every time an entity is added to a relation. The event contains information with the name of the relation, the id of the entity that was added and the list of ids of the other entities previously involved in the relation.

- **Entity Removed from Relation:** these events are raised each time that an entity is removed from a relation. The event contains information with the name of the relation, the id of the entity that was removed as well as a list with the ids of the other entities involved in the relation.

- **Relation Property Changed:** these events are raised whenever the value of a property of a relation is changed. The event contains information with the name of the relation, the name of the property that changed, the new and old values, and a list with the ids of the entities involved in the relation.

- **Relation Entity Property Changed:** these events are raised whenever the value of an property that characterizes an entity in a relation is changed. The event contains information with the name of the relation, the id of the entity, the name of the property, the new and old values, and a list with the ids of all the entities involved in the relation.

6.2.4 Sensors

Sensors represent the mechanisms that agents use to process the information coming from the world. Their role is to translate the information from the abstract world events into meaningful and relevant perceptions for the agent's external control modules (e.g. its mind).

When first created, sensors register themselves in the abstract world for the type of events (check section 6.2.3) that they want to receive. Each sensor may be registered in more than one type of event. Afterwards, every time an event of the desired type takes place in the abstract world, the sensors are notified and are requested to handle the event.

Sensors are always enabled by default, however, they may be disabled if needed. Disabled sensors stop receiving event notifications from the abstract world.

The activity of a sensor starts when it is notified with a specific event. Then, the event is processed in a sequence of procedures that create the sensor output, its perceptions. Figure 6.7 shows this sequence, which can be described in the following steps:

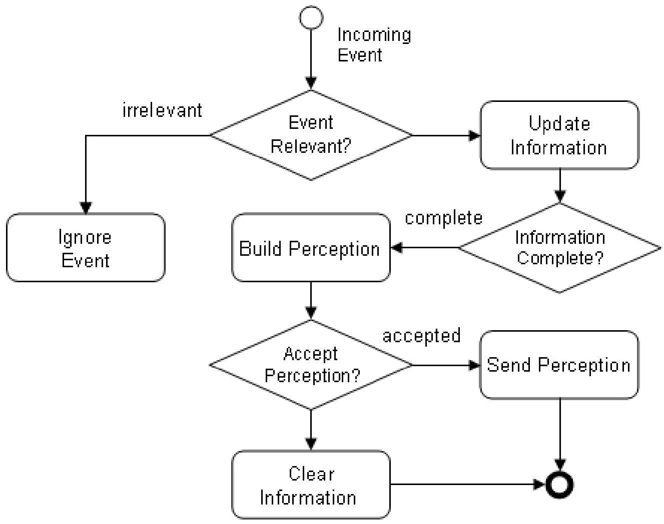

Figure 6.7: The sensors activity diagram.

1. **Check Event:** First the sensor checks if the event is relevant according to the current focus of the sensor. For example, a sensor that is registered to receive property changed events may focus only on a given property and will, therefore, ignore the events related to the other properties (e.g. a colour sensor that monitors only the colour of the blocks in the blocks world example).

2. **Update Information:** If the event is relevant then the sensor updates its local information that supports the construction of its perceptions.

3. **Check Information:** The sensor may require information from more than one event, which means that the information that it gathers from the current event may not be enough to generate a perception. For example, the sensor may be built with the purpose to notify patterns of actions, like for instance, a sensor that counts the number of times that one agent repeats a given action (e.g. counts the number of movements that successively place a block in the table).

4. **Build Perception:** If the information is sufficient, the sensor generates the perception. For example, if an agent has successively moved three blocks to the table and then it moves one block to the top of another, the "move to table" sequence has ended and the sensor may generate the perception that informs the successive execution of three movements to the table.

5. **Check Perception:** However, before sending the perception, the sensor filters its contents and checks if it is acceptable. For example, the sensor that counts successive block movements to the table may be built with the restriction to only notify the perceptions that involve an odd number of executions. Thus, it will not notify, for instance, the occurrence of four successive movements to the table.

6. **Send Perception:** Finally, if the perception is acceptable it is trusted to the Communication Manager that will send it to its addressee (e.g. the perception is sent to the agent's mind).

Furthermore, the framework defines a set of basic sensors that, together, handle all the events presented in section 6.2.3:

- The *Effector Execution Sensor* handles the events related to the state of the effectors' execution (Effector Started, Effector Finished, Effector Stopped, Effector Failed) and creates a correspondent perception for each one. The perception includes the same information provided by the event.

- The *Property Sensor* handles the Property Changed event. Thus, it creates a perception each time a property of an entity is set to a different value. The perception includes the information about the entity, the property, and the old and new values.

- The *Relation Sensor* handles the events that are connected to relations (Entity Added to Relation, Entity Removed from Relation, Relation Property Changed, Relation Entity Property Changed) and creates a correspondent perception any time each occurs. The perception includes the same information provided by the event.

- The *World Content Sensor* handles the events that are related to the addition and removal of entities from the abstract world (Entity Added, Entity Removed). Thus, it generates a perception any time a new entity is added or removed from the abstract world. The perception includes information with the id of the entity.

6.2.5 Effectors

Effectors constitute the mechanism that gives the agents the ability to make things happen in the abstract world. They can change the abstract world's contents by adding new entities, change the value of their properties, create and change relations between the entities or create new events.

To start the execution of an effector, a *start execution* request must be addressed to the agent that controls the effector. This request must provide the specific arguments necessary for the execution. Once started, the effector's execution can be stopped at any time through a *stop execution* request.

The effectors that have started, become part of the list of active effectors that are updated in the abstract world cycle as discussed in section 6.2.2. The effectors remain active until their execution has finished, failed or has been stopped. Figure 6.8 shows a diagram that resumes the procedures executed on each update cycle of an effector. These procedures follow this specific sequence:

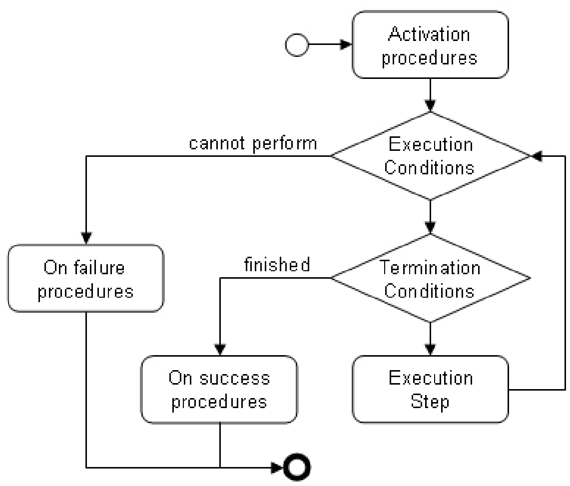

Figure 6.8: The effectors activity diagram.

1. **Activation Procedures:** As soon as the effector execution is requested to start and, just before the placing of the effector in the active effectors list, the effector's activation procedures are executed. These procedures make the effector ready to start its execution. For example, in the case of a "move-block" effector that allows an agent to move one block in the blocks world example, the activation procedures may get, from the abstract world, the current "onTopOf" relations that involve the block that is going to be moved and the destination of the movement.

2. **Check Execution Conditions:** The first thing to do on each update cycle is to check

the effector execution conditions, which determine if the effector can continue its current execution. For example, in the "move-block" effector, the execution conditions determine if both the moving block and the destination block have their tops cleared (e.g. they do not take part in any "onTopOf" relation with the property "position" set to "bottom").

3. **Failure Procedures:** If the execution conditions are not met, then the effector execution fails. Upon failure, some recovery procedures are executed. For example, imagine that an agent has requested the execution of its "move-block" effector to move block A to the top of block B. The effector may start by picking up block A, for instance, by removing its "onTopOf" relations. However, if at the same time another agent moves a third block C to the top of block B then the movement of block A can no longer take place. Thus, the effector recovery procedures restore the initial state of block A, for example, by placing it back in its initial position (e.g. by restoring its old "onTopOf" relation).

4. **Check Termination Conditions:** If the effector may proceed with its execution, then its termination conditions are checked. The termination conditions determine if the effector has successfully completed its execution. For example, in the case of the "move-block" effector, these conditions determine if all of the movement steps have been carried out. The termination conditions only apply to non atomic executions, which take more than one execution step to complete, otherwise the termination conditions are always evaluated as true.

5. **Success Procedures:** Once the termination conditions are reached, the effector's execution is ended and its finalization procedures are executed. For example, when the execution of the "move-block" effector is finished, it may set the "onTopOf" relations to its new values. In the case that where block A is moved to the top of block B, a new relation between the two blocks is created and the respective position properties are set. Block A's takes the value "top" while block B's takes the value "bottom".

6. **Execute step:** The execution of an effector can be a multi step procedure, and therefore, be divided into several update cycles. On each cycle, successive steps are performed until the termination conditions are met. For example, the "move-block" effector may be separated into two steps: the first clears the relations of the block that is going to be moved (e.g. picks up the block), and the second creates the new relation (e.g. moves the

block to the top of the target). In addition, if we consider the movement of the block a more continuous process, then the effector may take several steps, each one updating a bit of the block's current position (e.g. 2D coordinates).

In atomic executions, which take only one execution step, the *Execute Step* procedure is not executed, so the changes to the abstract world are usually performed on the *Success Procedures*.

The framework defines several basic effectors, which include:

- The *Add/Remove Entity Effectors* that add and remove entities from the abstract world (e.g. add a new block to the blocks world). These effectors take the entity id as an execution argument.
- The *Set Property Effector* that sets a new property or changes the value of a given property of an entity (e.g. change the colour of a block in the blocks world example). This effector takes the entity id, the property name and the property's new value as execution arguments.
- The *Add/Remove Relation Effector* that establishes or removes a relation between several entities (e.g. establish an "onTopOf" relation between two blocks). These effectors take the id of the entities and the name of the relation as execution arguments.
- The *Set Relation Property Effector* that sets a property on a specific relation (e.g. sets the strength of the magnetic attraction between two magnetized blocks). This effector takes the relation name, the id of one of the entities involved in the relation, the property name and the new value as execution arguments.
- The *Set Relation Entity Property Effector* that sets a property of an entity in a given relation (e.g. set the "position" property of the blocks involved on the "onTopOf" relation). This effector takes the relation name, the entity id, the property name and the new value as execution arguments.

6.2.6 The View

In the beginning of the description of this framework we presented a conceptual distinction between the several components of a multi-agent virtual environment, namely the distinction between the View components (*the presentation layer*) and the Control components (*the behaviour layer*).

However, although the distinction between view and control supported the development of this framework, this distinction is not explicit in its implementation. Thus, control modules and view modules are equally seen as an external component of the agent (see figure 6.9). This perspective merges the two concepts of view and control in the notion of agent. In fact, each agent has its own view of the abstract world that it builds through its sensors, and has its own participation in the control of the abstract world that it achieves through the execution of its effectors.

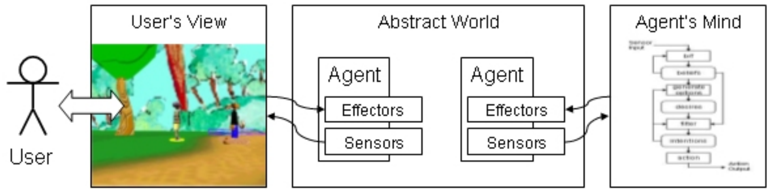

Figure 6.9: The view and the control components as agents in the abstract world.

Nevertheless, the view has a distinct role in the system, since it has the important responsibility to show the abstract world simulation in a meaningful way to the user (e.g. in a 3D visualization) and to enable the user to interact with the simulated world. Thus, the view can be seen as an agent in the abstract world that represents the user.

On the one hand (*the view's perspective*), the view receives, through its sensors, the information about the execution of the agents' actions and about the values of the entities' properties and relations. Then, with that information, it displays the correspondent animations on the synthetic characters that represent the agents and changes the appearance of the graphical entities according to their properties and relations in the abstract world (e.g. change the graphical objects' colour).

On the other hand (*the control perspective*), the view captures and interprets the users' input from their interactions with the representation that they get from the view (e.g. the system's user interface) and translates them into changes in the abstract world, through the use of its effectors. For example, by interacting with the view, the user may add new entities to the abstract world, or change their relations and properties. In addition, the user may also communicate and request actions from the other agents in the system.

Finally, we would like to add that by using this approach, that sees the view as an ordinary agent in the system, there is not a compromise of the view, or, the presentation components, to be visual at all. Thus, for example, the user's interface may be supported only by sound and tactile components as used by unsighted people. Furthermore, several distinct presentation components may be developed for the same system to support the interaction of people with different needs. Which means that, for instance, we may develop two different views to support the interaction of both sighted and unsighted people, and thus, allow these two kind of people to interact in the same virtual environment.

6.3 Supporting Groups of Agents

The multi-agent framework presented in the previous section (6.2) was first developed without considering the possibility of the agents to perform jointly in a group. Nevertheless, it was easily extended with a few new concepts that we describe in this section.

First, we needed a mechanism to define groups. For this, we used the ability to create relations that the framework already supported. Thus, a group is defined as a relation established between several agents, which, additionally, and in order to identify it as a group, must have a property named "group" set with the name of the group (see equation 6.5).

$$R(rName, \{P(group, groupName)\}, \{(agent1, [])...(agentN, [])\}) \tag{6.5}$$

Furthermore, public information related to the group can be stored in the form of properties of the group relation. For example, these properties may store information about the collaborative goal of the group. Likewise, these properties may characterize the state of the group, for instance, they can assess its cohesiveness.

In addition, the properties that characterize the agents in the relation convey their attitude and their role in the group. For example, this type of properties may be used to express the agent's motivation and attraction for the group. However, we would like to alert that the information stored in the relation is shared with all the other entities of the abstract world, including the other group member's. Therefore, only common knowledge should be stored in this way. For example, the motivation in the group is a subjective and personal state of each agent, thus, if the agent does not explicitly desires to share its motivation state it should not be stored as a property in the relation.

As an example, imagine a group formed by two agents, agent A and agent B, that have the common goal to collaboratively arrange the blocks world example in figure 6.2 in a way that blocks A, B and C get on top of each other in the following order: C on the table, B on top of C and A on top of B. In addition, agent A is highly motivated to perform the task while agent B is not motivated at all. One possible way to define the group described above is as follows:

R(r1,{P(group, "the block team"), P(goal, [C,B,A])},
{(agentA, [P(motivation,high)]), (agentB, [P(motivation,low)])})

Additionally, we have defined five basic effectors to handle groups in the framework, which are extensions of the basic effectors that handle relations (e.g. Add Relation Effector, Set Relation Property Effector and Set Relation Entity Property Effector). These effectors allow the agents to create groups, add new members to groups, and set properties to groups, among other things. A description of these five effectors is presented below:

- The *Create Group Effector* allows an agent to create a group with itself as its single member. This effector needs the name of the group as an execution argument. The execution results in the creation of a relation with a property "group" set to the name specified in the execution arguments and the agent as its unique member. The relation id is automatically generated by the system, since it is irrelevant for the definition of the group.

- The *Add Group Member Effector* allows an agent to add a new member to one of its groups. This effector takes the name of the group and the id of the agent that will become the new member as execution arguments. Note that only an agent that already belongs to a group may add new members to it.

- The *Remove Group Member Effector* allows an agent to remove one agent from one of its groups. This effector takes the name of the group and the id of the agent that will be removed as execution arguments. Note that only an agent that belongs to the group may remove members from the group. In addition, an agent may remove itself from one of its groups.

- The *Set Group Property Effector* allows one agent to set a property in the group relation. This effector takes the name of the group, the name of the property and the value to be set as execution arguments. Note that only an agent that belongs to the group may set properties on it.

- The *Set Group Member Property Effector* allows an agent to set a property that characterizes itself in the group. This effector takes the name of the group, the name of the property and the value to be set as execution arguments. Note that the agent may only set a property that characterizes itself and not the others.

Another feature needed in the context of groups is the agents' ability to perform joint executions. This feature is essential in order to engage agents in the resolution of conjunctive and additive tasks (see section 4.7.1), like for example to allow two agents to move a heavy object from one place to another together. Thus, to fulfil this need we have extended the concept of effector presented in section 6.2.5 to define a special type of effector, the *Joint Effector*.

Joint effectors represent actions that can be performed simultaneously by several agents. Their execution is a combination of several individual contributions and can not usually start if some of these contributions are not present. For example, the joint execution for carry a heavy object from one place to another can not start until two distinct agents volunteer themselves to perform the task.

A joint execution always has an associated commitment phase, where the several participants declare their intentions to perform the joint action. This commitment phase follows the same ideas defended on the joint intentions framework proposed by Jennings [85] [86].

Moreover, agents, when committing to a joint execution, must also declare the effector that they are going to use during such execution, which can be different from agent to agent. For example, in the case of the two agents moving a heavy object from one room to another, they are committed to use their "carry-object" effector to transport the object. However, a third agent may also join in the execution and commit itself to use its "open-door" effector, which helps the execution by opening the door that separates the two rooms. Thus, the three agents are committed to a joint execution but with different individual contributions. For the purpose of this discussion the effectors that are committed by the agents are referred to as the joint effector's child effectors.

To support the commitment process we have defined two auxiliary effectors in the framework:

1. The *Join to Execution Effector* that an agent may use to commit itself to a joint execution. This effector needs the id of one of the agents already committed to the joint effector and the id of the joint effector in order to identify the correct joint execution. And, additionally, it needs the id of the effector that the agent is going to use in the joint execution. If no other agent have yet previously committed to the execution, then the agent can use its own id to identify the joint execution, which creates a new instance of a joint execution.

2. The *Withdraw from Execution Effector* that conversely to the *Join to Execution Effector* allows an agent to remove its commitment from a joint execution. This effector needs the id of one of the agents committed to the joint effector and the id of the joint effector in order to identify the correct joint execution.

The joint effector may be started by any of the agents that has committed to its execution. Furthermore, while executing, other agents may as well join or withdraw their commitments. However, if these changes do not comply with the execution's specific constraints, the joint execution may be compromised. For example, when two agents are carrying a heavy object, neither one can lose the object without making it fall on the floor.

When the joint effector is started it is placed, as any other effector, in the active effectors list of the abstract world to be updated in the world cycle. However, its child effectors committed are not included in this list. Their update cycle will be conducted by the update cycle of the joint effector. Therefore, they are started and stopped whenever the joint effector is started of stopped.

Furthermore, the joint effector follows the same execution procedures that are discussed in section 6.2.5, but with some nuances that we describe below:

1. **Activation Procedures:** when the activation procedures of the joint effector are executed, it will first call the activation procedures of its child effectors.

2. **Check Execution Conditions:** when considering the execution conditions to determine if the joint effector's execution fails, the child effectors' execution conditions are also checked. Thus, if any of the child effectors fails, the joint execution will consequently fail. On the other hand, if the joint effector fails, either because one of its child effectors failed or because its own execution conditions were not satisfied, all its child effectors will also fail.

3. **Failure Procedures:** when the failure procedures of the joint effector are executed, it will fist call the failure procedures of its child effectors.

4. **Check Termination Conditions:** to determine if the joint execution has successfully finished, the child effectors' termination conditions as well as the joint effector's own termination conditions must be satisfied. Therefore, the joint execution can only finish if all of its child effectors have already finished.

5. **Success Procedures:** when the success procedures of the joint effector are executed, it will first call the success procedures of its child effectors.

