

Norm Negotiation in Online Multi-Player Games

Guido Boella¹

Patrice Caire²

Leendert van der Torre²

¹ Dipartimento di Informatica, Università di Torino, Italy.

² University of Luxembourg - Luxembourg.

Abstract. In this paper we introduce an agent communication protocol and speech acts for norm negotiation. The protocol creates individual or contractual obligations to fulfil goals of the agents based on the so-called social delegation cycle. First, agents communicate their individual goals and powers. Second, they propose social goals which can be accepted or rejected by other agents. Third, they propose obligations and sanctions to achieve the social goal, which can again be accepted or rejected. Finally, the agents accept the new norm by indicating which of their communicated individual goals the norm achieves. The semantics of the speech acts is based on a commitment to public mental attitudes. The norm negotiation model is illustrated by an example of norm negotiation in multi-player online gaming.

Keywords: Multi-player online games, normative multi-agent systems, dependence networks, social delegation cycle, agent communication protocols.

1. Introduction

Online multi-player games are online spaces in which human and artificial agents interact with each other every day [4]. For example, EverQuest has more than four hundred thousand subscribers and has developed into a complex system of social dynamics, like The Sims Online or Second Life, something more than just a game. Massively multi-player online games have now become among the most complex and sophisticated online social spaces in existence [34]. One aspect in which online games move towards

Received xxx

Revised xxx

Accepted xxx

social spaces is that players become more independent and autonomous with respect to the game designers, in the sense that they start to play in ways unforeseen by the game designers.

“Why, you might ask, would anyone waste four hours of their life doing this? Because a game said it couldn’t be done.

This is like the Quake freaks that fire their rocket launchers at their own feet to propel themselves up so they can jump straight to the exit and skip 90% of the level and finish in 2 seconds. Someone probably told them they couldn’t finish in less than a minute.

Games are about challenges, about hurdles or puzzles or fights overcome. To some players, the biggest hurdle or challenge is how to do what you (the designer) said couldn’t happen. If you are making a game, accept this.” [34]

An important aspect of autonomy, both for individuals and groups, is the creation of one’s own norms, as reflected by the literal meaning of autonomy: auto-nomos. In online games, players can create their own norms to coordinate their activities, as we illustrate by a running example in this paper. Moreover, autonomy is a central concept in agent theory, and therefore multi-agent systems technology can be used to develop online games. However, existing social mechanisms for norm creation and change in multi-agent systems are restricted in the sense that either norms are not explicitly represented, or there are no communication primitives, protocols and social mechanisms to create or change norms. Normative multi-agent systems [21, 13, 8, 9] provide agents with abilities to automatically devise organizations and societies coordinating their behavior via obligations, norms and social laws. A distinguishing feature from group planning is that also sanctions and control systems for the individual or contractual obligations can be created. Since agents may have conflicting goals with respect to the norms that emerge, they can negotiate amongst each other which norm will be created. In [10] we assume that the social goals emerge, and in [11] we propose a two step negotiation process:

- Propose the social goal and argue about that (instead of merging), and
- Negotiate the norm and sanction.

The research question of this paper is how to define a communication protocol for online multiplayer games and social spaces based on the social delegation cycle [11, 10] explaining the negotiation of new social norms from cognitive agent goals in three steps. First individual agents or their representatives negotiate social goals, then a social goal is negotiated in a social norm, and finally the social norm is accepted by an agent [20] when it recognizes it as a norm, the norm contributes to the goals of the agent, and it is obeyed by the other agents. A model of norm negotiation explains also what it means, for example, to recognize or to obey a norm, and how new norms interact with existing ones. This breaks down in the following three sub-questions:

1. How to use agent technology for the social delegation cycle referring to the individual goals of the agents, which are not accessible for other agents and which are not represented in the agent environment? We propose an additional step before the social delegation cycle, in which the agents communicate their individual goals and powers. In our communication model, the social delegation cycle does not refer to the individual private goals of the agents, but to the communicated and therefore public goals.
2. How do agents publicly accept a norm in the social delegation cycle based on public mental attitudes? We assume that each agent has to commit itself to the social norm by indicating which of its publicly communicated goals the norm achieves.