6. **Execute step:** when the joint effector performs a step of its execution, it will additionally request each of its child effectors to perform the next step of its own execution.

6.4 Concluding Remarks

In this chapter we have presented a framework that supports the creation of multi-agent virtual environments based on the fact that these systems have different components with two distinct functions: those responsible for the presentation of the simulated world to the user and those responsible for the behaviour of the simulated world.

The framework integrates these two distinct components as agents in the same abstract world and defines a clean interface between the two. The abstract world defines the relevant entities of the system in terms of their properties and relations. The agents are special entities that have the capability to sense and change the abstract world's state through the use of sensors and effectors.

Additionally, we presented mechanisms built on the framework to support the creation of groups of agents. These mechanisms allowed the definition of groups and their characterization and supported the execution of joint actions.

The framework was successfully used in the integration of the behaviour of the synthetic characters and the visual interface of the application that we developed for the evaluation of our hypothesis, see chapter 7.

Chapter 7

Case Study: Perfect Circle, a Collaborative Game

7.1 Introduction

This chapter describes our case study that was developed in order to evaluate the effects of the SGD Model in the interaction of users with a group of synthetic characters.

We start by presenting the results of an essay that identifies the crucial issues that should be taken into account when defining the task to be performed by the group. Then, taking these issues into consideration, we describe a computer game that we developed to create the context of our groups and to support the execution of the collaborative task. This game places the user in the world of *FantasyA* where s/he, together with a group of autonomous synthetic characters, must find a secret item.

We also present some details regarding the development of the game, describing the architecture of the autonomous members of the group and how the SGD Model was integrated to guide their behaviour.

7.2 Choosing the Right Task

According to Mennecke and Wheeler [112], the group's task has a very important role in the group process and should, therefore, be carefully designed before it can be given to the group. This is particularly important for groups used in research studies, because the choice of the task may influence the results. For example, certain types of tasks are more suitable to be solved by groups (e.g. conjunctive and associative tasks) and, therefore, encourage the group interactions. Thus, if we expect the group to perform as such, we should use this type of tasks in our studies and not those that subvert the group process. Furthermore, the group definition itself should drive the task definition since not every type of group works well with every type of task. In addition, the task should be adequate to the group's context. For example, many experiments fail because their designers do not consider the structure of organizations, the task problems interrelatedness, or the differences in the power and status of the group members.

For all these reasons, Mennecke and Wheeler have compiled an essay to guide researchers in the design and selection of their experiment tasks. This essay presents six critical issues for the definition of a group's task:

1. **Appropriateness of the Task to the Subjects:** The task should present a situation within the realm of most of the subjects's experiences and capabilities. The task should be designed so that subjects have or can acquire the skills or knowledge needed to complete the task.

2. **Subjects' Intellectual Engagement:** The task should engage the subjects to complete and solve the task. An effective way to generate high involvement is to give the subjects the perception that they have a stake in the task's outcome.

3. **Control of the Differences in the Subjects' Preferences, Needs and Experiences:** The task definition should create an environment that minimizes the differences between the subjects' individual needs and preferences. The task instructions and stimuli should be given to the subjects in a clear and consistent way, so that individual interpretations of the task do not differ too much.

4. **The Level of Task Complexity:** According to Nunamaker et al. [119] the tasks that subjects solve in laboratory settings are much simpler compared to the tasks that are solved in reality. Thus, the task should be complex enough in order to support the generalization of the experiment's results outside its controlled environment. Additionally, the task complexity is a way to achieve the desired level of engagement of the subjects since tasks that are either too simple or too difficult can have negative effects on the subjects' motivation. One way to control the task complexity is to vary the quantity and quality of information given to the subjects to process the task.

5. **Conjunctive Versus Disjunctive Tasks:** Disjunctive tasks may often be solved by an individual as well or as easily as by a group. Conjunctive tasks, on the other hand, require the participation of all group members in order for the group to successfully find a solution, and, therefore, should be preferred to disjunctive tasks.

6. **Task Typology and Measurement of the Solution Quality:** The task typology must be taken into account when designing the experiment. If the quality of the task solution is one of the important variables to be measured, the task should be designed in a way that facilitates the solution evaluation. For example, following McGrath [111] typology, performance or intellective tasks, which have a correct solution, should be used.

7.3 The Game

7.3.1 Background Story

We have designed our task to be a computer game, the *Perfect Circle*[1], which, as any good computer game, has an introductory background story. This background story is the first factor to potentially increase the user's engagement in the task, which, as we have discussed in the previous section (section 7.2), is one of the important issues when designing an experiment on group dynamics.

Our story was inspired by the initial concept created by Carlos Martinho on the world of *FantasyA*. This concept has already been used as a basis for the creation of the *FantasyA* computer game that was developed as a test tool for the SenToy [135, 74].

FantasyA is a world where gods have been banned to imprisonment into the essence of gemstones. The gemstones have, then, been shattered and scattered all over the world, rendering the god's powers to their bare minimum. The gods barely have the strength to be in their eternal sleep in a far elemental plane. Therefore, their existence was unseen to Men.

In our game the user is teleported to *FantasyA* in the age where some Men are already aware of the powers within the gemstones and have learned the first steps of their manipulation. These men, known as "The Alchemists", have learned the secrets of the gemstones from ancient writings, older than any memory of the Men's existence.

The writings described the *Perfect Stone*, a circle-shaped stone also known as the *Rainbow*

[1]The game is available to download in *http://web.tagus.ist.utl.pt/˜ rui.prada/perfect-circle/*

Pearl, which bears incredible and unimaginable powers, as it merges the powers of all essences. For ages, the Alchemists have been looking for the Rainbow Pearl, some because they believe that the stone can reveal more of the gemstones' hidden secrets, others seek a source for the increase of their own powers, while some others just want to collect a reward given by wealthier Alchemists.

According to the writings, the *Perfect Stone* is hidden in one of the elemental planes. There are portals that open passageways between these planes, and the Alchemists have found ways of using the stones to activate these bridges. Using certain stones in a particular sequence activates these portals and enables the Alchemists to travel between planes. As this knowledge is spread among Alchemists many have undertaken the quest for the Rainbow Pearl.

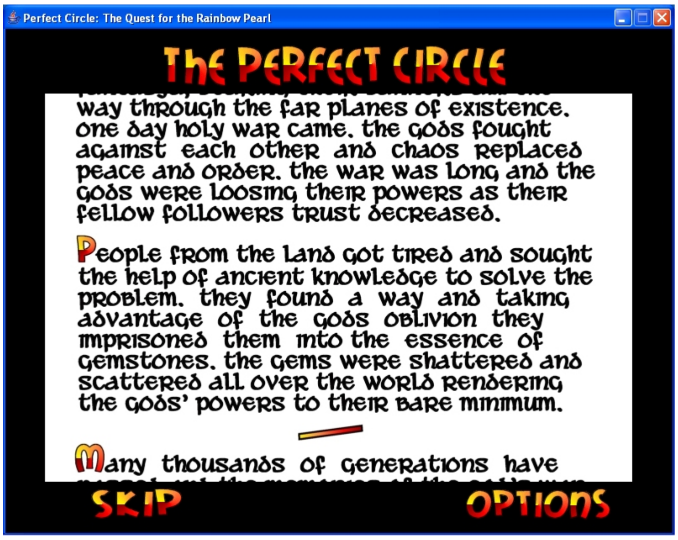

Figure 7.1: Snapshot of the game that shows the screen that presents the story to the player.

7.3.2 Gemstones

The concept of the game is centred on the manipulation of gemstones made from the essence of the gods. The stones are made of seven different essences that correspond to the seven rainbow colours (see figure 7.2) that were the colours that each of the seven gods used in their banners ages ago, when they freely walked on the earth, before being banned to the elemental planes. The seven essences are known to Men as:

1. **The Ruby:** which corresponds to the first rainbow colour, Red.

2. **The Topaz:** which corresponds to the second rainbow colour, Orange.
3. **The Citrine:** which corresponds to the third rainbow colour, Yellow.
4. **The Emerald:** which corresponds to the fourth rainbow colour, Green.
5. **The Sapphire:** which corresponds to the fifth rainbow colour, Blue.
6. **The Amethyst:** which corresponds to the sixth rainbow colour, Purple.
7. **The Iolite:** which corresponds to the seventh and last rainbow colour, Violet.

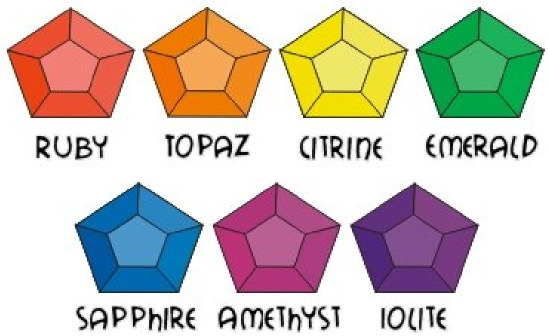

Figure 7.2: The seven essences.

Moreover, the stones can be found in *FantasyA* in different shapes according to their purity state. They are shaped as solid polygons, which can have any number of edges. The number of edges reflects the stone purity. The less pure and, therefore, more common gemstones only have a few edges while the more pure and more rare have many edges making them closer to the perfect shape, the circle.

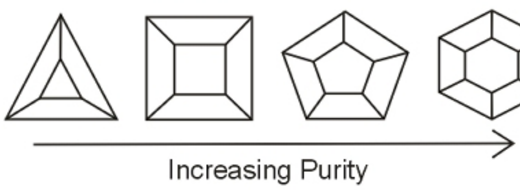

Figure 7.3: Gemstones and their different states of purity.

The stones can also be found in different sizes, reflecting the amount of essence that they hold and, consequently, their inner powers. Smaller stones are more frequent than bigger stones.

7.3.3 Game Play

Users play the role of Alchemists in their quest for the Perfect Stone. In the beginning of the game users are introduced to the background story and, then, are requested to create their characters. In the character creation process users choose the character's individual characteristics, such as its name and abilities (see figure 7.4). The user is given an amount of development points that s/he can distribute among the available abilities that define the character's initial skills. This process can be automatically done for the user. When the user's character is ready, the game proceeds.

Figure 7.4: The character creation screen.

7.3.3.1 The Group

In the next step, it is explained to the user that s/he (the character) was chosen to be part of a group of Alchemists, selected among the most prominent in the region, formed in order to search for the Rainbow Pearl. Then the other members of the group are introduced to the user one by one. In this description, the user can see the other characters' abilities and their personal information (name and personality traits). The personality traits reflect the values of *agreeableness* and *extraversion* and are presented to the user as words that describe the character (see figure 7.5). For example: this character is friendly, extroverted, very unfriendly, etc.

After inspecting the other Alchemists, the user decides if s/he will accepts the invitation to be part of the group and undertake the quest. If so, the group will be transported to the virtual environment that shows the surroundings of the first portal and the game begins.

This phase between the character's creation and the beginning of the game is important to give the user the sense that the task is going to be performed in a group.

Additionally, we have explicitly given no control to the user in the creation of the group and the other characters. This is often not the case in other games where the user interacts with groups of characters, like Role Playing Games in general. By taking this control from the user, we are stressing the autonomy of the characters. We did not want to give the user the idea that the other characters are in her/his direct control.

Figure 7.5: The group screen.

7.3.3.2 The Task

Once in the virtual environment, the group must interact in order to activate the portal and move further to the next plane (see figure 7.6). Portals are activated when Alchemists apply certain stone powers in a proper sequence. The sequence and type of stones required to activate the portal is inscribed in the portal, and is readable by any common Alchemist.

Thus, the main task of the group is to collect and combine stones until they get the necessary ones to open the portal. To achieve this goal, characters need to apply their individual abilities of gemstone manipulation. For instance, if the group has two small rubies but it needs one medium-sized ruby, a character can use its ability to merge the small stones into a bigger one (see section 7.3.4 for more details on the abilities). Two or more characters can combine their efforts to use an ability if they all have at least one development point on the ability. The effect will be the same if a single character has the skill level equal to the sum of all the individual levels of skills. Therefore, the probability of success of the actions is higher.

Figure 7.6: The group is engaged in the resolution of the task.

Furthermore, there is a wide range of possibilities to reach to the solution, since gemstones can be acquired in many different ways. The individual members of the group must, therefore, follow a similar strategy and coordinate their efforts. For that reason, the group engages in a negotiation process for every required action. This negotiation process is explained in detail in section 7.3.5.

In addition, the group should activate the portal as soon as the necessary stones are acquired, because time is an important factor, as other Alchemists are also on the run for the Perfect Stone. If the group takes too long to activate and to move across the portal, characters miss the chance to use the portal and must restart the quest in the current plane.

After entering the portal, the group will be teleported to another plane where they will find a similar challenge. The challenges become more and more difficult as the group approaches the plane where the *Perfect Stone* can be found.

When the group solves one of the puzzles and successfully goes through the portal, some development points are awarded to the group members. The amount of awarded development points depends on the group's performance. Then, characters may spend the new points to develop new abilities or improve the ones that they already have. The game ends when the *Rainbow Pearl* is found.

7.3.4 The Abilities: Gemstone Manipulation

This section describes the abilities that characters can develop in order to manipulate the stones. There are seven different abilities:

- **Merge stone:** using this ability, the character can join two small stones to create a bigger one. The two stones must be from the same essence and have the same shape and size. As a result, the character gets a new stone from the essence of the two smaller ones with the same shape and one degree of size bigger. The two smaller stones are consumed in the process.

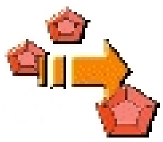

Figure 7.7: Merge stone.

- **Split stone:** this ability has the opposite effect of the merge stone ability. Using this ability, the character can split a stone into two smaller ones. The resultant stones are one degree of size smaller than the initial one and have the same essence and shape of the initial stone.

- **Transmute stone:** using this ability, the character can change the essence of a stone. The stone size and form of the stone stay the same. The change can only be made to one of the essences that are close to the initial essence according to the rainbow sequence. For example, a topaz (orange) can only be transmuted into a ruby (red) or a citrine (yellow). In the case of rubies (red) and iolites (violet) the change can only be made to

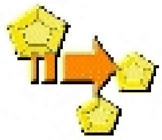

Figure 7.8: Split stone.

a topaz (orange) and amethyst (purple), respectively. The transmutation process is very unstable making it almost impossible to predict the essence of the resultant stone, except for the cases of rubies and iolites that are always transmuted into the same essence.

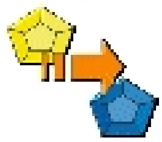

Figure 7.9: Transmute stone.

- **Purify stone:** using this ability, the character can increase a stone's purity state and, consequently, change its shape. The resultant stone gets one more edge and its size decreases one degree, the essence stays the same. The decrease in the size is due to the contraction of the essence that is needed to increase the stone's state of purity.

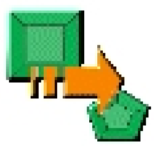

Figure 7.10: Purify stone.

- **Taint stone:** this ability has the opposite effect of the purify stone ability. By using this ability, the character can decrease a stone's purity state consequently removing one edge from its shape. The stone's essence stays the same but it increases. The increase in the size is due to the expansion that occurs with the decrease in the stone's purity.

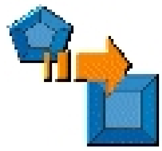

Figure 7.11: Taint stone.

- **Use stone:** this ability allows the character to use a stone's powers to activate the portal, the stone is consumed in the process.

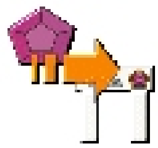

Figure 7.12: Use stone

- **Harvest stone:** using this ability, the character searches the ground for essence and extracts a stone of the desired essence. The stone may be of any form or size. The more frequent stones are those of small sizes and with low levels of purity.

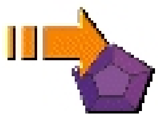

Figure 7.13: Harvest stone.

The ways to manipulate each essence are very different, and, therefore, each of the abilities described above needs to be specialized for each of the seven essences. This means that characters must develop each particular ability for each essence. For example, characters need to develop their ability to merge rubies separately from the ability to merge emeralds, or sapphires, or any other essence. Thus, there are 49 possible abilities for stone manipulation.

The difficulty that a character has in successfully using any of these abilities highly depends on the size and state of purity of the stones used with the ability. Its easier to process small and less pure stones than bigger and more pure ones. Thus, the more the character develops its abilities, the higher is the probability of success while using these abilities on bigger and more pure stones.

7.3.5 The Group's Interaction

Every character in the group is engaged in the same goal, thus trying to solve the same task. There are many ways to reach a solution, but if each one of the characters follows its own ideas, the group may never solve the task. Thus, characters have to coordinate their actions in order to follow a similar strategy in the search for the correct stones to activate the portal.

We need an interactive protocol for engaging all the members in group on each decision concerning the task that is suitable to involve the system-controlled characters and the user at the same time; the protocol should therefore be as simple as possible. We have searched for group decision strategies in human and autonomous agents studies and found inspiration in one of the simplest decision making protocols: the voting protocol.

Thus, every action that is performed in the group concerning the resolution of the task, is discussed by the group before it can be executed. However, instead of a formal vote, the other members only express their opinion, and the decision to perform or not the action is always taken by the character that proposed it in the first place. This decision is made according to the position that the other members have in the group, in the characters' point of view. For example, opinions expressed by members with more influence in the group are taken more seriously in the decision. Therefore, if two members in the group give an opinion against the action while just one agrees with it, this does not necessarily means that the action is not going to be executed. If the member that agrees with the action has a better position in the group than the other two, then the character will feel supported and will probably execute the action.

The interaction protocol has three different steps:

1. First, a character declares that s/he wants to take a certain action (e.g. *"I think that it will be best if I merge these two sapphires"*).

2. The other characters can respond to the proposal with one of the following:

    (a) ***Agree*** *with the course of action.* A character agrees with a proposal if the action is within the character's own plan and/or if the social relations built with the proposing character are favourable (e.g. the character has positive social attraction for the proposing character or is in her/his circle of social influence).

(b) ***Join*** *the action and help in the execution.* If a character agrees with a proposal and has her/himself the ability to execute the action s/he can decide to join efforts with the proposing character to perform the action together. The character can only take this position if s/he is not executing an action yet.

(c) ***Disagree*** *with the course of action.* A character disagrees with a proposal if the action is not in the character's current plans and/or if the social relation with the proposing character is not favourable (e.g. the character has negative social attraction for the proposing agent or is not in her/his circle of social influence).

3. Based on the opinions expressed by the group, the character decides to proceed with the execution of the action or to withdraw the proposal. If s/he decides to proceed with the action then s/he starts its execution. All of the other characters that have decided to join in the action start their contributions to the joint execution.

The group interaction is not restricted to the execution of the task. Each member can, at any time, engage in social-emotional interactions, by expressing their opinion about other members of the group. A character can encourage or discourage other members by directing some words to them (e.g. *"You're doing fine."*, *"What do you think you are doing?!"*).

7.4 *Case Study* Development Details

This section describes some details concerning the development of the *Perfect Circle* game. It describes: (1) the architecture of the autonomous agents that play the role of the synthetic members of the group, (2) how the SGD Model was used in such architecture and (3) how the complete system was integrated using the agent framework described in chapter 6.

7.4.1 The General System

The *Perfect Circle* game was developed in Java, Java3D[164] and JESS[2][56] [57]. The game's logics and user interface was implemented in Java. The 3D visualization of the synthetic characters and the environment was made in Java3D. For this, we have used some work that was

[2]JESS is a shell for Java that supports the creation of rule-based systems.

previously developed in the context of Teatrix [104][134]. This work includes some 3D scenes and character animations. Furthermore, for the development of the synthetic characters' minds we have used JESS because it facilitates the creation of a knowledge base and the creation of inference and production rules.

Moreover, the system that constitutes the game was integrated using the multi-agent framework described in chapter 6. Therefore, we have developed an *abstract world* that represents the game elements and rules. The players' interaction is achieved through a set of *agents* that were defined in the *abstract world*. There are two different types of *agents*: (1) a type that represents the group of players, the group *agents*, (2) and another type that represents the *agents* that were used to support the integration of the system, the system *agents*. Figure 7.14 summarizes the system's integration scheme. It shows the different *agents* and their relations with the system components.

1. **System Agents:** these agents are used for the control of the game and the visualization of the virtual environment:

    - *The Game Agent:* this agent is responsible for the control of the game progress. It gives the challenges to the group and notifies the group's members when to start or stop their resolution.
    - *The View Agent:* this agent gathers information from the abstract world and updates the view that is presented to the user. For example, it starts an animation when a member of the group starts the execution of an action.

2. **Group Agents:** these agents represent the members of the group:

    - *The User Agent:* this agent represents the user in the game. Thus, it translates the user's actions into correspondent executions in the *abstract world*. For example, if the user decides to start harvesting new gems, this *agent* will propose the action to the group on behalf of the user.
    - *The Synthetic Members:* these agents implement the behaviour of the autonomous members of the group.

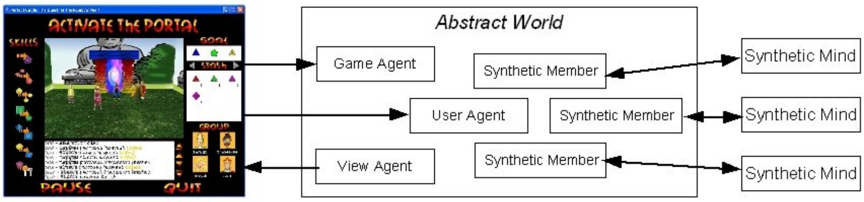

Figure 7.14: The System's Integration Scheme.

7.4.2 The Abstract World

7.4.2.1 Entities

Each character in the game is represented by an *agent* in the *abstract world*. These *agents* have a set of *properties* that defines the character's name, personality and abilities. These *properties* are set as follows:

- **Property(”name”, *char-name*):** a property named ”name” is set with the name of the corresponding character.

- **Property(”extraversion”, *extr-level*):** a property named ”extraversion” is set with the level of the extraversion dimension defined in the character's personality. We chose to model the extraversion using an integer with values from 1 to 100.

- **Property(”agreeableness”, *agr-level*):** a similar property is set for the other dimension in the character's personality, the ”agreeableness”. It also ranges from 1 to 100.

- **Property(”abilities”, *abilities-list*):** a property named ”abilities” is set with a list of the names of the character's abilities.

- **Property(*ability*:*essence*, *skill-level*):** a property is set for each of the character's abilities. The name of the property is a concatenation of the name of the ability with the character ':' and the ability's essence specialization. The value of the property is set according to the character's skill level for that ability. For example, if a character has the ability to merge rubies developed five skill levels, then a property *P(”Merge:Ruby”, 5)* is set in the corresponding agent.

The *properties* related to the personality are only set in the agents that represent the synthetic members, because the personality of the user is unknown.

7.4.2.2 The Group

The group is modelled as a *relation* that is established between the *agents* that represent the members of the group. As discussed in section 6.3, this relation has a property named "group" set with the name of the group, which is, in this case, "PerfectCircle".

The relation is created in the beginning of the game, and, upon creation, each of its members sets some properties that characterize his/her initial relation with the group. These properties correspond to the initial motivation and position in the group:

- **Property("group-position", *gp-value*):** each agent sets a property named "group-position" with the value they have computed for their initial position in the group.

- **Property("motivation", *motivation-value*):** each agent sets a property named "motivation" with the value of their initial motivation in the group.

By setting these *properties*, in the *abstract world*, members make their motivation and position in the group public to all the other *agents*. However, these values are only used by the *View Agent* in order to update the user interface. The autonomous members of the group will, in their minds, model the motivation and position in the group of all the members and use this internal value in their computations.

In the beginning of each challenge the *Game Agent* sets three properties in the relation that models the group, which define (1) the set of gems that the group possesses, (2) the goal of the group for the current challenge and (3) resolution state of the goal:

- **Property("Stash", *set-gems*):** a property named "Stash" is set in the relation. Its value contains the set of gems that the group currently possesses. This set is shared among all the members and represents the set of gems that can be manipulated.

- **Property("Goal", *set-gems*):** a property named "Goal" is set in the relation. Its value contains the set of gems needed to open the portal and proceed to the next challenge.

- **Property("ActiveGems", *set-gems*):** a property named "ActiveGems" is set in the relation. Its value contains the set of gems that were successfully used in the portal in order to open it. The challenge is complete when this set is equal to the set defined in the property "Goal".

An additional property is set in the group relation that represents the proposal being discussed. This property is set by the agents that want to make a new proposal to the group and is removed after the proposal has been discussed. I.e. after every member has given its opinion.

- **Property("CurrentProposal", *proposal*):** a property named "CurrentProposal" may be set in the relation. Its value defines the proposal that is under discussion in the group. The proposal contains information regarding the action to be performed (the ability and the gems used) and the member that started the proposal.

7.4.2.3 Sensors and Effectors

Each type of *agent* has a set of *sensors* and *effectors* that defines its capabilities to monitor and change the *abstract world*. To keep this description simple and focused, we will only describe the *sensors* and *effectors* defined for the *agents* that represent the members of the group.

1. **Sensors:** the group agents use some of the sensors that are pre-defined in the agent framework (see section 6.2.4) and define a new one to monitor the messages sent by the game agent concerning the game's state.
   - *Property Sensor:* monitors changes in the properties of the entities. E.g. changes in the characters' abilities levels.
   - *Relation Sensor:* monitors changes in the relation that defines the group. For example, it monitors if new members are added to or leave the group and monitors the properties that determine the group's goal.
   - *Effector Execution Sensor:* notifies the agent about the execution of effectors. For example, it notifies if the attempts to use the gems in the portal are successful or not.
   - *Message Sensor:* receives the control messages sent by the game agent. For example, these messages inform the members when to start to solve the challenge.
2. **Effectors:**

- *Skill Effector:* this effector allows the agent to use one of its skills in order to manipulate the gems. The effector receives arguments that determine the ability and the gems to use. If it succeeds, it removes the gems used from the group "stash" and place the ones produced. The success of the execution depends on the level of the skill of the agents involved in the execution. More than one agent can be involved since this is a joint effector (see section 6.3).

- *Use Gem Effector:* this effector allows the agent to use a gem's powers to activate the portal. It performs very similarly to the *Skill Effector*. However, in the case of success, it places the gem used in the "ActiveGems" set.

- *Start Proposal Effector:* this effector allows the agent to start the discussion of a new proposal in the group. It receives arguments that determine the nature of the proposal, e.g. the ability and gems to use. The execution is successful if there is not, at the time of the execution, another proposal in the group. If successful, the effector sets the "CurrentProposal" property in the group's relation.

- *Vote Effector:* this effector allows the agent to give its opinion regarding the group's current proposal. It takes an argument that specifies the opinion, which should be one of: *Agree*, *Disagree* or *Join*. The effector fails if the group is not discussing a proposal (e.g. the "CurrentProposal" property is not set).

- *Close Proposal Effector:* this effector allows the agent to close the discussion of the current proposal. It fails if the group is not discussing a proposal. In case of success, it removes the "CurrentProposal" property from the group's relation.

- *Join Execution Effector:* this effector allows the agent to help another agent in the execution of its *Skill Effector* or *Use Gem Effector*. It receives arguments that identify the agent and the effector to perform jointly.

- *Encourage Effector:* this effector allows the agent to encourage another agent. It receives an argument that identifies the agent that will be the target of the encouragement.

- *Discourage Effector:* this effector allows the agent to discourage another agent. It receives an argument that identifies the agent that will be the target of the discouragement.

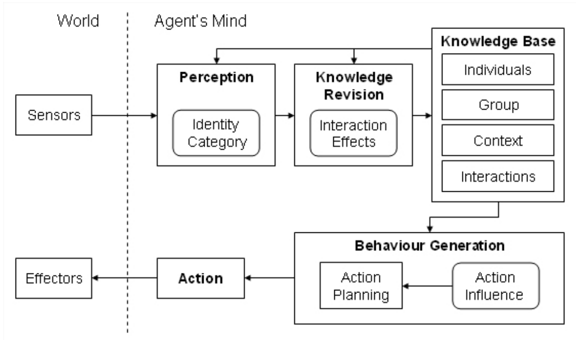

Figure 7.15: The agents' architecture.

7.4.3 The Agent's Mind

The mind of the autonomous members of the group was created according to the architecture shown in figure 7.15. It consists of five main modules: (1) the *perception module* that translates the events captured by the sensors into meaningful perceptions for the mind, (2) the *knowledge base module* that stores the beliefs of the agent, (3) the *knowledge revision module* that is responsible for the belief revision process, (4) the *behaviour module* that is responsible for the deliberation process, thus, generating the actions that the agent should perform, and (5) the *action module* that translates the action requests of the *behaviour module* into concrete executions of the agent's effectors.

7.4.3.1 The Knowledge Base Module

The knowledge base module stores facts that represent the agent's beliefs about the world. These include facts that define the members of the group, the group's state, the group's interactions and the group's task.

The knowledge is divided into three different sets: (1) the general knowledge that is directly derived from the perceptions generated by the agent's sensors, (2) the knowledge built regarding the group's task and (3) the knowledge built for the SGD Model.

The general knowledge contains facts that represent the concepts found in the abstract world (see section 6.2), such as properties and relations. A particularly relevant type of facts found in this set are those related to the execution of the agents' effectors. Figure 7.16 shows an example of a fact that states that agent *A1* has started to use its *Harvest* skill at time *10.*

(*action* (*action* UseSkill) (*agent* A1) (*args* Skill Harvest Essence Emerald) (*event* started) (*time* 10))

Figure 7.16: Example of an *action* fact.

The knowledge related to the group's task determines the agent's beliefs concerning the current goal of the group. Therefore, it includes, for example, facts that define the existing gems and that describe the group's current goal. Figure 7.17 shows an example of facts that state that the group's current goal is to get a *Ruby* and a *Sapphire* both at the minimum state of purity and with a very small size.

(*goal-gem* (*essence* Ruby) (*size* 1) (*purity* 1))
(*goal-gem* (*essence* Sapphire) (*size* 1) (*purity* 1))

Figure 7.17: Example of facts that define the group's current goal.

The knowledge built for the SGD Model contains facts that define the group, its members and the relations that support the group's structure (social influence and social attraction). In figure 7.18 we can see some examples of facts that state that *A1* and *A2* belong to the same group *PerfectCircle*, that *A1* is agreeable although a bit shy, and that *A1* likes member *A2* despite his belief that member *A2* does not like him. The example also states that *A1* has a better position in the group than *A2*. Note that the values of the personality dimensions vary between 1 and 100 and that the values of the social relations vary from -100 to 100. The values expressed in the 'group-position' facts represent the percentage of influence of the agent in the group, thus, they vary between 1 and 100. For example, in a balanced group of 4, each member will have 25 percent of influence in the group.

In addition, this set of knowledge defines the history of the interactions that occurred in the group, which includes the proposals discussed and the opinions of each member. In figure 7.19 we can see an example that states that *A1* was responsible for an *Encourage* interaction involving *A2* at time *20*.

Note that the three different sets, described above, constitute the knowledge that an autonomous member needs in order to perform in the group when it uses the SGD Model.

(*group-member* (*agent* A1) (*group* PerfectCircle))
(*group-member* (*agent* A2) (*group* PerfectCircle))
(*personality* (*member* A1) (*agreeableness* 70) (*extraversion* 23))
(*social-attraction* (*source* A1) (*target* A2) (*value* 30))
(*social-attraction* (*source* A2) (*target* A1) (*value* -50))
(*group-position* (*member* A1) (*group* PerfectCircle) (*value* 70))
(*group-position* (*member* A2) (*group* PerfectCircle) (*value* 30))

Figure 7.18: Example of facts that define the group's members and their social relations.

(*group-interaction* (*interaction* ENCOURAGE) (*agents* A1) (*targets* A2) (*time* 20))

Figure 7.19: Example of a fact that defines an interaction in the group.

However, the knowledge contained in the first two sets (general and task) is sufficient for a member that performs in the group's task without the guidance of the SGD Model.

7.4.3.2 The Perception Module

The perception module is responsible for filtering the information that is gathered by the agent's sensors and translates it into facts in the knowledge base. Thus, it keeps the knowledge base updated regarding the changes that occur in the abstract world. It records, for example, changes in the value of properties and keeps record of the execution of actions and the group's task current state.

This module has an additional functionality concerning the SGD Model. It is responsible for the identification of the group interactions as defined in section 5.6.1.

The interactions are classified according to the actions that are observed. In the case of the socio-emotional categories (*Encourage*, *Discourage*, *Agree* and *Disagree*) the identification is directly carried out from the events of execution of the correspondent effectors (Encourage, Discourage and Vote Effectors). Thus, every time that the execution of one of these effectors is successful a new 'group-interaction' fact is created in the knowledge base. The strength of the interaction in the group corresponds to the position of the performer in the group.

In turn, the identification of the instrumental interactions depends on the agent's current plan and the success of the task actions. Therefore, an action will be interpreted as positive for the group - *Facilitate Problem* - if it succeeds, and it will be interpreted as negative - *Obstruct*

Problem - if it fails. The strength of the instrumental interactions in the group depends not only on the position of the members that execute it and its supporters in the group, but also on the relevance of the action for the group's task (see equation 7.1 for details). Actions that the agent has considered in its planning process are defined as relevant for the task, in the agent's perspective, while any other action is not relevant. Furthermore, more relevant actions will produce interactions with higher strength. Thus, for example, failing an action that is not relevant for the task leads to a less significant loss in the social influence than failing a relevant action. The relevance of an action is measured in a value that goes from 0 to 100. In this implementation all relevant actions have a relevance value of 100 while the non relevant have a value of 0.

$$\begin{gathered}\forall G, I : Group(G) \wedge InstrInteraction(I), \\ Strength(I,G) = 0.3 * Relevance(I,G) + \\ 0.7 * (\sum_{p \in Performers(I)}^{p} Position(p,G) + \sum_{s \in Supporters(I)}^{s} \frac{Position(s,G)}{2}) \end{gathered} \tag{7.1}$$

7.4.3.3 The Knowledge Revision Module

The perception module only produces new facts in the knowledge base that are directly connected to the perceptions of the agent. Thus, it will not make further deduction like inferring the emergent social relations in the group. This kind of deduction is performed by the knowledge revision module.

The main functionality of this module is to apply the effects of the interaction interactions on the group state. Therefore, for each new interaction this module computes the new values for the variables that the interaction influences (see section 5.6.2). For example, for *Encourage* interactions this module computes the change on the values of the target's *motivation* and *social attraction* for the performer.

The amount of change to apply to the values is computed using a formula that depends on five arguments: (1) the current value; (2) the intensity of the change; (3) the minimum value; (4) the maximum value; and (5) the direction of the change (increment or decrement). The change is not linear since it depends on the distance of the current value to the neutral value (e.g. 0). This is used to simulate the tendency of the values to stay neutral. Thus, it is harder to produce a change in the value from the neutral state than it is from its higher values (e.g. 100). Furthermore, the intensity is a value between 0 and 100 and corresponds to the strength of the interaction in the group. Additionally, the new value will never be out of the boundaries defined by the minimum and maximum values. The function used to compute the absolute value of the amount of change is shown in equation 7.2.

$$change = \frac{10 + value^2}{320} * \frac{50 + intensity}{100} \tag{7.2}$$

This function is composed of two factors: the first one computes the change in relation to the current value while the second computes the influence of the intensity in the change. It follows the properties discussed above regarding the dependence of the change in relation to the current value and the intensity. Thus, for example, if the current value is 50, then the change will be between 9 and 27 according to the strength of the change. On the other hand, if the current value is 0, then the change will be lower (from 5 to 15).

Moreover, if the revision process produces a change on the social relations then the members' position in the group is recalculated using the formula 5.7 presented in section 5.5. Note that after this calculation the values of the positions in the group are normalized in order to express the influence of each member in the group represented by a percentage.

7.4.3.4 The Behaviour Module

The behaviour module manages the decisions concerning the agents' acting. It consists of two distinct processes: one related to the agent's task planning, and another that determines its reactions to the group interactions, according to the rules described in section 5.6.2.1.

The Planning Process

The task's planning is a deliberative process that runs on regular intervals. The length of these intervals depends on the variables that influence the frequency of interaction in general, i.e. the agent's extraversion, motivation and position in the group (see section 5.6.2.1). The value of the interval is computed using a liner combination of the three variables that are shown in equation 7.3. It determines the number of mind's update cycles where the planning process is asleep.

$$length = 0.2 * (100 - motivation) + 0.1 * (100 - position) + 0.3 * (100 - extraversion) \tag{7.3}$$

In the end of each idle interval the agent may start its planning procedure. However, it first determines if it has a predisposition to interact. This predisposition also depends on the three variables that influence the general frequency of interaction. They determine a threshold that is checked by a probabilistic function in order to decide if the agent is motivated or not to act. Equation 7.4 shows the formula that is used to compute this threshold, it returns a value between 0 and 100. The probabilistic decision function generates a random number between 1 and 100 and checks if it is bellow the threshold. If so, the agent is motivated.

$$threshold = 0.4 * motivation + 0.2 * position * groupsize + 0.4 * extraversion \tag{7.4}$$

If the agent is motivated to act, the planning algorithm is started. It computes a partial plan for the current state of the task involving all the members of the group. It consists in a backward search that follows the algorithm described below:

1. First, the goal nodes are generated, one for each gem that is in the goal set and that was not used in the portal. These nodes correspond to the action "use gem" that activates a gem on the portal. The performer of the action is chosen among those that have the ability to use gems of the correspondent essence. Only the best performer according to the *heuristic* presented in equation 7.5 is used. The goal nodes are placed in the open list and the algorithm proceeds to the second step.

2. In this step, the best node in the open list (the one with the lower cost) is chosen and the current group's stash is checked to meet the node's prerequisites. The prerequisites consist in the gems that are needed for the action specified by the node. If they are present in the group's current stash, then the planning algorithm is ended and the action is returned as the result.

3. If the prerequisites of the node are not met, then new search nodes are created to achieve them. These new nodes correspond to the actions that create the missing gems. For example, if a size 2 ruby is needed, then three actions are generated: one that merges two size 1 rubies, one that splits a size 3 ruby and one that harvests the ruby. Only the actions that have at least on member in the group that is able to perform them will be generated. If more than one member can perform the action the node is generated using the best performer according to the *heuristic* (see equation 7.5).

4. The new nodes are included in the open list and the algorithm proceeds to step 2 until an executable action is found or the open list is empty, which means that it was not possible to achieve a plan that solves the goal. In addition, the search was limited to plans with depth 6. Therefore, at this level, the algorithm only produces "harvest" actions (if there is at least one member that can perform them). Which means that if no solution is found for the problem, given the current goal and group stash, the group will try to get new gems from the soil in order to increase its stash.