3. How to use the communication protocol in multi-player online games like EverQuest and social spaces like The Sims Online or Second Life, where the kind of norms that must be used cannot be anticipated beforehand, norms are created by human players and users, and we cannot expect that users are able to specify logical formulas as used in most agent communication languages. We use an abstract power and dependence based representation of the social delegation cycle.

The semantics of the speech acts is based on public mental attitudes. To keep the formal details of this paper within limits, we do not detail the operational semantics of the speech acts, but sketch informally their effects and illustrate them by a running example. Formalizations of speech acts with public mental attitudes can be found in [26, 39, 25].

An important subtask of the social delegation cycle is to break down goals in sub-goals, using a goal hierarchy. However, to focus on the main social mechanism, in this paper we do not detail the sub-protocol for subgoaling. Therefore, we assume that the individual goals communicated by the agents are such that they can be seen to by other agents. The extension with a protocol for subgoaling is left for further research.

The layout of this paper is as follows. In Section 2 we discuss norm creation in online games and social spaces, and we introduce our running example. In Section 3 we define our social-cognitive conceptual model of multiagent systems and in Section 4 we discuss the kind of agent communication language in which we study and formalize the social delegation cycle. In Section 5 we formalize individual goal and power communication. In Section 6 we define the negotiation protocol, which we use in Section 7 to formalize goal negotiation, and in Section 8 we formalize norm negotiation. Finally, in Section 9 we formalize the acceptance relation. In Section 10 we discuss related research of an abstract model of the social delegation cycle in Tennenholtz' game-theoretic artificial social systems, our own more detailed social-cognitive model, and we discuss why these models are less suitable for communication models in online games and social spaces than our power and dependence based model.

2. Multi player online games

As a running example we use in this paper an example discussed by Peter Ludlow [34] from Sony's EverQuest.

2.1. EverQuest

EverQuest is a multiplayer online game where gamers are supposed to fight each other in a world of snakes, dragons, gods and the Sleeper. Sony intended the Sleeper to be unkillable and gave it an extreme high hit points. However, a combined combat of close to 200 players nearly succeeded to kill the animal. Unfortunately, Sony decided to intervene and rescue the monster. Most of the discussion on this example has highlighted the decrease in trust of the game players in Sony, despite the fact that the next day Sony let the game players beat the Sleeper. However, in this paper we would like to highlight what this story tells us about the desires of game players, and its consequences for necessary technology in games. The following quote illustrates the excitement in killing the Sleeper.

A supposedly [player-vs.-player] server banded together 200 people. The chat channels across the server were ablaze, as no less than 5,000 of us listened in, with OMG theyre

attempting the Sleeper! Good luck d00dz! Everyone clustered near their screens, sharing the thrill of the fight, the nobility of the attempt and the courage of those brave 200. Play slowed to a crawl on every server as whispers turned to shouts, as naysayers predicted, It cant be done or It will drop a rusty level 1 sword and most of us just held our breath, silently urging them forward. Rumors abounded: If they win, the whole EQ world stops and you get the text from the end of Wizardry 1, or If they win, the president of Sony will log on and congratulate them. With thousands watching and waiting, the Sleepers health inched ever downward.

...

[EverQuest player] Ghenwivar writes, On Monday, November 17th, in the most amazing and exciting battle ever, [EverQuest guilds] Ascending Dawn, Wudan and Magus Imperialis Magicus defeated Kerafyrm, also known as The Sleeper, for the first time ever on an EverQuest server. The fight lasted approximately three hours and about 170180 players from [EverQuest server] Rallos Zeks top three guilds were involved. Hats off to everyone who made this possible and put aside their differences in order to accomplish the impossible. Congratulations RZ!!!" [34]

The example illustrates that the game had been so well wrought that a real community of players had formed, one that was able to set aside its differences, at least for a night, in pursuit of a common goal. This was not intended or foreseen by Sony, and getting two hundred people to focus on accomplishing the same task is a challenge. In this paper we study multiagent technology to support these emerging cooperation in online games.