The *heuristic* used to drive the search computes the cost of the action, considering the performer of the action and its skill level as well as the distance to the goal and the type of action. The agent considers two type of actions, "harvest" and "transmute", to be more costly than the others, and it prefers the actions that are performed by itself. The formula to compute the *heuristic* value of a node is shown in equation 7.5.

$$\begin{aligned} heuristic = {} & ActionBaseCost(action) + PerformerCost(performer) + \\ & MergeCost(action, gem) + 5 * Distance(gem, goal) - \\ & 5 * SkillLevel(performer, action) - UseCost(action, gem) \end{aligned} \tag{7.5}$$

The *ActionBaseCost* returns the base cost for each type of action. These costs are predefined and equal to: 3000 in the case of "harvest"; 2000 in the case of "transmute"; and 600 in the case of any other action. The *PerformerCost* returns 10000 if the performer is not the agent itself. The *Distance* function compares the gem used in the action with the gem in the top goal node. It computes the differences in terms of size, purity state, and the index of the essence in the rainbow (e.g. Ruby(red) has a index of 1 and Iolite(violet) has an index of 7).

The *MergeCost* function returns an additional cost to the merge action if the group does not have at least one of the required gems (note that the merge action needs two gems). The *UseCost* function reduces the *heuristic* cost value if the action is to use a gem in the portal and the group already has the required gem. This function is used to privilege actions that directly fulfil a goal. The *SkillLevel* function returns the level of expertise of the performer for the specified action.

The planning process has three possible outcomes: (1) no suitable action is found; (2) an action to be performed by the agent itself; (3) an action to be performed by another agent. In the second case the agent proposes the action to the group and waits for its members' opinion, while in the third case the agent will encourage the agent that he believes should perform the action.

Moreover, the decision to perform or not the action that was proposed is made according to the opinions and positions in the group of the other members. Thus, opinions expressed by members with more influence in the group are taken more seriously in the decision. For example, if two members in the group give an opinion against the action while just one agrees with it, this does not necessarily mean that the action is not going to be executed. If the member that agreed with the action has a better position in the group than the other two together, then the character will feel supported and will execute the action. Equation 7.6 shows the formula that agents use to compute the value that supports their decision concerning the execution of actions they propose in the group. If this value is above 0 then they proceed with the action, otherwise, they drop it.

$$\forall G, P, A : Group(G) \wedge Proposal(P) \wedge A \in Members(G) \wedge Proposes(A, P, G),$$
$$value = Position(A) + \sum_{m:Agrees(m,P)} Position(m, G) - \sum_{m:\neg Agrees(m,P)} Position(m, G) \quad (7.6)$$

This formula in only applied in the agents that use the SGD Model. In the case of the agents that do not use the SGD Model the decision is made through a simple majority rule. Note that, after this consideration, the planning algorithm returns to its sleeping state. Thus, it will stay inactive for a number of cycles in the mind correspondent to the value returned by the formula presented in equation 7.3.

The Reactive Process

The behaviour module has a second process that deals with actions that are not proactive. This process runs in parallel with the proactive behaviour, and it also depends on the fact that the agent is motivated to act. This means that if the agent is not motivated as determined by the formula in equation 7.4, then the reactive process will not be executed.

The actions requested by this process are response to two different situations: (1) reactions to some of the group interactions, i. e., socio-emotional reactions and (2) the expression of opinions about the proposals of the other agents.

Socio-emotional reactions may occur in three different situations: (1) in reaction to instrumental interactions (i.e success and failure of task actions), (2) in reaction to encouragements and discouragements directed to the self, and (3) in reaction to encouragements and discouragements directed to other members. Therefore, upon the occurrence of any of these 3 situations, group members have to decide if they will or not actively intervene in the group. This decision depends on the strength of the interaction in the group (e.g. how important the success of an action was to the resolution of the group's task) as well as the member's personality and the social structure of the group.

In the first situation, when instrumental interactions take place, agents make their decision using a probability rule based on the formula presented in equation 7.7. This formula computes a value that defines the probability of reaction with an encouragement to an instrumental interaction. This value increases with the strength of the interaction in the group if the interaction is positive (Facilitate Problem) and decreases if the interaction is negative (Obstruct Problem). This means that positive instrumental interactions have higher probability of generating an encouragement than negative instrumental interactions. On the other hand, negative instrumental interactions have a better chance to generate a discouragement. However, the probability also depends on the agent's agreeableness, the difference between the position in the group of the interaction's performer and the agent, and the social attraction of the agent to the interaction's performer.

If the value computed in the formula is positive then there is a probability that the agent will perform an encouragement action. Otherwise, if the value is negative, then the probability is that it will perform a discouragement action. The value represents a threshold that is checked against a randomized number between 1 and 100. If the generated number is below the absolute value of the threshold, then the agent decides to perform the correspondent action.

$$\forall G, I, A, P : Group(G) \wedge InstrInteraction(I) \wedge \{A, P\} \subset Members(G) \wedge Performs(P, I, G),$$
$$threshold = 0.5 * Strength(I) + (Agreeableness(A) - 50) +$$
$$(Position(A, G) - Position(P, G)) + 0.3 * SocAttraction(A, P) \tag{7.7}$$

In the second situation, when there is a reaction to encouragements and/or discouragements directed to the self, the reaction is always reciprocal. This means that agents only react to encouragements with encouragements and to discouragements with discouragements. The decision to reply to the encouragement is made according to the formula presented in equation 7.8. This formula computes a value between 1 to 100 that represents the probability of responding with an encouragement action. This probability increases according to the value of the agreeableness of the agent and the social attraction of the agent to the interaction's performer. In addition, it depends on the difference between the agent's position in the group and the performer's position in the group. I.e. it increases if the position of the performer is higher than the position of the agent and decreases otherwise. The probability of replying to a discouragement is computed using a similar formula shown in equation 7.9.

$$\forall G, I, A, P : Group(G) \wedge Encourage(I) \wedge \{A, P\} \subset Members(G) \wedge$$
$$Performs(P, I, G) \wedge Target(I, A),$$
$$threshold = 20 + (Agreeableness(A) - 50) +$$
$$(Position(P, G) - Position(A, G)) + 0.5 * SocAttraction(A, P) \tag{7.8}$$

$$\forall G, I, A, P : Group(G) \wedge Discourage(I) \wedge \{A, P\} \subset Members(G) \wedge$$
$$Performs(P, I, G) \wedge Target(I, A),$$
$$threshold = 20 - (Agreeableness(A) - 50) -$$
$$(Position(P, G) - Position(A, G)) - 0.5 * SocAttraction(A, P) \tag{7.9}$$

In the third situation, the case of reactions to encouragements and/or discouragements that are directed to other members, the decision is made according to a set of probabilistic rules based on formulas similar to the one presented in equation 7.10. There are four different reaction rules that were defined following the ideas found in Heider's Balance Theory (see section 4.6.2). Two of these rules are applied if the agent has a positive social attraction to the other member that was the target of the interaction (encouragement or discouragement). If so, these rules check the probability of the agent to reply to the performer of the interaction with an encouragement action if he has encouraged the other member, or with a discouragement action if he has discouraged the other member. The other two rules are applied if the agent has a negative social attraction to the target of the interaction. In this case, the rules decide if the agent encourages the performer if he has discouraged the target or discourages him if he has encouraged the target. Table 7.1 summarises the possible reactions of the agent given each of the four conditions. Furthermore, the formula presented in equation 7.10 computes the probability of the execution of an encouragement action as a reaction to an observed positive socio-emotional interaction, when the agent has a positive attraction to the target of the interaction.

$$\begin{aligned} \forall G, I, A, P, T : Group(G) \wedge Encourage(I) \wedge \{A, P, T\} \subset Members(G) \\ Performs(P, I, G) \wedge Target(I, T), \\ threshold = 0.5 * (Agreeableness(A) - 50) + \\ (Position(P, G) - Position(A, G)) + 0.5 * SocAttraction(A, T) \end{aligned} \tag{7.10}$$

SocialAttraction(A, T)	**Encourage(P, T)**	**Discourage(P, T)**
positive	*Encourage(A, P)*	*Discourage(A, P)*
negative	*Discourage(A, P)*	*Encourage(A, P)*

Table 7.1: Possible reactions of the agent when he observes two other members engaged in a socio-emotional interaction. A denotes the agent, P denotes the performer of the interaction and T denotes the target of the interaction.

The other kind of reactive behaviour is related to the discussion of the actions that other members propose in the group. Thus, whenever a new proposal is made, the agent has to consider its relevance and decide if it should agree or not with the proposal. When the agents do not use the SGD Model, this decision is made simply by verifying if the action was considered in the planning algorithm. If so, the agent agrees with the proposal. Otherwise, it disagrees.

However, agents that use the SGD Model extend this simple decision algorithm by considering the values of their agreeableness and their social relations with the member that proposes the action. For example, if the agent is very agreeable, or if it has a positive social attraction to the member or even if the member has a better position in the group than the agent, then the agent is more likely to agree with the proposal. The formula presented in equation 7.11 is used by the agent in this decision. If the value computed is positive, then the agent will agree with the proposal, but, if the value is negative, it will disagree. The *base* value referred to in the equation corresponds to the relevance of the action based on the planning algorithm. It will take the value of 50 if the action was considered during planning and -50 if not.

$$\forall G, P, A, B : Group(G) \land Proposal(P) \land \{A, B\} \subset Members(G) \land Proposes(B, P, G)$$
$$vote = base + 0.5 * (Agreeableness(A) - 50)) +$$
$$0.5 * (Position(B, G) - Position(A, G)) + 0.5 * SocAttraction(A, B) \quad (7.11)$$

Moreover, if the agent agrees with a proposal and it has the ability to perform the action, it will help the agent in its execution.

7.4.3.5 The Action Module

The action module translates the action requests from the behaviour module into concrete executions of the agent's effectors. In our example, agents may execute four different types of actions:

1. **Propose(A,G1, G2):** this is an act that starts the discussion of a proposal for action in the group. A proposal describes the ability *A* and the gems to be used *G1* and *G2*. E.g. "I propose to *merge* these *two rubies*.".

2. **Agree(P), Disagree(P), and Join(P):** these are acts that agents use to express their opinions concerning the current proposal *P*.

3. **Execute(A,G1, G2):** this starts the execution of a task action using an ability *A* with some gems *G1* and *G2* (*G2* is optional). For example, it starts to merge two gems or start to use a gem on the magic portal.

4. **Encourage(C), Discourage(C):** these are acts that express encouragement or discouragement towards another character *C*. For example, "Keep up the good work." or "Stop messing up.".

7.5 Analysing the Task

We developed our task taking into account, whenever possible, the considerations made by Mennecke and Wheeler. Thus, we will present some comments, concerning our task, on each of the six issues presented in their essay.

1. **Appropriateness of the Task to the Subjects:** all the necessary knowledge to execute the task is within the game and all the elements of the group have access to it.

2. **Subjects' Intellectual Engagement:** the game presents a background story with the purpose to engage the users in the game. Additionally, the group is rewarded with development points after the correct execution of each puzzle that can be applied to further development of each character. These techniques are often used in computer games as a means of attaining users' engagement.

3. **Control of the Differences in the Subjects' Preferences, Needs and Experiences:** The environment is shared between all the members and the task's instructions are clearly defined, so the experience is similar for all members. In the case of computer controlled characters these issues do not apply.

4. **The Level of Task Complexity:** The task is not trivial and requires a certain engagement of the user to be perfectly understood, but it is not difficult because it does not require any special skills besides a perfect understanding of the abilities. The difficulty increases as the group gains experience in the task resolution (e.g. the puzzles get more difficult), this can also be seen as a factor for intellectual engagement.

5. **Conjunctive versus Disjunctive Tasks:** The task is conjunctive for two different reasons. First, the task defines a high number of different skills, see section 7.3.4 for details, and, therefore, no character can successfully develop all of them. Consequently, a single character can not perform all the necessary actions to solve the puzzle. Second, several characters will solve the task faster than a single character and time is an important factor on the task reward.

6. **Task Typology and Measurement of the Solution Quality:** Our goal is not to evaluate the quality of the solution reached by the group, it is centred on the quality of the group interaction. Thus, this issue is not relevant for our task.

7.6 Concluding Remarks

In this chapter we presented the task that was developed for the evaluation study, described in chapter 8, designed to evaluate the model for synthetic group dynamics that we propose in this research (see chapter 5). The design of this task followed some principles in order to make it more suitable to be used in a group and, therefore, did not compromise the results of the study.

The task was developed as a computer game, the *Perfect Circle*. This game promotes the collaboration of the user with several synthetic characters as a group of adventurers that are searching for a sacred item hidden in the world of *FantasyA*. To perform this task, the group has to, collaboratively, collect and manipulate the essence of several gemstones in order to get the correct ones to open portals, making it possible to travel between several planes.

While performing the task, the group discusses the course of actions by sharing their opinions about each of the actions that should be taken in order to complete the task. Additionally, the members of the group may engage in socio-emotional interactions.

Moreover, we have described some details regarding the development of the game. We started our description with the general scheme of integration of the different components and how it was done using the agent framework presented in chapter 6. Then, we have focused on the creation of the mind components of the synthetic members of the group explaining how the SGD Model was integrated to generate their behaviour. These components were developed in a modular way that allowed the implementation of some agents that did not use the SGD Model. This was important to facilitate the creation of the different scenarios used for the conditions of the evaluation study described in chapter 8.

Chapter 8

Evaluation

8.1 Introduction

This chapter describes the experiment that we have conducted in order to evaluate the SGD Model. Its main goal was to test the hypothesis formulated in the beginning of this document (see section 1.2), that is:

> *"If the interactions in a group of synthetic characters follow similar dynamics as the interactions in human groups do, the synthetic group will become more believable for a user that participates in the group, which, consequently improves the user's interaction experience."*

In this experiment we have used the *Perfect Circle* game, described in chapter 7, to engage users with a group of four synthetic characters in the resolution of a collaborative task. The experiment was conducted at IST - Technical University of Lisbon with several students and the results showed that the SGD model had some positive effects on the users' interaction experience, namely on their trust and identification with the group.

8.2 Independent Variables

The experiment was conducted with two main control conditions: (1) the fact that the group's interactions follow, or not, the SGD Model, and (2) the initial structure and consequent cohesion level of the group. These two control conditions allowed us to draw some conclusions on whether the presence of the model would lead to the predicted results concerning the believability of the group, and, at the same time, study the effect of the cohesion of the group in relation to the studied variables. See a description of the two conditions below:

1. **The Use of the SGD Model:** two different versions of the game were built, one where the characters followed the SDG Model and other where they did not. Thus, the first condition determines whether the subjects play with or without the believable group dynamics component. This control variable is driven directly from the hypothesis mentioned in the previous section.

2. **The Group Initial Structure:** subjects can start the game in a group with non neutral social relations, which means that the initial group can have levels of cohesion that may be either very high or very low. Two different scenarios were considered: one where the group had neutral social relations and a second one where the members of the group disliked each other leading the levels of cohesion in the group to become very low. Note that this condition could only be applied when the game was run with the believable group dynamics component.

   We do not have a strong hypothesis on this second condition. Our goal was to check how the initial cohesion level of the group could affect the dependent variables.

In addition to the two mentioned independent variables, we defined three other variables, that we believed would help in the interpretation of the results: the subjects' gender, personality and gaming experience.

1. **Gender:** there are studies that state that female users and male users have different gaming preferences [83][82][30], for example, males prefer action games while females prefer social games. By verifying the gender variable we wanted to check if this difference had any effect on the subjects' evaluation. For example, do female players benefit more from the use of the SGD Model given that it enhances the game's social component?

2. **Personality:** it was previously discussed in chapter 4 that personality is one of the factors that influence the behaviour of people when interacting in a group. Thus, the subjects personality may have effect on the results. For example, would more cooperative subjects benefit more from the use of a believable group dynamics? Or would it be otherwise?

3. **Gaming experience:** it was considered that subjects' gaming experience, and their attraction to gaming, may change their perception of the game, and, therefore, may influence their evaluation. For example, subjects that play computer games more often or that like computer games a lot, may give less importance to the details of the group interaction and be more concerned with the details of the task. Thus, do "novice" players gain more from a believable group dynamics than "veteran" players? Or will it be the other way around?

8.3 Dependent Variables

The main goal of this experiment was to assess the quality of the subjects' interaction experience while playing the *Perfect Circle* game. To assess this experience, we have considered the subjects's interaction at two different levels: one exclusively concerning the interaction with the group and another concerning the interaction with the game in general.

Concerning the group interaction it was decided to observe two different variables: the subjects' identification with the group and their trust in the group. We based our choice on the work of Allen et al. [4] that have used these two variables on their studies to assess the satisfaction of people when interacting in a group.

1. **Group Trust:** people's trust on a group has a positive effect on their perceptions about their experience in the group [42], which consequently leads to a more satisfactory interaction [6].

2. **Group Identification:** according to Ashforth and Mael [7] social identification is, in addition to social trust, one of the factors that foster the members of a group to be more engaged and more satisfied with the group.

Supported by these findings, one can conclude that the level of the users' identification with the group of synthetic characters along with their levels of trust make good indicators for the quality of their interaction with the group. Furthermore, for the evaluation of the general interaction with the game we have considered two additional variables: the subjects' social presence in the virtual environment and their satisfaction with the game.

- **Social Presence:** in chapter 3 we discuss the importance of presence and immersion in a virtual environment. Presence is associated with the believability of a virtual environment and, consequently, the higher the degree of presence, the better the interaction experience for the user. Furthermore, if the virtual environment is shared with other users, the physical presence is not enough to convey a believable reality, thus, a sense of social presence must also be conveyed to the user. In our experiment, users interact with a group of synthetic characters that simulates the interactions in human groups. We expect to enhance the users' social presence by making them experience the sense of being "there" interacting in a group.

- **Satisfaction with the Game:** computer games are supposed to be fun, thus, the user should enjoy every moment that s/he spends with the game. Hence, the improvement in the interaction experience, as stated in the initial hypothesis, would imply also the increase in the user's fun.

8.4 Measures

To measure the variables discussed in the previous sections (8.2 and 8.3), we have adapted, whenever possible, questionnaires found in the literature and previously applied in other studies. However, in some cases we had to developed our own.

These questionnaires present several sentences, which describe people's behaviours and opinions, that the subjects are asked to classify. This classification rates the questionnaire's items in terms of their accuracy to describe the subjects' own behaviours, opinions and experience. The rating goes from 1, meaning completely inaccurate, to 7, meaning completely accurate (see table 8.1).

Furthermore, the questionnaires' items can express positive or negative sentences in relation to the topic that they evaluate. For example, in our gaming questionnaire one of the items states *"I love to play computer games"*, which is a positive sentence in relation to gaming experience, while another item states *"Playing on the computer is a waste of time"*, which is a negative one.

To assess the value of each variable we compute the average value of the subject's answers in all the items of the variable's questionnaire. However, while the values of positive items are used directly in this calculation the value of the negative items are inverted to their positive correspondent. E.g. the minimum value of a negative item corresponds to the maximum item of a positive item and the maximum value of a negative item corresponds to the minimum of

Value	Meaning
1	Completely Inaccurate
2	Very Inaccurate
3	Moderately Inaccurate
4	Neither Inaccurate nor Accurate
5	Moderately Accurate
6	Very Accurate
7	Completely Accurate

Table 8.1: The questionnaire items' rating.

a positive item. Thus, for example, if the value of a negative item is 5 then its positive value is *8 - 5 = 3*.

In the next sections we will present and discuss each of the questionnaires used to measure the value of the variables, namely the questionnaires for the *personality*, the *gaming experience*, the *group trust*, the *group identification*, the *social presence* and the general *satisfaction with the game*.

Moreover, the questionnaire given to the subjects was composed of all the items that belonged these questionnaires and, in addition, two items that inquired the age and gender of the subjects, and one open item where the subjects could freely write their comments about their experience.

8.4.1 Personality

To measure the subjects' personality we have used a fifty item questionnaire developed by the International Personality Item Pool [62] [84] that is based on the Big Five model of personality [61].

This questionnaire contains ten different items (five positive and five negative) for each of the five factors described in the Big Five model (see table 8.2). Thus, the output of the questionnaire measures not one, but five different variables.

In addition, despite the fact that the classification of each of the items in the questionnaire is given in values from 1 to 7, the measures of the five factors are presented in correspondent values that range from 1 to 100, which is a more common practice in these types of questionnaires.

Factor	Designation
I	Surgency or Extraversion
II	Agreeableness
III	Conscientiousness
IV	Emotional Stability
V	Intellect or Imagination

Table 8.2: The Big Five factors.

The complete list of items of this questionnaire is presented in Appendix A.

8.4.2 Gaming Experience

Concerning the computer gaming experience, we did not find any relevant studies or questionnaires that could be used in our study. Therefore, we had to design our own questionnaire. To do so, we defined some items that are related to the ideas behind the definition of this variable, which should measure if the users like to play computer games, if they play them often and if they believe that playing computer games is important.

As a result, we have designed a questionnaire with six different items (three positive and three negative) that are shown below. Each item is presented in English and Portuguese and has an indication that identifies its positive or negative nature:

1. I love to play computer games (positive)
   Gosto de jogar jogos de computador.

2. Playing on the computer is a waste of time. (negative)
   Jogar no computador é uma perda de tempo.

3. I play computer games frequently. (positive)
   Jogo frequentemente jogos de computador.

4. The computer should only be used as a working tool. (negative)
   O computador deve ser usado apenas com ferramenta de trabalho.

5. I play several hours a day on the computer. (positive)
   Jogo algumas horas por dia no computador.

6. I do not understand those who play computer games. (negative)
Não compreendo quem gosta de jogar computador.

8.4.3 Group Trust

In the case of group trust we relied on the questionnaires that Allen et al.[4] used in their studies. They proposed a seven item questionnaire with five positive items and two negative ones. However, for consistency purposes we have only used six of these items and changed one of the positive sentences into a negative one in order to keep a balance between the negative and positive items, thus, having three items of each kind. The items of this questionnaire are presented below:

1. Most people on this team are basically honest and can be trusted. (positive)
Os outros membros do grupo agem de boa vontade.

2. On this team, team members are always interested only in their own welfare. (negative)
Neste grupo os elementos apenas se interessam pelo seu próprio proveito.

3. Members in this team are always trustworthy. (positive)
Posso sempre confiar nos membros deste grupo.

4. In this team, one has to be alert or someone is likely to take advantage of you. (negative)
Neste grupo é preciso estar em alerta caso contrário os outros abusam da confiança.

5. If I have a problem, there is always someone to help me. (positive)
Se eu tiver um problema o grupo está sempre pronto a ajudar-me.

6. Nobody in the group is willing to help me with my tasks. (negative)
Ninguém no grupo me ajuda nas minhas tarefas.

8.4.4 Group Identification

To build the questionnaire that measures the identification with the group, we relied on the work of Allen et al.[4] again, since they also proposed a questionnaire to measure this variable. Their questionnaire is composed of five different elements with a positive nature. These items formed the base of our questionnaire, however, with two significant changes: first, some of the sentences were changed to meet our gaming scenario, and second, three of the items were changed to negative. In addition, a new positive item was added, to complete the set of six. The resultant questionnaire is shown below:

1. I feel strong ties with the member of this group. (positive)
   Criei laços fortes com os membros do grupo.

2. I did not enjoy to play with this group. (negative)
   Não gostei de jogar com este grupo.

3. I feel accepted as a member of this group. (positive)
   Sinto que fui bem aceite pelos outros elementos do grupo.

4. I experience a sense of not belonging to this group. (negative)
   Sinto que não pertenço a este grupo.

5. If I play again I would like to play with the same group. (positive)
   Se voltar a jogar o jogo gostaria de jogar com o mesmo grupo.

6. I am not sufficiently acknowledged in this group for my expertise. (negative)
   O grupo não valorizou as minhas competências.

8.4.5 Social Presence

In order to find a way to measure the social presence of the subjects in the game we have consulted several resources through the Presence-Research portal [136], namely the Igroup Presence Questionnaire (IPQ) [149] and the IPO Social Presence Questionnaire (IPO-SPQ) [39].

These two questionnaires respectively define 14 and 17 different items each. All of positive nature. However, for the purpose of keeping our questionnaires consistent we only used six of these items and changed three of them to negative sentences. The resultant questionnaire is shown below:

1. The interaction with the group in the virtual world was consistent with my experience in the real world. (positive)
   A interacção com o grupo no mundo virtual foi consistente com a minha experiência no mundo real.

2. I was always aware of the environment in the room. (negative)
   Não me consegui abstrair do ambiente da sala.

3. I felt close to the members of the group (positive)
   Senti-me próximo dos membros do grupo.

4. I felt that the characters were not aware of my actions. (negative)
   Senti que os personagens não estavam atentos às minhas acções.

5. The characters acted and reacted according to my expectations. (positive)
   As acções e reacções dos membros do grupo estavam de acordo com as minhas expectativas.

6. I felt that the characters were just lifeless static images. (negative)
   Senti que os personagens eram apenas imagens estáticas sem vida própria.

8.4.6 Satisfaction with the Game

We have found several questionnaires in the literature to measure the users' satisfaction with computer systems, such as the End User Computing Satisfaction questionnaire [40]. However, these questionnaires focus on questions related to the system's accuracy, ease of use and effectiveness, and do not take into account the user's joy during the experience. In fact, as stated by Wiberg [177], these classical measures are not completely appropriate to attain the users' satisfaction in entertainment systems. For example, if the user spends a lot of time on a particular task this is not necessarily a bad sign, since this may happen because the user is having fun with the task.

Nevertheless, we did not find any reference questionnaires concerning the users' satisfaction in entertainment systems, thus, we have developed our own questionnaire based on the notion of the user's fun. The items of this questionnaire are shown below:

1. I loved playing this game. (positive)
   Adorei jogar este jogo.

2. I felt bored while playing the game. (negative)
   Senti-me aborrecido ao jogar o jogo.

3. The game was very interesting. (positive)
   O jogo era bastante interessante.

4. I would not suggest this game to anybody. (negative)
   Não aconselhava este jogo a ninguém.

5. I would like to play this game again. (positive)
   Gostaria de voltar a jogar o jogo.

6. The game was too complex. (negative)
   O jogo era demasiado complexo.

8.5 Method

The experiment was conducted in IST - Technical University of Lisbon, with a sample of twenty four students, twenty of them were male and four, female. It was run in a laboratory with six computers for a whole day. The experiment was divided into four sessions of two hours each with six subjects at a time.

The six computers in the room had the *Perfect Circle* game installed with three different conditions (two computers for each condition):

1. In the first condition the game was installed without the SGD Model.

2. In the second condition the game was installed with the SGD Model and the group had neutral social relations in the beginning of the game.

3. In the third condition the game was installed with the SGD Model but the members of the group started with negative social attraction relations, thus, the level of cohesion of the group was very low.

Furthermore, apart from the differences that were mentioned, all the other details were similar in the three conditions. The four synthetic characters had the same name, the same appearance, the same personality and the same skills. In addition, the sequence of the game puzzles was predefined and the same for all the subjects. This sequence was randomly generated beforehand.

The subjects were selected on the fly in the beginning of each session and they freely chose which computer to use. In the first half-hour, the subjects read the game instructions that were previously distributed and filled the first part of the questionnaire, which included the items related to their gender, age, gaming experience and personality.

After filling the first part of the questionnaire, the subjects could start playing the game. They first created their own character and played the game for one hour. Then, after playing the game, the subjects were asked to fill the second part of the questionnaire, which included the items related to the group trust, the group identification, the social presence, the satisfaction and the open question.

This process was repeated in the four sessions, which, in the end, gave a sample of eight subjects for each of the conditions.

8.6 Results

From the experiment conducted as described in the previous section (8.5) we obtained twenty four complete questionnaires. These questionnaires were compiled in order to find the values that assess the several variables of the study. The results of this compilation are summarised in three different tables, each one containing the results of one of the conditions of the experiment: table 8.3 for the first condition, table 8.4 for the second one and table 8.5 for the third one.

Each row in these tables represents one of the eight subjects that were involved in that particular condition and their columns represent the different variables of the experiment: the subjects' id (*Subj*), their age (*Age*), their gender (*Gen*), their gaming experience (*Gam*), the five factors of their personality (*I - Surgency or Extraversion; II - Agreeableness; III - Conscientiousness; IV - Emotional Stability; V - Intellect or Imagination*), their trust in the group (*Trust*), their identification with the group (*Ident*), their sense of social presence (*Pres*) and, finally, their levels of satisfaction (*Satis*).

In the next sections we will present the analysis of these results concerning the four dependent variables: the group trust, the group identification, the social presence and the satisfaction with the game. In addition, we will present the analysis of the sample that we performed to verify its levels of homogeneity among the different conditions.

All the analyses were performed using non-parametric tests, such as the Kruskal-Wallis and the Spearman's Rank Correlation [159] tests, because they are more reliable for small size statistical samples and in scenarios with subjective scales, which is the case of this one.

Subj	Age	Gen	Gam	I	II	III	IV	V	Trust	Ident	Pres	Satis
1	23	M	6,00	50	60	47	55	60	4,33	3,33	3,67	5,83
2	23	M	4,17	45	52	55	55	57	3,83	4,17	4,33	4,50
3	20	M	5,83	37	75	67	75	70	4,50	4,50	4,67	6,33
4	20	F	5,17	58	52	48	58	45	3,50	3,83	4,50	5,00
5	20	M	6,17	55	85	88	88	67	6,33	6,33	5,83	6,00
6	21	F	5,00	47	87	77	55	50	3,00	3,33	4,00	5,50
7	18	M	5,00	53	75	70	27	53	3,50	3,83	3,83	5,50
8	20	M	4,83	93	77	42	47	70	3,83	3,17	4,17	3,83

Table 8.3: The results in the first condition - No SGD Model.

Subj	Age	Gen	Gam	I	II	III	IV	V	Trust	Ident	Pres	Satis
9	32	M	6,33	75	83	78	83	85	6,50	4,50	4,00	5,83
10	28	M	6,50	57	63	60	60	78	5,83	4,33	3,67	5,67
11	20	M	5,83	43	50	40	72	47	6,00	5,67	4,83	5,33
12	23	M	3,83	30	68	52	72	58	4,33	5,00	3,50	4,67
13	20	M	5,17	53	70	68	60	63	5,17	4,33	5,00	5,17
14	20	M	5,67	17	55	48	47	15	3,67	4,50	5,00	4,33
15	20	M	5,00	53	82	63	87	62	5,50	6,17	5,00	4,33
16	21	F	4,17	58	65	58	63	55	4,50	3,33	3,33	3,33

Table 8.4: The results in the second condition - Neutral Group.

8.6.1 Sample's Homogeneity

Before analysing the dependent variables, we analysed the sample in order to identify possible significant differences across the experiment's conditions. We wanted to assure that the three groups of samples were homogeneous enough so that we could confidently guarantee that the differences on the dependent variables were only due to the control conditions and not due to specificities of the sample.

Therefore, we have compared the subjects' ages, personalities and gaming experience across the three conditions: C1 - No SGD Model; C2 - Neutral Group; C3 - Low Cohesion Group. Regarding the subjects' gender, since there were only four female subjects, two of them were in the first condition while one was in the second and the other one was in the third, we do not expect any significant differences and, thus, we can state that the sample is homogeneous

Subj	Age	Gen	Gam	I	II	III	IV	V	Trust	Ident	Pres	Satis
17	22	M	7,00	32	63	58	25	48	5,00	5,00	5,00	6,00
18	21	F	4,83	58	62	60	55	55	5,17	5,33	4,50	5,33
19	19	M	6,33	47	78	45	48	57	6,17	6,17	4,17	5,50
20	19	M	4,00	35	60	37	60	62	4,83	4,67	4,33	5,67
21	24	M	5,50	72	85	90	87	70	4,33	3,33	4,50	5,00
22	20	M	5,67	58	65	85	45	63	5,83	6,50	5,67	6,17
23	18	M	5,00	72	65	68	45	70	5,00	5,50	3,83	5,50
24	18	M	5,67	77	65	53	78	83	4,83	5,33	6,17	6,67

Table 8.5: The results in the third condition - Low Cohesion Group.

in terms of gender.

Cond.	Mean	N	Std. Dev
1	20,625	8	1,68502
2	23,000	8	4,56696
3	20,125	8	2,10017
Total	21,250	24	3,19306

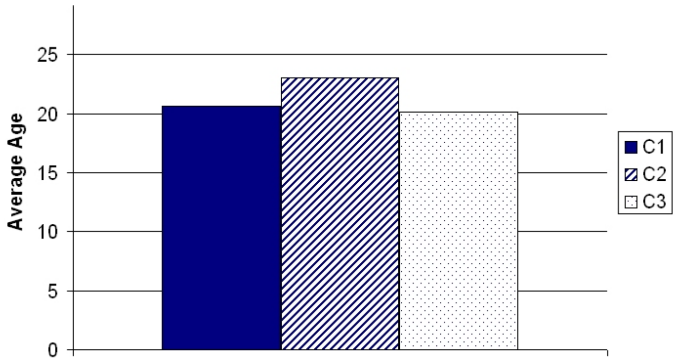

Figure 8.1: The subjects' ages across the three different conditions of the experiment.

The comparison of the subjects' ages across the three conditions is presented in figure 8.1. From the graph in this figure we can see that the differences are not significant, since the average age in each of the three conditions is very similar.

Figure 8.2 shows the comparison in terms of the subjects' gaming experience and, as we can see from the graph, the differences are not significant either.

Cond.	N	Mean Rank
1	8	11,75
2	8	12,69
3	8	13,06

Chi-square: ,147; *Asymp. Sig:* ,929

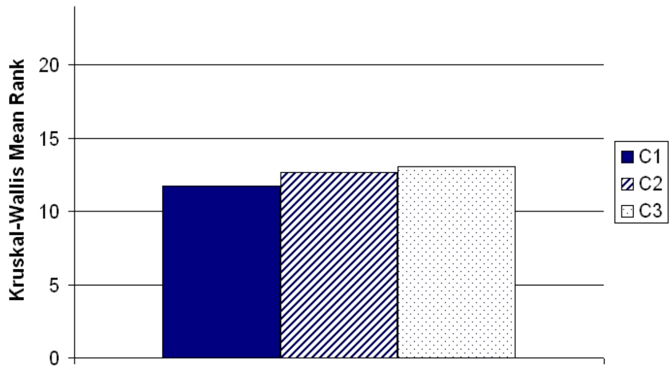

Figure 8.2: The subjects' gaming experience across the three different conditions of the experiment.

However, in figure 8.3 we can see that there are some differences in relation to the subjects' personalities. The one that is worth mentioning is the difference in the Emotional Stability (factor IV) which is much higher in the sample of the second condition.

In conclusion, we can consider that the sample is, in general, homogeneous across the three conditions, with only a small exception, which was clearly identified and that we will consider in the analysis of the dependent variables.

8.6.2 Group Trust

Concerning the subjects' trust in the group we have reached some significant results. As shown in figure 8.4, the subjects who played the game with the SGD Model had more trust in the group than those who played without the model. Furthermore, the fact that the group was initially cohesive, or not, did not influence the final levels of trust.

From these results we can conclude that the use of the SGD Model to drive the behaviour of the synthetic group has a positive influence on the users' trust in the group.

In addition, we have analysed the correlations of the group trust with the other variables of the experiment. These correlations were computed using the Spearman's Rank Correlation

Cond.	N	P-I	P-II	P-III	P-IV	P-V
1	8	12,25	13,50	12,81	11,38	11,31
2	8	10,94	11,88	11,88	16,00	12,19
3	8	14,31	12,13	12,81	10,13	14,00
Chi-square		,933	,247	,094	3,089	,605
Asymp. Sig.		,627	,884	,954	,213	,739

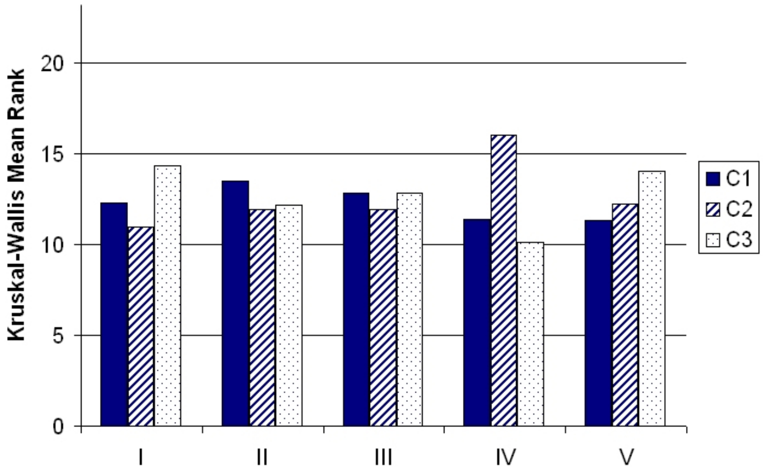

Figure 8.3: The subjects' personalities across the three different conditions of the experiment.

test [159], which is a non-parametric test for the strength of the relationship between pairs of variables. This test calculates a correlation coefficient between two variables and its relative significance. This coefficient is measured in values that range from 0 (not correlated) to 1 (completely correlated), and usually conveys a meaningful correlation when this value reach 0,5. The significance is also measured from 0 to 1 and determines the confidence in the correlation. Its values should be as low as possible, and usually only values below 0,05 are considered to assert strong significance.

The computed correlations are presented in table 8.6. This table shows the variables that are significantly correlated in ***bold and italic***, and the variables that are just slightly correlated in *italic* (e.g. variables with significance values slightly above 0,05).

We can see significant correlations between the subjects trust in the group and three other variables:

- the condition used: what confirms our first conclusion is that the SGD model has a positive effect on the users' trust in the group;

Cond.	N	Mean Rank
1	8	7,31
2	8	15,19
3	8	15,00

Chi-square: 6,492; *Asymp. Sig:* ,039

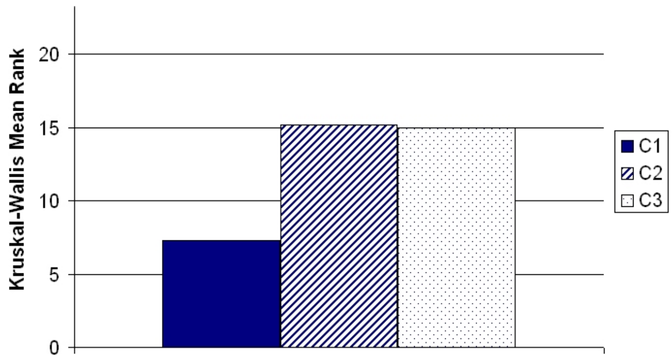

Figure 8.4: The subjects' trust in the group across the three different conditions of the experiment.