2.2. Multiagent systems

Computer games have traditionally implemented empirical solutions to many AI problems and are now turning to more traditional AI algorithms. Cavazza [18] introduces the role of AI in gameplay, reviews the main techniques used in current computer games such as Finite-State Transition Networks, rule-based systems and search algorithms, describes the implementation of AI in several commercial computer games, as well as academic research in AI targeting computer games applications, and discusses future trends and proposing research directions.

"Another important challenge for AI in computer games is the development of on-line and multi-player gaming. In theory, the availability of human opponents around the clock would seal the fate of AI in on-line games. However, this is unlikely to be the case, as for instance autonomous actors would still be needed for the most mundane tasks in large-scale simulations and on-line games, as human players are unwilling to play such parts.

The recent advances in graphic rendering will leave more room for the future development of AI in games. With the level of realism currently achieved, further improvements in realism will come from physics and AI, The development of AI should also foster the development of new game genres based on more sophisticated user interaction with artificial actors." [18]

Laird and van Lent [30] propose that artificial intelligence for interactive computer games is an emerging application area in which this goal of human-level AI can successfully be pursued. Interactive computer games have increasingly complex and realistic worlds and increasingly complex and intelligent computer-controlled characters. They motivate their proposal of using interactive computer games, review previous research

on AI and games, present the different game genres and the roles that human level AI could play within these genres, describe the research issues and AI techniques that are relevant to each of these roles. Their conclusion is that interactive computer games provide a rich environment for incremental research on human-level AI.

“From a researcher’s perspective, even if you are not interested in human-level AI, computer games offer interesting and challenging environments for many, more isolated, research problems in AI. We are most interested in human-level AI, and wish to leverage computer games to rally support for research in human-level AI. One attractive aspect of working in computer games is that there is no need to attempt a Manhattan Project approach with a monolithic project that attempts to create human-level intelligence all at once. Computer games provide an environment for continual, steady advancement and a series of increasingly difficult challenges. Just as computers have inexorably gotten faster, computer game environments are becoming more and more realistic worlds, requiring more and more complex behavior from their characters. Now is the time for AI researchers to jump in and ride the wave of computer games.” [30]

2.3. Agent communication protocol

We use the EverQuest example in this paper to illustrate our communication model. It consists of the following four steps.

1. Agents communicate their individual goals and powers. They say which monster they would like to kill, how they want it to be killed, and which tasks they are able to see to given their powers.
2. They propose social goals which can be accepted or rejected by other agents, based on the communicated goals of the agents. The negotiation can be done by a subset of all players involved.
3. They propose obligations and sanctions to achieve the social goal, which can again be accepted or rejected. Again, the negotiation can be done by a subset of the agents involved in the coordinated attack.
4. Finally, the agents who like to participate in the coordinated attack accept the new norm by indicating which of their communicated individual goals the norm achieves.

3. Power viewpoint on normative multiagent systems

In this paper we follow the definition of power as the ability of agents to achieve goals. Thus, an agent is more powerful than another agent if it can achieve more goals.

For example, in the so-called power view on multi-agent systems [5], a multi-agent system consists of a set of agents (A), a set of goals (G), a function that associates with each agent the goals the agent desires to achieve (*goals*), and a function that associates with each agent the sets of goals it can achieve (*power*). To be precise, since goals can be conflicting in the sense that achieving some goals may make it impossible to achieve other goals, the function *goals* returns a set of set of goals for each set of agents. Such abstract structures have been studied as qualitative games by Wooldridge and Dunne [53], though they do not call the ability of agents to achieve goals their power. To model trade-offs among goals of agents, we introduce a priority relation among goals.

Definition 1. Let a multiagent system be a tuple $\langle A, G, goals, power, \geq \rangle$ where:

- the set of agents A and the set of goals G are two finite disjoint sets;
- $goals : A \rightarrow 2^G$ is a function that associates with each agent the goals the agent desires to achieve;
- $power : 2^A \rightarrow 2^{2^G}$ is a function that associates with each set of agents the sets of goals the set of agents can achieve;
- $\succeq : A \rightarrow \subseteq 2^G \times 2^G$ is a function that associates with each agent a partial pre-ordering on the sets of his goals;

The following example illustrates the notion of a multiagent system for the running example.