- the subjects' gaming experience, which shows that more experienced gamers have more trust in the group, thus suggesting that this type of subject benefits more from the use of the SGD Model;

- the identification with the group, which confirms the relation between trust and identification in a group (also reported in other studies [7]).

In addition, the two correlations with the subjects' satisfaction with the game and their Intellect dimension of personality, although not as strong, may suggest that:

- subjects with a higher Intellect, which relates to their openness to new experiences, are more receptive to the synthetic group, and are less reluctant to trust in its synthetic members.

- the subjects' satisfaction relates to their trust in the group, which confirms the results presented by Driscoll [42].

VARIABLE	CORRELATION	SIGNIFICANCE
Condition	***0,454***	***0,026***
Age	0,008	0,971
Gender	-0,316	0,133
Gaming	***0,505***	***0,012***
Pers F-I	0,158	0,461
Pers F-II	0,137	0,524
Pers F-III	0,173	0,419
Pers F-IV	0,319	0,129
Pers F-V	*0,395*	*0,056*
Identification	***0,714***	***0,000***
Social Presence	0,277	0,190
Satisfaction	*0,389*	*0,061*

Table 8.6: Correlations of the experiment variables with the group trust.

8.6.3 Group Identification

The results regarding the subjects' identification with the group are similar to those verified concerning the group trust. Which means that the SGD Model also has effect on the subjects' social identification with the group. Figure 8.5 shows these results and, as we can see, the identification with the group is higher in the two conditions where the synthetic characters used the SGD Model to drive their behaviours.

In addition, there are differences in this identification concerning the initial group cohesion. It seems that the most cohesive group induced lower levels of identification in the subjects. We believe that this effect may be related to the fact that the socio-emotional interactions in the highest cohesion group are essentially positive, which is probably less believable than a scenario where both, positive and negative, socio-emotional interactions occur, as in the case of the third condition.

Furthermore, with the analysis of the correlations between the subjects' identification with the group and the other variables in the experiment, that are shown in table 8.7, we can draw some more conclusions:

- First of all, the strong correlation between the identification with the group and the experiment's control condition confirms our conclusion that the SGD model has a positive

Cond.	N	Mean Rank
1	8	8,00
2	8	12,94
3	8	16,56

Chi-square: 5,960; *Asymp. Sig:* ,051

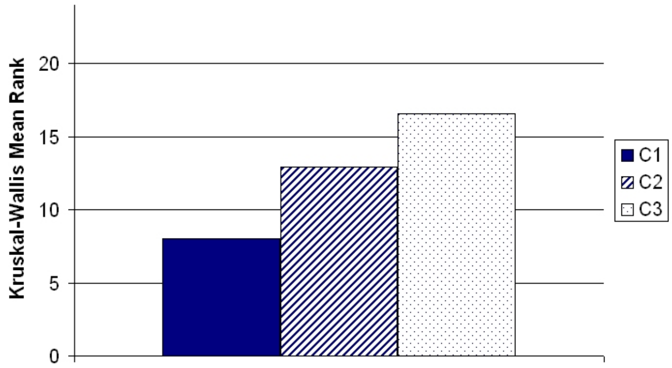

Figure 8.5: Subjects' identification with the group across the three different conditions of the experiment.

effect on the users' identification with the group.

- The strong correlation between identification and trust, as stated before, complies with the results reported in other studies [7].
- The correlation with the sense of social presence suggests that achieving a strong identification with a group of synthetic characters is a good way of increasing the social presence of the user in a virtual environment, which confirms the basis of our hypothesis.
- In addition, there seems to be a correlation between the subjects' age and their identification with the group, which indicates that the older subjects are more reluctant to identify themselves with the synthetic group. However, the age range of our sample (from 19 to 30) is very small in order to confirm such conclusion.
- Finally, although only slightly significant, there is a correlation between the subjects' identification with the group and their satisfaction with the game, which suggests that the users have more fun with the game if they are able to identify themselves with the group.

VARIABLE	CORRELATION	SIGNIFICANCE
Condition	***0,507***	***0,011***
Age	*-0,373*	*0,073*
Gender	-0,308	0,143
Gaming	0,258	0,224
Pers F-I	-0,135	0,531
Pers F-II	-0,022	0,919
Pers F-III	0,066	0,761
Pers F-IV	0,107	0,620
Pers F-V	0,074	0,730
Trust	***0,714***	***0,000***
Social Presence	***0,516***	***0,010***
Satisfaction	*0,376*	*0,070*

Table 8.7: Correlations of the experiment variables with the identification with the group.

We can additionally state, recalling the fact that the sample in the second condition had higher levels of Emotional Stability in general (factor IV), that this dimension of the subjects' personality may affect their perception of the group and explain the differences between the results of conditions two (Neutral Group) and three (Low Cohesion Group). However, the statistical significance of this hypothesis seems to be very low, since the correlation between the fourth factor of personality and the identification with the group is only 0,107.

8.6.4 Social Presence

The comparison of the levels of social presence in the subjects across the three conditions are not very conclusive. As we can see from figure 8.6 there is a positive effect in the third condition, but not in the second one. We believe that this fact is related to the lower levels of identification verified in the second condition. In fact, the analysis of the correlations between the social presence and the other variables of the experiment (see table 8.8) confirms this hypothesis since there is a strong positive correlation between these two variables, social presence and identification.

Therefore, we can conclude that because the SGD Model has a positive effect on the subjects' identification with the group, it can indirectly influence their sense of social presence.

Furthermore, the analysis of the correlations did not identify any other correlation with

Cond.	N	Mean Rank
1	8	11,31
2	8	11,19
3	8	15,00

Chi-square: 1,514; *Asymp. Sig:* ,469

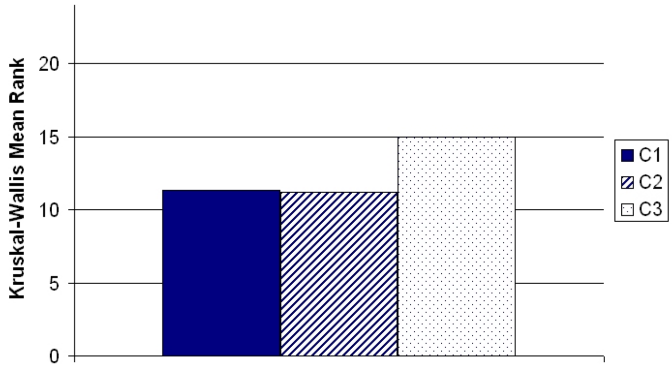

Figure 8.6: Subjects' sense of social presence across the three different conditions of the experiment.

the social presence other than the identification with the group.

8.6.5 Satisfaction with the Game

Concerning the subjects' satisfaction with the game, we have had some interesting results. As figure 8.7 shows, the general satisfaction was the highest in the case of the third condition and it was the lowest in the case of the second condition.

This effect is surprising as it contradicts, in a certain way, the other results. Regarding the other three variables, the effects of the SGD Model were always positive, although in the case of the social presence the effects were not very explicit for the second condition. Nevertheless, in the third condition, where the group was initially non cohesive, the positive effect still applies to the satisfaction with the game. So there was a particular element of the second condition that did not please the players. Our hypothesis for this is that since the socio-emotional interactions in a cohesive group are more likely to be positive, the subjects did not find the group itself to be challenging, and therefore were more bored with the group interactions. However, with the current variables of this study we can not confirm this hypothesis with confidence.

Furthermore, with the analysis of the correlations between the subjects' satisfaction with

VARIABLE	CORRELATION	SIGNIFICANCE
Condition	0,218	0,305
Age	-0,329	0,116
Gender	-0,219	0,304
Gaming	0,297	0,159
Pers F-I	-0,030	0,889
Pers F-II	-0,022	0,920
Pers F-III	0,143	0,506
Pers F-IV	0,172	0,421
Pers F-V	0,011	0,960
Trust	0,277	0,190
Identification	***0,516***	***0,010***
Satisfaction	0,276	0,191

Table 8.8: Correlations of the experiment variables with the social presence.

the game and the other variables in the experiment, that are shown in table 8.9, we can additionally state that:

- The correlation between the subjects' gaming experience and their satisfaction with the game suggested that subjects with more experience and interested in computer games enjoyed the game more. This fact alone, may not be very significant since this is the type of subject that would probably like the game anyway, but on the other hand this may also suggest that more experienced gamers benefit more from the SGD Model. However, again, we cannot confirm this with the current study.

- There is also a slight correlation with the subjects' personality on the Intellect dimension (factor V). A similar correlation was already affirmed to have some effects on the subjects' trust, which in turn affects the subjects' satisfaction.

- In addition, there is a correlation with the subjects' identification with the group, which was already discussed in section 8.6.3. It suggests that the users have more fun with the game if they are able to identify themselves with the group.

Cond.	N	Mean Rank
1	8	12,94
2	8	8,56
3	8	16,00

Chi-square: 4,503; *Asymp. Sig:* ,105

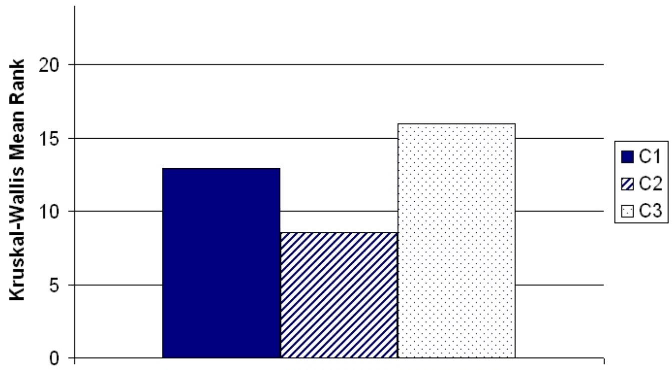

Figure 8.7: Subjects' general satisfaction with the game across the three different conditions of the experiment.

8.7 Concluding Remarks

In this chapter we have described the experiment that was conducted in order to evaluate the effects of the model, designed to support the believability of the interactions of groups of synthetic characters, in the interaction experience of users with such groups. This experience was conducted with several subjects that used the *Perfect Circle* game to engage in collaboration with a group of four autonomous synthetic characters.

The users' interaction experience was measured in terms of four different variables: their trust in the group, their identification with the group, their sense of social presence in the virtual environment and their general satisfaction with the game. Additionally, the experiment was run with three control conditions that, besides controlling the use of the SGD Model, also controlled the initial cohesion of the group.

The results have shown that the use of the SGD Model to drive the behaviour of the synthetic characters in the group had a positive effect on the users' interaction, which confirms the hypothesis that grounded this research:

"If the interactions in a group of synthetic characters follow similar dynamics as

VARIABLE	**CORRELATION**	**SIGNIFICANCE**
Condition	0,181	0,397
Age	-0,135	0,528
Gender	-0,284	0,179
Gaming	***0,604***	***0,002***
Pers F-I	0,019	0,929
Pers F-II	0,095	0,660
Pers F-III	0,264	0,213
Pers F-IV	0,022	0,917
Pers F-V	*0,381*	*0,066*
Trust	*0,389*	*0,061*
Identification	*0,376*	*0,070*
Social Presence	0,276	0,191

Table 8.9: Correlations of the experiment variables with the general satisfaction with the game.

the interactions in human groups do then the users' interaction experience is in fact improved."

Furthermore, several particular conclusions could be achieved from this study. We can highlight the most important ones in the following items:

- The SGD Model improves the users' trust in the synthetic group, despite any differences in the level of cohesion of the group.
- The SGD Model induces a higher level of identification with the group in the users. However this identification seems to depend on the level of cohesion of the group.
- The users' sense of social presence can be achieved through their identification with the group and can, therefore, be improved by the SGD Model.
- Trust and identification with the group are connected to the satisfaction with the group interactions. Therefore, the SGD Model may also have a positive effect on the users' satisfaction with the game. However, this satisfaction seems to highly depend on the level of cohesion of the group.
- Very cohesive groups induce lower levels of identification with the group and, consequently, lower levels of social presence and satisfaction with the game.

Chapter 9

Conclusions

9.1 Synthesis

With the emergence of synthetic characters, collaborative virtual environments can now be populated with synthetic characters and users at the same time, all interacting, collaborating or competing with each other. However, the users' interaction with the synthetic characters is not always the best, and it is only positive if the characters are able to show a coherent and believable behaviour. Therefore, in scenarios where these characters engage as a group in a collaborative task, they should be able to show believable group dynamics.

Furthermore, most of the research conducted on the believability of synthetic characters is centred on the interactions between a user and a single character without considering the user within the group and without a common collaborative task. Thus, we found a clear need for research on scenarios that engage the user in collaboration with a group of synthetic characters.

This work addresses the problem of creating groups of synthetic characters and promoting the collaboration of such groups with users, while making the emergent interactions and collaboration, natural and believable. To achieve this, we grounded our developments on the hypothesis that the interactions and dynamics of the synthetic group should resemble the collaboration and interactions that emerge in humans groups. Doing so, we believed that the synthetic group would become more believable and consequently improve the interaction experience of a user that actively participates in the group.

Following this idea we have designed a model (Synthetic Group Dynamics Model - SGD Model) to support the behaviour of a cognitive synthetic agent in a group scenario, inspired by several theories of human group dynamics developed in socio-psychological studies. This model defines the group process as a set of interactions that occur between the several agents, which are grounded on the knowledge that each agent builds regarding the group, the other members and the social context.

One important facet of this knowledge concerns the social relations that the members establish with each other and that constitute the main structuring mechanism of the group. Each member of the group should model two different kinds of social interactions: the social attraction relations, which are related to the liking and disliking attitudes, and the social influence relations, which are related to the social power.

In addition, the SGD Model defines some rules for the group dynamics that determine, on one hand, the conditions for the occurrence of the interactions and, on the other hand, the effects that such interactions have on the group. These dynamics are grounded on a categorization of the several kinds of interactions that occur in the group that divide the interactions into two main classes: the socio-emotional interactions and the instrumental interactions (related to the task).

The SGD Model was implemented in the behaviour of the synthetic characters that are part of a game, the *Perfect Circle*, that engages the players in a group of four synthetic characters. These characters play the role of adventurers that wander in a fantasy world and face several challenges in their quest for a sacred item. The group must interact and collaborate in order to overcome each consecutive challenge. Thus, they can propose actions, manipulate objects, express their opinions about the current proposals and actions, and can encourage or discourage the other members.

Moreover, the integration of the behaviour of these characters in the virtual environment, that supported the game, was achieved through a framework that we also developed in the context of this work. The main goal of this framework is to promote a clean integration of the several components that usually constitute the system that simulates a virtual environment, namely the integration of the visual and the behavioural components, using a multi-agent approach

The *Perfect Circle* game was developed with the main purpose of evaluating the SGD Model. Therefore, it was used in an experiment conducted in our university, in order to evaluate the effects of the SGD Model on the interaction experience of a user with a synthetic group. This experiment was run with three different control conditions that, besides controlling the use of the SDG Model, also controlled the initial cohesion of the group.

The users' interaction experience was measured in terms of four different variables: their trust in the group, their identification with the group, their sense of social presence in the virtual environment and their general satisfaction with the game.

The experiment's final results show that the use of the SGD Model to drive the behaviour of the synthetic characters in the group had a positive effect on the users' interactions, specially on the users' trust and identification with the synthetic group. These results confirm the hypothesis that grounded this research and, therefore, we can conclude that if the interactions in a group of synthetic characters follow similar dynamics as the interactions in human groups do then the users' interaction experience is in fact improved.

9.2 Contributions

This research has produced some contributions that we believe are relevant for the research fields of multi-agent systems and synthetic characters.

First, we proposed a framework for the integration of the behavioural components and the visual components of a virtual environment through the use of agents in a multi-agent system. This framework is flexible and can be used, for example, to support the integration of several different views on a simulated world in order to support the collaboration of people with different needs, e.g. between sighted and unsighted people.

In addition, we have developed a model to support the dynamics of the interactions in groups with synthetic and human participants. We have successfully used the model and confirmed that it has some positive results on the users' interaction in the group. Thus, we believe that this model can be applied to other contexts where the synthetic agents engage in a collaborative task.

Finally, we think that this work draws attention to some issues concerning user interaction with synthetic characters that have not been explored by the research community before, in particular, the interaction in collaborative groups. We believe that we have highlighted the importance of the believability of group interactions in synthetic characters and established some bases for future work in the field.

9.3 Future Work

The SGD model, that constitutes the basis of this work, used a very simplified view of the group process. Thus, there are aspects of the group dynamics that could be further explored. And, in addition, the results of the experiment that we have conducted also suggest some interesting issues for future research.

The following items describe some of the open issues that we believe are relevant and, thus, deserve some attention as future work:

- **Dynamics of the group composition:** this work did not address the dynamics of the group regarding its composition. For example, we did not establish rules for the acceptance of new members in the group, or rules to support the decision of a member to leave the group. This, among other things, is related to the individual expectations concerning the group (e.g. the goal that led the individual to join the group) which is not considered in our model, and, thus, should be addressed in the future.

- **More complex group structures:** we have also simplified the structure of our groups. For example, the structure of communication was not considered and, additionally, there are not specific organizational roles. The model can be extended to include more structured groups, in particular to include leadership roles.

- **Context factors:** there are many open issues concerning the context factors, and our model does not deal with any particular task model or social norm. Thus, the SGD Model can be extended with a particular task model built for certain types of tasks, or with some mechanisms for handling social norms.

- **Relations with other groups:** another aspect that was not handled by this model was the inter-group relations. The model was centred on the interaction of the group members with each other without concerning their interaction with external groups. This, would be a very interesting extension and would be essential for scenarios where several different groups engage in collaboration or competition.

- **Further evaluation:** we believe that there are several issues that should be further evaluated with the current version of the model:

  - **Leadership:** despite the fact that the concept of leadership is not explicitly modelled, spontaneous leaders can emerge with the current model. Since these can be simply defined as the members that interact more with the group and that achieve the best position in the group, the frequency of interaction and group position are two concepts included in the model. Therefore, we believe that the study of these emergent leaderships may lead to some interesting results.
  - **The users personality:** the evaluation results have shown some correlations between the fifth dimension of personality (Intellect) of the users and their acceptance of the group, namely in their satisfaction and trust in the group. We think that this correlation should be further studied.
  - **The user's cultural context:** we believe that an interesting study could be the comparison of results from samples of different societies that have distinct cultures concerning collaboration. For example, there is some evidence that the western societies are less cooperative than eastern societies. We could verify if this fact has an impact on the user's experience with the SGD Model.
  - **Group cohesion:** another interesting result from the evaluation was the fact that users seem to prefer the interactions of groups with a lower level of cohesion. We believe that this is also an interesting issue to be further researched.