Example 1. The set of agents are the players that participate in the coordinated attack. We assume that the individual agents have various ways to attack the Sleeper, and they disagree about which is the best way to kill the beast. The goals of the agents are different individual ways to attack the animal. The powers of the agents are the goals each agent can achieve. The priorities of the agents are their preferences for the ways to attack the animal.

To model the role of power in norm negotiation, we extend the basic power view in a couple of ways. To model obligations we introduce a set of norms, we associate with each norm the set of agents that has to fulfill it, and for each norm we represent how to fulfill it, and what happens when it is not fulfilled. In particular, we relate norms to goals in the following two ways.

- First, we associate with each norm n a set of goals $O(n) \subseteq G$. Achieving these normative goals $O(n)$ means that the norm n has been fulfilled; not achieving these goals means that the norm is violated. We assume that every normative goal can be achieved by the group, i.e., that the group has the power to achieve it.
- Second, we associate with each norm a set of goals $V(n)$ which will not be achieved if the norm is violated (i.e., when its goals are not achieved), this is the sanction associated with the norm. We assume that the group of agents does not have the power to achieve these goals.

Since we accept norms without sanctions, we do not assume that the sanction affects at least one goal of each agent of the group the obligation belongs to.

Definition 2. Let a normative multi-agent system be a tuple $\langle MAS, N, O, V \rangle$ extending a multiagent system $MAS = \langle A, G, goals, power, \succeq \rangle$ where:

- the set of norms N is a finite set disjoint from A and G ;
- $O : N \times A \rightarrow 2^G$ is a function that associates with each norm and agent the goals the agent must achieve to fulfill the norm; We assume for all $n \in N$ and $a \in A$ that $O(n, a) \in power(\{a\})$;
- $V : N \times A \rightarrow 2^G$ is a function that associates with each norm and agent the goals that will not be achieved if the norm is violated by agent a ; We assume for each $B \subseteq A$ and $H \in power(B)$ that $(\cup_{a \in A} V(n, a)) \cap H = \emptyset$.

The normative multiagent system is illustrated by our running example.

Example 2. The obligations of each agent are the attacks each individual player must make. Sanctions can be that someone is removed from the coalition, has to pay a penalty (if this is possible in the game), gets a lower reputation, and so on.

An alternative way to represent normative multiagent systems replaces the function

power by a function representing dependencies between agents. For example, a function of minimal dependence can be defined as follows. Agent a depends on agent set $B \subseteq A$ regarding the goal g if $g \in \text{goals}(a)$, $g \notin \text{power}(\{a\})$, $g \in \text{power}(B)$, and there is no $C \subset B$ such that $g \in \text{power}(C)$. Note that dependence defined in this way is more abstract than power, in the sense that we have defined dependence in terms of power, but we cannot define power in terms of dependence.

4. Agent communication languages

Agent communication languages share the idea that agents communicate by performing speech acts [3, 40], and that these actions can be modeled as some types of planning operators with preconditions and effects [19, 2]. Longer stretches of dialogue can be explained by plans which structure attempts of participants to reach their joint goals, e.g., [32]. However, many variations of agent communication languages have been designed on these common speech act foundations, thus making the standardization effort difficult.

There is a distinction in terms of communication primitives between two main traditions in agent communication. The agent's *mental attitudes* like beliefs and intentions are included in preconditions and effects of speech acts used in information seeking dialogues using speech acts such as inform and request. This side has been popularized by the standardization efforts of the Foundation for Intelligent Physical Agents [23], though the concepts of *desire*, *goal* and *intention* are not used uniformly, and *desire* and *goal* are often used interchangeably. Below the agent social semantics tradition, developed as a reaction to FIPA, is depicted [45]. The agent can either be a creditor or a debtor of his social commitments. The semantics of speech acts used in negotiation or persuasion refer to these commitments. Though both kinds of dialogue use the terminology of commitment, requesting an action in negotiation means something else than defending a proposition in persuasion after a challenge. We therefore distinguish action commitment from propositional commitment, where the former is used to give a semantics of speech acts like *request* or *propose* in negotiation [45, 24], and the latter is used for speech acts like *assert* or *challenge* in persuasion [28, 52].

In this paper we use an approach bridging the two traditions based on a reinterpretation of the beliefs and intentions. The idea that meaning cannot be private, but is inter-subjective or based on a common ground, has been accepted for a long time in the philosophy of language. This idea of public meaning has been discussed, for example, by [31] or [46], who stress the importance of the common ground. The crucial point is that a private semantics does not make it possible for a language user to be objectively wrong about an interpretation from a third person point of view. In contrast to FIPA, beliefs and intentions are not interpreted as private mental attitudes, but as some kind of public mental attitudes, for example as grounded beliefs [26] or as ostensible beliefs or public opinions [39].

5. Individual goal and power communication

The first phase consists of communicating the individual goals and powers of the agents. Since existing agent communication languages do not cover inform speech acts whose content is a public goal, we introduce the speech act *desire*. The semantics of the speech act

$$\text{desire}(a, \phi)$$

is that agent a has the public goal ϕ . It means that he is committed to this public goal, and, for example, he cannot say later that he desires the opposite.

If hundreds of agents have to communicate their goals, then the communication will become long and cumbersome. So, when an agent informs the others that he desires ϕ , the other agents can respond by a simple ‘Me too!’. Moreover, we do not have only speech acts implying that the speaking agent has a goal, but also speech acts implying that all agents have the same goal. Agent a says that everyone has the desire for ϕ :

desireall(a, ϕ)

The agents can respond to this speech act by saying ‘Not me!’. If none of them answers in this way, then the *desireall* act is equivalent to agent a stating his goal for ϕ , and all the other agents saying ‘Me too!’. The principle that by default all agents agree, is sometimes called ‘silence means consent’.

An optional parameter of these two speech acts is the priority of the individual goals. In this running example we assume that there are five priorities: essential, important, intermediate, useful, possible.

Example 3. Suppose that three agents $\{a, b, c\}$ are going to negotiate the attack strategy to kill the Sleeper. Moreover, we assume there are various ways in which the individual players can attack the animal: hitting the feet, throwing stones at the animal, shooting between the eyes, and so on. The three agents are negotiating for three kinds of game players with specific powers. For example, agent a is negotiating for all agents that have the power to hit the feet, throw stones, and have guns. In this example, each negotiator communicates goals they cannot see to by themselves. For example, agent a will communicate the desire to scare the Sleeper to death.

The negotiators begin by communicating their individual goals. Each player can type individual or shared goals by entering a text string into his computer. In the text box is also a list of goal priorities, which is default at intermediate. When someone types an individual goal, other agents can press the ‘Me too!’ button, When an agent enters a shared goal, other agents can press the ‘Not me!’ button, or change the suggested priority of the goal. For example, when agent a enters his goal and agent b accepts it too, and agent c communicates a shared goal but agent d refuses it, then the system translates it into the following speech acts.

desire(b, ‘hitting the feet’)
desire(b, ‘hit him with stones’, intermediate)
desire me too(c, ‘hit him with stones’, important)
desire(c, ‘shoot him in the eyes’, possible)
desire(a, ‘show him a picture of M.’)
desireall(a, ‘scare him to death’)
desire not me(c, ‘scare him to death’)
...

Each agent can indicate which goals it can see to. We use the new speech act *power(a, ϕ)*, indicating that agent a communicates that he has the power to see to goal ϕ . As before, the other agents can respond by ‘Me too!’ and there is a speech *powerall* which agents can respond to by saying ‘Not me!’.

Example 4. Consider the following three speech acts.

power(a, ‘hitting the feet’)
powerall(a, ‘hit him with stones’)
power not me(c, ‘hit him with stones’)
powerall(c, ‘shoot him in the eyes’)
...

# goal	speaker	agents and priorities	goal description	powers
1	<i>b</i>	<i>b, c</i>	hitting the feet	<i>a</i>
2	<i>b</i>	<i>b:intermediate, c:important</i>	hit him with stones	<i>a, b</i>
3	<i>c</i>	<i>c:possible</i>	shoot him in the eyes	<i>a, b, c</i>
4	<i>a</i>	<i>a</i>	show him a picture of M.	<i>b, c</i>
5	<i>a</i>	<i>a, b</i>	scare him to death	<i>c</i>
...