- **Apply the model to different applications:** we believe that the SGD Model can be used in different applications. We explore the possible ideas below:

  - **A tool for training leadership:** a possible application of this work could be in systems for training team leaders. For example, managing a group of people requires more than managing the members' abilities and coordinating them in the resolution process of the tasks. There are always issues concerning members' motivation and social relations to be handled. In addition, people that manage working teams are not often aware of these issues and do not have the proper social skills to handle them. We believe that they could improve these skills if they could interact with a believable simulation of a group where socio-emotional issues are important. The SGD Model could support the creation of such simulations.
  - **A tool for social studies:** the SGD Model could support the creation of an application to be used by sociologists in their studies of group dynamics. Studies

that evaluate a single subject often use several other people to play a given role in the group just to create the evaluation scenario. However, it is not usually easy to find and train such human resources. Therefore, these types of studies are very difficult and expensive to conduct. The ideas of the SGD Model could be used to support the creation of an application that uses synthetic characters to play the role of the people that are usually required to simulate the group and conduct the study.

Appendix A

Personality Questionnaire

A.1 Introduction

This appendix presents the complete list of the items included in the IPIP personality questionnaire. Each item is presented in two languages: the original version in English and the version translated to Portuguese. In addition, each item has an indication that identifies the factor that it assess and the positive or negative nature of the sentence. For example the following item: *"Am the life of the party. (I+)"* is related to the first factor (I) and expresses a positive comment about it (+).

A.2 The Questionnaire Items

1. Am the life of the party. (I+)
   Sou a alegria de uma festa.

2. Feel little concern for others. (II-)
   Preocupo-me pouco com os outros.

3. Am always prepared. (III+)
   Estou sempre preparado.

4. Get stressed out easily. (IV-)
   Entro em stress facilmente.

5. Have a rich vocabulary. (V+)
   O meu vocabulário é rico.

6. Don't talk a lot. (I-)

   Não sou muito falador.

7. Am interested in people. (II+)

   Interesso-me pelas pessoas.

8. Leave my belongings around. (III-)

   Largo os meu pertences onde calha.

9. Am relaxed most of the time. (IV+)

   Sou calmo na maioria do tempo.

10. Have difficulty understanding abstract ideas. (V+)

    Tenho dificuldade em compreender ideias abstractas.

11. Feel comfortable around people. (I+)

    Sinto-me confortável com outras pessoas.

12. Insult people. (II-)

    Se achar necessário não tenho problemas em insultar uma pessoa.

13. Pay attention to details. (III+)

    Presto atenção aos pormenores.

14. Worry about things. (IV-)

    Preocupo-me com as coisas.

15. Have a vivid imagination. (V+)

    Tenho uma imaginação fértil.

16. Keep in the background. (I-)

    Mantenho-me longe das atenções.

17. Sympathize with others' feelings. (II+)

    Simpatizo com os sentimentos dos outros.

18. Make a mess of things. (III-)

    Faço uma confusão das coisas.

19. Seldom feel blue. (IV+)

    Raramente me sinto deprimido.

20. Am not interested in abstract ideas. (V-)

    As ideias abstractas não me interessam.

21. Start conversations. (I+)

    Inicio conversas.

22. Am not interested in other people's problems. (II-)

    Os problemas dos outros não me interessam.

23. Get chores done right away. (III+)

    Executo as minhas tarefas prontamente.

24. Am easily disturbed. (IV-)

    Sou facilmente perturbado.

25. Have excellent ideas. (V+)

    Tenho ideias excelentes.

26. Have little to say. (I-)

    Tenho pouco a dizer.

27. Have a soft heart. (II+)

    Tenho um coração mole.

28. Often forget to put things back in their proper place. (III-)

    Esqueço-me frequentemente de guardar as coisas no seu lugar.

29. Get upset easily. (IV-)

    Fico chateado facilmente.

30. Do not have a good imagination. (V-)

    Não tenho uma boa imaginação.

31. Talk to a lot of different people at parties. (I+)

    Falo com todo o tipo de pessoas numa festa.

32. Am not really interested in others. (II-)

    Não me interesso pelos outros.

33. Like order. (III+)

    Adoro a ordem.

34. Change my mood a lot. (IV-)

    Mudo frequentemente de disposição.

35. Am quick to understand things. (V+)

    Sou perspicaz a entender as coisas.

36. Don't like to draw attention to myself. (I-)

    Não gosto de atrair as atenções.

37. Take time out for others. (II+)

    Tenho tempo para os outros.

38. Shirk my duties. (III-)

    Evito os meus deveres.

39. Have frequent mood swings. (IV-)

    Tenho fequentes mudanças de humor.

40. Use difficult words. (V+)

    Uso palavras caras.

41. Don't mind being the center of attention. (I+)

    Não me importo de ser o centro das atenções.

42. Feel others' emotions. (II+)

    Sinto as emoções dos outros.

43. Follow a schedule. (III+)

    Sigo um horário.

44. Get irritated easily. (IV-)

    Irrito-me facilmente.

45. Spend time reflecting on things. (V+)

    Passo tempo a pensar nas coisas.

46. Am quiet around strangers. (I-)

    Fico calado perto de estranhos.

47. Make people feel at ease. (II+)

    Faço as pessoas sentir-se bem.

48. Am exacting in my work. (III+)

    Faço o meu trabalho rigorosamente.

49. Often feel blue. (IV-)

    Sinto-me deprimido frequentemente.

50. Am full of ideas. (V+)

    Sou cheio de ideias.

Bibliography

[1] The activeworlds. *(online) http://www.activeworlds.com*, last access on Jan 2005.

[2] Scott Acton. Great ideas in personality - theory and research. *(online) http://www.personalityresearch.org/bigfive.html*, last access on Jan 2005.

[3] H. Aebli. *A Didactics founded on the psychology of Jean Piaget*. Kapelusz, 1951.

[4] Kathleen Allen, Richard Bergin, and Kenneth Pickar. Exploring trust, group satisfaction, and performance in geographically dispersed and co-located university technology commercialization teams. In *Proceedings of the NCIIA 8th Annual Meeting: Education that Works*, March 18-20, 2004.

[5] E. Andre, J. Muller, and T. Rist. The ppp persona: a multipurpose animated presentation agent. In *Proceedings of the International Workshop on Advanced Visual Interfaces*, pages 245–247, Gubbio, Italy, 1996. ACM Press.

[6] J. Ang and P. H. Soh. User information satisfaction, job satisfaction and computer background: an exploratory study. *Information and Management*, 32:255–266, 1997.

[7] B. E. Ashforth and F. A. Mael. Social identity and the organization. *Academy of Management Review*, 14:2–39, 1989.

[8] John L. Austin. *How to do things with words*. Oxford University Press, 1962.

[9] Robert Axelrod and William D. Hamilton. The evolution of cooperation. *Science*, 211:1390–1396, 1981.

[10] K. W. Back. Influence through social communication. *Journal of Abnormal and Social Psychology*, 46:9–23, 1951.

[11] Norman I. Badler, C. B. Phillips, and B. L. Webber. Simulating humans. *Computer Graphics Animation and Control*, 1993.

[12] Robert Freed Bales. *Interaction Process Analysis.* The University of Chicago Press, Chicago, 1950.

[13] Robert Freed Bales and P. Slater. *Family, socialization and interaction process*, chapter Role differentiation in small decision-making groups. Free Press, Glencoe, Ill, 1955.

[14] Robert Freed Bales, F. L. Strodtbeck, T. M. Millis, and M. E. Roseborough. Channels of communication in small groups. *American Sociology Review*, (16):461–468, 1951.

[15] B. M. Bass and F. T. M. Norton. Group size and leaderless discussions. *Journal of Applied Psychology*, (35):397–450, 1951.

[16] Josheph Bates. Virtual reality, art and entertainment. *Presence: teleoperators and Virtual Environments*, 1992.

[17] Josheph Bates. The role of emotions in believable characters. *Communications of the ACM*, 37(7):122–125, 1994.

[18] S. Benford, C. Greenhalgh, and D. Lloyd. Crowded collaborative virtual environments. In *Proceedings of the Conference on Human Factors in Computing Systems - CHI'97*, Atlanta, Georgia, USA, 1997.

[19] Timothy Bickmore and Justine Cassell. Relational agents: a model and implementation of building user trust. In *Proceedings of the Conference on Human Factors in Computing Systems - CHI'2001*, Seattle, USA, 2001. ACM Press.

[20] W. R. Bion. Experiences in groups, i-vi. *Human Relations*, pages 314–320(1), 487–496(2), 13–22(2), 295–303(2), 3–14(3), 295–402(3), 1948-1950.

[21] Bioware. Star wars: knights of the old republic. *(online) http://www.lucasarts.com/products/swkotor/*, 2003.

[22] Bruce Blumberg. (void*): A cast of characters. In *Visual proceedings of the conference on SIGGRAPH 99*, page 169. ACM Press, 1999.

[23] Michael E. Bratman. *Intentions, Plans and Practical Reasoning*. Harvard University Press, Cambridge, Cambridge, 1987.

[24] Rodney Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, pages 14–23, 1986.

[25] D. Cartwright and F. Harary. A generalization of heider's theory. *Psychological Review*, (63):277–292, 1956.

[26] Dorwin Cartwright and Alvin Zander. *Group Dynamics: research and Theory*. Harper and Row, New York, 1968.

[27] Justine Cassell, M. Ananny, A. Basu, Timothy Bickmore, P. Chong, D. Mellis, K. Ryokai, J. Smith, H. Vilhjálmsson, and H. Yan. Shared reality: physical collaboration with a virtual peer. In *Proceedings of the Conference on Human Factors in Computing Systems - CHI'2000*, The Hague, The Netherlands, 2000. ACM Press.

[28] Justine Cassell, T. Bickmore, M. Billinghurst, L. Campbell, K. Chang, H. Vilhjálmsson, and H. Yan. Embodiment in conversational interfaces: Rea. In *Proceedings of the Conference on Human Factors in Computing Systems - CHI'99*, pages 520–527, Pittsburgh, USA, 1999. ACM Press.

[29] Justine Cassell, T. Bickmore, L. Campbell, H. Vilhjalmsson, and H. Yan. *Embodied Conversational Agents*, chapter Human Conversation as a System Framework: designing Embodied Conversational Agents. MIT Press, 2000.

[30] Justine Cassell and H. Jenkins. *From Barbie to Mortal Kombat: Gender and Computer Games*. MIT Press, 1998.

[31] Justine Cassell, Joseph Sullivan, Scott Prevost, and Elizabeth Churchill. *Embodied Conversational Agents*. MIT Press, 2000.

[32] Cristiano Castelfranchi. Guaranties for autonomy in cognitive agent architectures. In *Intelligent Agents - ECAI-94 workshop on Agent Theories, Architectures, and languages*, pages 56–70, The Netherlands, 1994. SpringerVerlag.

[33] R. B. Cattell. Concepts and methods in the measurement of group syntality. *Psychological Review*, (55):48–63, 1948.

[34] B. E. Collins and B. H. Raven. *The Handbook of Social Psychology*, volume 4, chapter Group Structure: attraction, coalitions, communication, and power. Addison-Wesley, Reading, Massachusetts, 2000.

[35] M. K. D. Coomans and H. J. P. Timmermanns. Towards a taxonomy of virtual reality user interfaces. In *Proceedings of the International Conference on Information Visualisation*, London, 1997.

[36] C. Cruz-Neira, D. Sandin, and T. DeFasti. Surround-screen projection based virtual reality. the design and implementation of the cave. In *Proceedings of the SIGGRAPH*, 1993.

[37] A. Damásio. *Descartes' Error: Emotion, Reason, and the Human Brain.* G. P. Putnam's Sons, 1994.

[38] Kerstin Dautenhahn, Alan H. Bond, Lola Cañamero, and Bruce Edmonds. *Socially Intelligent Agents: Creating Relationships with Computers and Robots.* Kluwer Academic Publishers, 2002.

[39] P. DeGreef and W. A.I Jsselsteijn. Social presence in a home tele-application. *CyberPsychology and Behaviour*, 4:307–316, 2001.

[40] William J. Doll and Gholamreza Torkzadeh. The measurement of end user computing satisfaction. *MIS Quarterly*, 12(2):259–274, 1988.

[41] P. Doreian, R. Kapuscinski, D. Krackhardt, and J. Szczypula. A brief history of balance through time. *Journal of Mathematical Sociology*, 21:113–131, 1996.

[42] J. W. Driscoll. Trust and participation in organizational decision making as predictors of satisfaction. *Academy of Management Journal*, 21:44–56, 1978.

[43] EA-Games. The sims homepage. *(online) http://www.thesims.com*, 2000.

[44] EA-Games. The sims2 homepage. *(online) http://www.thesims2.com*, 2004.

[45] Michael E. Bratman et al. Plans and resource-bounded practical reasoning. *Computational Intelligence*, 1988.

[46] Nick R. Jennings et al. Automated negotiation: prospects, methods and challenges. *International Journal of Group Decision and Negotiation*, 2001.

[47] R. S. Patil et al. The darpa knowledge sharing effort: progress report. *Knowledge Representation and Reasoning*, 1992.

[48] S. Parsons et al. Agents that reason and negotiate by arguing. *Journal of Logic and Computation*, 1998.

[49] T. Finn et al. Specification of the kqml agent communication language. Technical report, DARPA knowledge sharing initiative external interfaces working group, 1993.

[50] L. Festinger. *A Theory of Cognitive Dissonance*. Stanford University Press, Stanford, CA, 1957.

[51] FIPA. Specification part 2 - agent communication language. Technical report, Foundation for Intelligent Physical Agents, 1999.

[52] FIPA. The foundation for intelligent physical agents. *(online) http://www.fipa.org*, last access on Jan 2005.

[53] John R. P. French. *Group Dynamics: Research and Theory*, chapter A Formal Theory of Social Power. Harper and Row, New York, 1968.

[54] John R. P. French and Bertram H. Raven. *Group Dynamics: Research and Theory*, chapter Bases of Social Power. Harper and Row, New York, 1968.

[55] S. Freud. *Group Psychology and the Analysis of the Ego*. Hogarth, London, 1922.

[56] Ernest Friedman-Hill. *Jess in Action. Java Rule-based Systems*. Manning Publications, 2003.

[57] Ernest Friedman-Hill. Jess, the rule engine for the java. *(online) http://herzberg.ca.sandia.gov/jess/*, 2005.

[58] M. R. Genesereth, M. Ginsberg, and J. S. Rosenschein. Cooperation without communication. In *5th National Conference on Artificial Intelligence*, Philadelphia, PA - USA, 1986.

[59] M. Georgeff and A. Lansky. Reactive reasoning and planning. In *Sixth National Conference on Artificial Intelligence*, 1987.

[60] M. Gilbert. Multi-modal argumentation. *Philosophy of the Social Sciences*, 1994.

[61] L. R. Goldberg. An alternative "description of personality": the big-five factor structure. *Journal of Personality and Social Psychology*, 1990.

[62] Lewis R. Goldberg. *Personality Psychology in Europe*, chapter A broad-bandwidth, public domain, personality inventory measuring the lower-level facets of several five-factor models, pages 7–28. Tilburg University Press, Tilburg, The Netherlands, 1999.

[63] C. Greenhalgh and S. Benford. Massive: a distributed virtual reality system incorporating spatial trading. In *15th International Conference on Distributed Computing Systems (ICDCS'95)*, Vancouver, Canada, 1995.

[64] E. Gross. Primary functions of the small group. *American Journal of Sociology*, (60):24–30, 1954.

[65] J. R. Hackman. Effects of task characteristics on group products. *Journal of Experimental Social Psychology*, (4):162–187, 1968.

[66] J. R. Hackman. Toward understanding the role of tasks in behavioral research. *Acta Psychologica*, (31):97–128, 1969.

[67] J. R. Hackman and C. G. Morris. *Group Processes*, chapter Group process and group effectiveness: a reappraisal. Academic Press, New York, 1978.

[68] A. Hare. *Handbook of Small Group Research.* Free Press, New York, 1976.

[69] Barbara Hayes-Roth. An architecture for adaptive intelligent systems. *Artificial Intelligence: special Issue on Agents and Interactivity*, 72:329–365, 1995.

[70] Fritz Heider. Attitudes and cognitive organization. *Journal of Psychology*, (21):107–112, 1946.

[71] Fritz Heider. *The Psychology of Interpersonal Relations.* Wiley, New York, 1958.

[72] J. K. Hemphill. *Group dimensions: a manual for their measurement.* Columbus: Ohio State University, Bureau of Business Research, 1956.

[73] Lawrence J. Hettinger. Virtual environment technology: its past, present, and future. *(online) http://nmrc.bu.edu/seminars/lectures/hettinger/*, 2000.

[74] Kristina Höök, Adrian Bullock, Ana Paiva, Marco Vala, Ricardo Chaves, and Rui Prada. Fantasya and sentoy. In *Extended Abstracts on the Conference on Human Factors in Computing System - CHI2003*. ACM Press, 2003.

[75] H. J. Holt. The social labouring effect: a study on the effect of social identity on group productivity in real and national groups using ringelman's methods (unplublished). *cited by R. J. Brown in Group Process*, 1987.

[76] C. G. Homans. *The Human Group*. Harcourt, Brace, New York, 1950.

[77] Ulrich Hoppe, Andreas Lingnau, Isabel Machado, Ana Paiva, Rui Prada, and Frank Tewissen. Supporting collaborative activities in computer integrated classrooms - the nimis approach. In *Proceedings of the Sixth International Workshop on Groupware*, Madeira, Portugal, 2000. IEEE Press.

[78] M. J. Huber and J. Leto. Jam: a bdi-theoretic mobile agent architecture. In *Proceedings of the Third International Conference on Autonomous Agents*, Seattle, Washington, USA, 1999. ACM Press.

[79] N. P. Hummon and T. J. Fararo. Assessing hierarchy and balance in dynamic network models. *Journal of Mathematical Sociology*, (20):145–159, 1995.

[80] Norman P. Hummon and Patrick Doreian. Some dynamics of social balance processes: bringing heider back into balance theory. *Social Networks*, (25):17–49, 2003.

[81] J. I. Hurwitz, R. F. Zander, and B. Hymovitch. *Group Dynamics: Research and Theory*, chapter Some effects of power on the relations among group members. Harper and Row, New York, 1968.

[82] K. Inkpen, K. S. Booth, M. Klawe, and R. Upitis. Playing together beats playing apart, especially for girls. In *Proceedings of Computer Support for Collaborative Learning - CSCL'95*. ACM Press, 1995.

[83] K. Inkpen, R. Upitis, M. Klawe, J. Lawry, A. Anderson, M. Ndunda, K. Sedighian, S. Leroux, and D. Hsu. We have never-forgetful flowers in our garden: girls' responses to electronic games. *Journal of Computers in Mathematics and Science Teaching*, 1994.

[84] IPIP. International personality item pool. a scientific collaboratory for the development of advanced measures of personality traits and other individual differences. *(online) http://ipip.ori.org*, last access on Jan 2005.

[85] Nick R. Jennings. Commitment and conventions: the foundation of coordination in multi-agent systems. *The knowledge Engineering Review*, 8(3):223–250, 1993.

[86] Nick R. Jennings. Controlling cooperative problem solving in industrial multi-agent systems using joint intentions. *Artificial Intelligence*, 75, 1995.

[87] Nick R. Jennings. *Foundations of Distributed Artificial Intelligence*, chapter Coordination Techniques for Distributed Artificial Intelligence, pages 187–210. John Wiley and Sons, 1996.

[88] Nick R. Jennings. On agent-based software engineering. *Artificial Intelligence*, 117:277–296, 2000.

[89] Jorge Correia Jesuíno. *Psicologia Social*, chapter Estrutura e processos de grupo: interacções e factores de eficácia. Fundação Calouste Gulbenkian, 2000.

[90] M. P. Johnson, A. Wilson, B. Blumberg, C. Kline, and A. Bobick. Sympathethic interfaces: using a plush toy to direct synthetic characters. In *Proceedings of the Conference on Human Factors in Computing Systems - CHI'99*, pages 152–158, Pittsburg, USA, 1999. ACM Press.

[91] R. S. Kalawsky. *The Science of Virtual Reality*. Addison-Wesley, 1993.

[92] D. Kipnis. Does power corrupt? *Journal of Personality and Social Psychology*, (24):33–41, 1972.

[93] D. Kipnis, P. J. Castell, M. Gergens, and D. Mauch. Metamorphic effects of power. *Journal of Applied Psychology*, (61):127–135, 1976.

[94] G. E. Krasner and S. T. Pope. A cookbook for using the model-view-controller user interface paradigm in smalltalk-80. *Journal of Object-Oriented Programming*, 1988.

[95] D. A. Kravitz and B. Martin. Ringelman rediscovered: the original article. *Journal of Personality and Social Psychology*, (50):936–941, 1986.

[96] J. E. Laird, A. Newell, and P. S. Rosenbloom. Soar: an architecture for general intelligence. *Artificial Intelligence*, 1(33):1–64, 1987.

[97] J. Lester, J. Voerman, S. Towns, and C. Callaway. Cosmo: a life-like animated pedagogical agent with deictic believability. In *Proceedings of the IJCAI97 Workshop on Animated Interface Agents: making them Intelligent*, Nagoya, 1997.

[98] H. J. Levesque, P. R. Cohen, and J. H. T. Nunes. On acting together. In *8th National Conference on Artificial Intelligence*, Boston, MA - USA, 1990.

[99] Kurt Lewin. Field theory and experiment in social psychology: concepts and methods. *American Journal of Sociology*, 44:868–897, 1939.

[100] Kurt Lewin and Dorwin Cartwright. *Field Theory in Social Science*. Harper and Row, New York, 1951.

[101] R. Lippitt, N. Polansky, F. Redl, and S. Rosen. The dynamics of power. *Human relations*, (5):37–64, 1952.

[102] A. J. Lott and B. E. Lott. Group cohesiveness, communication level and conformity. *Journal of Abnormal and Social Psychology*, (64):408–412, 1961.

[103] A. J. Lott and B. E. Lott. Group cohesiveness as interpersonal attraction: a review of relationships with antecedent and consequent variables. *Psychological Bulletin*, (64):259–303, 1965.

[104] Isabel Machado, Rui Prada, and Ana Paiva. Is the wolf angry or just hungry: inspecting, disclosing and modifying characters' minds. In *Proceedings of the 5th International Conference on Autonomous Agents*, Montreal, Canada, 2001. ACM Press.

[105] Pattie Maes. Intelligent software: programs that can act independently will ease the burdens that computers put on people. *IEEE Expert Systems*, 11(6):62–63, 1996.

[106] D. G. Marquis, H. Guetzkow, and R. W. Heyns. *Groups, Leadership and Men*, chapter A social psychological study of the decision-making conference. Cognitive Press, Pittsburgh, 1951.

[107] Carlos Martinho and Ana Paiva. Pathematic agents. In *Proceedings of the 3rd International Conference on Autonomous Agents*, Seattle, USA, 1999. ACM Press.

[108] D. McAllester and D. Rosenblitt. Systematic nonlinear planning. In *Proceedings of the Ninth National Conference on Artificial Intelligence*, pages 634–639, Menlo Park, CA, USA, 1991. AAAI Press.

[109] R. McCrae and P. Costa. *The five factor model of personality: theoretical perspectives*, chapter Toward a new generation of personality theories: theoretical contexts for the five factor model, pages 51–87. Guilford, New York, 1996.

[110] Joseph E. McGrath. *Social Psychology: a Brief Introduction.* Holt, New York, 1964.

[111] Joseph E. McGrath. *Groups: Interaction and Performance.* Prentice Hall, Englewood Cliffs, New Jersey, 1984.

[112] Brian E. Mennecke and Bradley C. Wheeler. An essay and resource guide for dyadic and group task selection and usage. *ISWorld Net Research Tasks. (online) http://kelley.iu.edu/bwheeler/ISWorld/index.cfm*, 2004.

[113] J. G. Miller. Toward a general theory for the behavioral sciences. *American Psychologist*, 10:513–531, 1955.

[114] J. L. Moreno. *Who Shall Survive?* Nervous and Mental Disease Publishing Co., Washington D.C., 1934.

[115] New-Line-Productions. The lord of the rings movies official site. *(online) http://www.lordoftherings.net*, 2001.

[116] T. M. Newcomb. *Social Psychology.* Dryden, New York, 1950.

[117] T. M. Newcomb. *The Acquaintance Process.* Holt, Rinchart and Winston, New York, 1961.

[118] R. Z. Norman and F. S. Roberts. *Sociological Theories in Progress*, chapter A measure of relative balance for social structures, pages 358–391. Houghton Mifflin, 1972.

[119] J. Nunamaker, D. Vogel, and B. Konsynski. Interaction of task and technology to support large groups. *Decision Support Systems*, 5(2):139–152, 1989.

[120] Haycinth S. Nwana, L. Lee, and Nick R. Jennings. Coordination in software agent systems. *BT Technology Journal*, 14(4):79–88, 1996.

[121] Ana Paiva, G. Andersson, Kristina Höök, Dário Mourão, Marco Costa, and Carlos Martinho. Sentoy in fantasya: designing an affective sympathetic interface to a computer game. *Personal and Ubiquitous Computing Journal*, 6:5–6, 2002.

[122] Ana Paiva, João Dias, Daniel Sobral, Ruth Aylett, Polly Sobreperez, Sarah Woods, Carsten Zoll, and Lynne Hall. Caring for agents and agents that care: Building empathic relations with synthetic agents. In *Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'04*, volume 1, pages 194–201, New York, USA, 2004.

[123] Ana Paiva, Rui Prada, Ricardo Chaves, Marco Vala, Adrian Bullock, G. Andersson, and Kristina Höök. Towards tangibility in gameplay: building a tangible affective interface for a computer game. In *Proceedings of the 5th International Conference on Multimodal Interfaces*, Vancouver, Canada, 2003.

[124] S. Pasquariello and Catherine Pelachaud. Greta: a simple facial animation engine. In *Sixth Online World Conference on Soft Computing in Industrial Appications*, 2001.

[125] A. Pepitone and R. Kleiner. The effects of threat and frustration on group cohesiveness. *Journal of Abnormal and Social Psychology*, (54):192–199, 1957.

[126] Jean Piaget. *The Construction of Reality in the Child.* MIT Press, 1955.

[127] R. Picard. *Affective Computing.* MIT Press, 1997.

[128] Rui Prada, Isabel Machado, and Ana Paiva. Teatrix: virtual environment for story creation. In *Proceedings of the Fifth International Conference on Intelligent Tutoring Systems*, Montreal, Canada, 2000. Springer.

[129] Rui Prada and Ana Paiva. Building believability in groups of synthetic characters. In *Proceedings of the Workshop on Coalition and Teams: Formation and Activity in the 3rd International Joint Conference on Autonomous Agents and Multi Agents Systems (AAMAS'2004)*, New York, USA, 2004.

[130] Rui Prada and Ana Paiva. A believable group in the synthetic mind. In *Proceedings of the Symposium on Mind-Minding Agents (AISB'2005)*, Hatfield, England, 2005.

[131] Rui Prada and Ana Paiva. Believable groups of synthetic characters. In *Proceedings of the 4th International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS'05)*, Utrecht, Holland, 2005. ACM Press.

[132] Rui Prada and Ana Paiva. Intelligent virtual agents in collaborative scenarios. In *Proceedings of the 5th International Working Conference on Intelligent Virtual Agents (IVA2005)*, Kos, Greece, 2005.

[133] Rui Prada and Ana Paiva. Synthetic group dynamics in entertainment scenarios: Creating believable interactions in groups of synthetic characters. In *ACM-SIGCHI digital library in the 2nd International Conference on Advances in Computer Entertainment Technology (ACE'05)*, Valencia, Spain, 2005. ACM-SIGCHI.

[134] Rui Prada, Ana Paiva, Isabel Machado, and Catarina Gouveia. You cannot use my broom! i'm the witch, you're the prince: collaboration in a virtual dramatic game. In *Proceedings of the Sixth International Conference on Intelligent Tutoring Systems*, Biarritz, France, 2002. Springer.

[135] Rui Prada, Marco Vala, Ana Paiva, Kristina Höök, and Adrian Bullock. Fantasya - the duel of emotions. In *Proceedings of the 4th International Working Conference on Intelligent Virtual Agents (IVA2003)*, Irsee, Germany, 2003.

[136] Presence-Research. The presence-research portal. *(online) http://www.presence-research.org*, last access on Jan 2005.

[137] Anand S. Rao and Michael P. Georgeff. Modeling rational agents within a {BDI}-architecture. In James Allen, Richard Fikes, and Erik Sandewall, editors, *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning*, pages 473–484. Morgan Kaufmann publishers Inc.: San Mateo, CA, USA, 1991.

[138] Anand S. Rao and Michael P. Georgeff. Bdi-agents: from theory to practice. In *First International Conference on MultiAgent Systems (ICMAS-95)*, pages 312–319, San Francisco, California - USA, 1995.

[139] Bertram H. Raven. *Current Studies in Social Psychology*, chapter Social Influence and Power. Holt, Rinehart and Winston, New York, 1965.

[140] Bryon Reeves and Clifford Nass. *The Media Equation. How People Treat Computers, Television, and New Media Like Real People and Places.* Cambridge University Press, 1998.

[141] C. W. Reynolds. Flocks, herds, and schools: a distributed behavioural model. In *Computer Graphics (SIGGRAPH '87 Conference Proceedings)*, pages 25–34, 1987.

[142] Jeff Rickel and W. Lewis Johnson. Integrating pedagogical capabilities in a virtual environment agent. In *Proceedings of the First International Conference on Autonomous Agents*. ACM Press, 1997.

[143] Jeff Rickel and W. Lewis Johnson. Intelligent tutoring in virtual reality: a preliminary report. In *Proceedings of the the Eighth World Conference on Artificial Intelligence in Education*, pages 294–301. IOS Press, 1997.

[144] Jeff Rickel and W. Lewis Johnson. Virtual humans for team training in virtual reality. In *Proceedings of the Ninth International Conference on AI in Education*, pages 578–585, 1999.

[145] Jeff Rickel, W. Lewis Johnson, and Marcus Thiebaux. Steve: a virtual human for individual and team training. In *Achieving Human-Like Behaviour in Interactive Animated Agents Workshop on The Fourth International Conference on Autonomous Agents*, Barcelona, Spain, 2000.

[146] J. S. Rosenschein and G. Zoltkin. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers.* MIT Press, Cambridge, MA - USA, 1994.

[147] Stuart Russel and Peter Norvig. *Artificial Intelligence: a Modern Approach.* Prentice Hall, 1995.

[148] Markus Schmitt and Thomas Rist. Avatar arena: virtual group-dynamics in multi-character negotiation scenarios. In *4th International Workshop on Intelligent Virtual Agents*, page 358, 2003.

[149] T. Schubert, F. Friedmann, and H. Regenbrecht. The experience of presence: factor analytic insights. *Presence: teleoperators and virtual environments*, 10:266–281, 2001.

[150] J. R. Searle. *Speech Acts: an Essay in the Philosophy of Language*. Cambridge University Press, Cambridge, 1969.

[151] P. Sengers, S. Perry, and J. Smith. Semi-autonomous avatars. 2000.

[152] Peter Pin shan Chen. The entity-relationship model-toward a unified view of data. *ACM Transactions on Database Systems*, 1976.

[153] M. E. Shaw. *Group Dynamics: the Psychology of Small Group Behaviour*. McGraw-Hill, New York, 1981.

[154] T. B. Sheridan. Musings on telepresence and virtual presence. *Presence: teleoperators and Virtual Environments*, 1992.

[155] Y. Shoham and M. Tennenholtz. Emergent conventions in multi-agent systems. *Knowledge Representation and Reasoning*, 1992.

[156] Y. Shoham and M. Tennenholtz. On the synthesis of useful social laws for artificial agent societies. In *10th National Conference on Artificial Intelligence*, San Diego, CA - USA, 1992.

[157] J. Slaney and S. Thiebaux. Blocks world revisited. *Artificial Intelligence*, (125):119–153, 2001.

[158] S. Smith, T. Marsh, D. Duk, and P. Wright. Drowning in immersion. *UK Virtual Reality Special Interest Group*, 1998.

[159] StatSoft. The statistics homepage. *(online) http://www.statsoft.com/textbook/stathome.html*, last access on Jan 2005.

[160] I. D. Steiner. *Group Process and Productivity*. Academic Press, New York, 1972.

[161] D. Stock and H. A. Thelen. *Emotional Dynamics and Group Culture*. New York University Press, New York, 1958.

[162] R. M. Stogdill. *Individual Behaviour and Group Achievement*. Oxford University Press, London, 1959.

[163] A. D. Stright and L. H. Supplee. Children's self-regulatory behaviors during teacher-directed, seat-work, and small-group instructional contexts. *Journal of Educational Research*, 95(4):235–244, 2002.

[164] Sun. The java3d homepage. *(online) http://java.sun.com/products/java-media/3D/*, 2005.

[165] I. E. Sutherland. The ultimate display. In *IFIP Congress*, 1965.

[166] K. P. Sycara. Argumentation: planning other agents' plans. In *11th International Joint Conference on Artificial Intelligence*, Detroit, MI - USA, 1989.

[167] The GAIPS Team. Intelligent agents and synthetic characters group. *(online) http://gaips.inesc-id.pt*, last access on December 2005.

[168] The VICTEC Team. Virtual ict with empathic characters. *(online) http://www.victec.org*, last access on December 2005.

[169] Nadia Thalmann and Daniel Thalmann. *Artificial Life and Virtual Reality*. John Wiley and Sons, Chichester, UK, 1994.

[170] F. Thomas and O. Johnson. *Disney Animation: the Illusion of Life*. Abbeville Press, New York, 1981.

[171] B. Tomlinson and Bruce Blumberg. Multiple conceptions of character-based interactive installations. 2001.

[172] Badger Tomlinson and Bruce Blumberg. Social synthetic characters. *Computer Graphics*, 26(2), May 2002.

[173] Troika. The temple of elemental evil. *(online) http://www.atari.com/toee/*, 2003.

[174] J. Vince. *Essential Virtual Reality fast*. Springer, 1998.

[175] Douglas N. Walton and Erik C. Krabbe. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. State University Of New York Press, Albany, NY - USA, 1995.

[176] W. F. Whyte. *Social Psychology at the Crossroads*, chapter Small Groups and Large Organizations. Harper, New York, 1951.

[177] Charlotte Wiberg. From ease of use to fun of use : usability evaluation guidelines for testing entertainment web sites. In *Proceedings of Conference on Affective Human Factors Design*, Singapore, 2001.

[178] Michael Wooldridge. *An Introduction to Multi-Agent Systems*. John Wiley and Sons, 2002.

[179] Michael Wooldridge and Nick R. Jennings. Intelligent agents: theory and practice. *The Knowledge Engineering Review*, 10(2):115–152, 1995.

Addendum

This addendum clarifies 2 issues that were not very clear in the main document. The first one, is related to the formulas presented in chapters 5, 6 and 7. These formulas are not expressed in common *first order logic*. We have used an extension of the *first order logic* that handles sets, as defined by Russell and Norvig [147], because we have defined the concept of group as a set of elements. In addition, we have used the special notations, taken from the set theory, that the same authors present for the common set operators (e.g. $\in$, $\cup$ and $\subset$). We also have used, in some formulas, a syntax extension that we have taken from JESS, the language used for the implementation of the model. In formulas 5.16 to 5.25 you can see an expression that follows the pattern *v:(condition)*, for example, *Motivation(A,G) = m2:(m2 > m1)*. This syntax pattern associates a condition to a value and it is equivalent to the conjunction of the expression that uses the value and the condition itself. Thus, the previous example is equivalent to *Motivation(A,G) = m2 ∧ (m2 > m1)*.

Furthermore, we would like to clarify that some of the symbols presented in the formulas are functions. There are some explicit references to the predicates, but the functions are never referred as such. This is the case of: *Members(g), SkillLevel(a,g), SocInfluence(a,b), SocAttraction(a,b), Extraversion(a), Agreeableness(a), Position(a,g), Motivation(a,g), Performers(i), Supporters(i)* and *Relevance(i,g)*.

For this reason, the formula 5.10 is wrong since it presents two problems. First of all, the function *Position(A,G)* was referenced as *GroupPosition(A,G)*. And secondly, the functions *Extraversion(A)* and *Position(A,G)* were treated as predicates. The formula should include the predicate *High(x)* for each function and should, therefore, be defined as follows:

$$\forall G, I, A : Group(G) \wedge Interaction(I) \wedge A \in Members(G),$$
$$High(Extraversion(A)) \wedge High(Position(A,G)) \wedge$$
$$High(Motivation(A,G)) \vdash Starts(A,I,G)$$

In addition, we would like to stress that for each of the concepts defined by the predicates presented in equations 5.1 to 5.9 we have defined a set of auxiliary predicates and functions. These are mentioned in the description of the equations, where they are used, but are never properly defined beforehand. For example, we have defined the concept of group as *Group(identity, members)* but use the auxiliary predicate *Group(identity)* and function

Members(identity) in the succeeding equations. Other relevant examples are the case of *Extraversion(member)* and *Agreeableness(member)* that refer to the personality of a member of the group and the predicates related to the interactions. For example, those that associate the interaction with a given type (*Interaction(i), InstrInteraction(i), SocEmotInteraction(i), Encourage(i), Discourage(i), Positive(i)* and *Negative(i)*) or those that are related to the occurrence of the interaction (*Starts(a,i,g)*, *Target(i, t)*, *Performs(a,i,g)* and *Starts(a,i,g,t)*).

We have also introduced some predicates, related to the discussion of proposals in the group, without a previous reference. They appear in some equations of chapter 7 (7.6 and 7.11). These predicates are: (1) *Proposal(P)* that identifies P as a proposal, (2) *Proposes(A,P,G)* that states that member A is responsible for the proposal in group G and (3) *Agrees(A,P)* that states that member A agrees with the proposal P.

The second issue is related to the implementation of the SGD Model in the synthetic characters used in the test case. As stated, the model, as well as the characters' minds, was implemented in JESS, a rule base system. Thus, the model was not implemented in logic but by facts and rules.

Facts represent the knowledge that the characters build regarding the model, for example, the social relations established between the members. This knowledge follows directly the definitions presented in the several formulas of sections 5.4, 5.5 and 5.6 apart from a change in the syntax. For example, the concept of group is expressed in logic as *Group(identity, members)* (see formula 5.4) and in JESS as *(group (identity _id) (members _members))*. The main difference is that, in JESS, the arguments of the predicates (that correspond to facts) are labelled.

Furthermore, for each of the rules of interaction expressed by the formulas in section 5.6.2 we have defined some production rules in the JESS system. In this case, the syntax between the logic formalization of the rules and their implementation in JESS is very different, but the behaviour of the implementation follows what is expressed in the logic formalization.

The main motivation behind the formalization of the SGD Model in logic, as presented in chapter 5, was to have a clear and unambiguous definition of the model independent from its implementation.