Table 1. Communicated goals and powers

The communicated goals and powers are represented in a single table for each player, as shown in Table 1. In case the agents are negotiating a norm for a large set of players, then not all individual names of the agents are given, but only aggregate information such as the number of players having the power to see to a goal.

6. Generic negotiation protocol

A negotiation protocol is described by a set of sequences of negotiation actions which either lead to success or failure. In this paper we only consider protocols in which the agents propose a so-called deal, and when an agent has made such a proposal, then the other agents can either accept or reject it (following an order \succ of the agents). Moreover, they can also end the negotiation process without any result.

Definition 3 (Negotiation Protocol). A negotiation protocol is a tuple $\langle Ag, deals, actions, valid, finished, broken, \succ \rangle$, where:

- the agents *Ag*, *deals* and *actions* are three disjoint sets, such that $actions = \{propose(a, d), accept(a, d), reject(a, d) \mid a \in Ag, d \in deals\} \cup \{breakit(a) \mid a \in Ag\}$.
- valid*, *finished*, *broken* are sets of finite sequences of *actions*.

We now instantiate this generic protocol for negotiations in normative multiagent systems. We assume that a sequence of actions (a history) is valid when each agent does an action respecting the order defined on agents. Then, after each proposal, the other agents have to accept or reject this proposal, again respecting the order, until they all accept it or one of them rejects it. When it is an agent's turn to make a proposal, it can also end the negotiation by breaking it. The history is *finished* when all agents have accepted the last deal, and *broken* when the last agent has ended the negotiations.

Definition 4 (NMA protocol). Given a normative multiagent system $\langle MAS, N, O, V \rangle$ extending a multiagent system $MAS = \langle A, G, goals, power, \geq \rangle$, a negotiation protocol for *NMAS* is a tuple $NP = \langle A, deals, actions, valid, finished, broken, \succ \rangle$, where:

- $\succ \subseteq A \times A$ is a total order on *A*,
- a history *h* is a sequence of actions, and *valid(h)* holds if:
 - the *propose()* and *breakit()* actions in the sequence respect \succ ,
 - each *propose()* is followed by a sequence of *accept()* or *reject()* actions respecting \succ until either all agents have accepted the deal or one agent has rejected it,
 - there is no double occurrence of a proposal *propose(a, d)* in the same deal by any agent $a \in Ag$, and

- the sequence h ends iff either all agents have accepted the last proposal ($finished(h)$) or the last agent has broken the negotiation ($broken(h)$) instead of making a new proposal.

In theory we can add additional penalties when agents break the negotiation. However, since it is in the interest of all agents to reach an agreement, we do not introduce such sanctions. In this respect norm negotiation differs from negotiation about obligation distribution [7], where it may be the interest of some agents to see to it that no agreement is reached. In such cases, sanctions must be added to the negotiation protocol to motivate the agents to reach an agreement.

Example 5. Assume three agents and the following history.

$action_1 : propose(a, d_1)$

$action_2 : accept(b, d_1)$

$action_3 : reject(c, d_1)$

$action_4 : propose(b, d_2)$

$action_5 : accept(c, d_2)$

$action_6 : accept(a, d_2)$

We have $valid(h)$, because the order of action respects \succeq , and we have $accepted(h)$, because the history ends with acceptance by all agents ($action_5$ and $action_6$) after a proposal ($action_4$).

The open issue of the generic negotiation protocol is the set of deals which can be proposed. They depend on the kind of negotiation. In social goal negotiation the deals represent a social goal, and in norm negotiation the deals contain the obligations of the agents and the associated control system based on sanctions. This is illustrated in our running example in the following sections.

7. Social goal negotiation

We characterize the allowed deals during goal negotiation as a set of goals which contains for each agent a goal it desires. Moreover, we add two restrictions. First, we only allow goals the agents have the power to achieve. Moreover, we have to consider the existing normative system, which may already contain the norms that look after the goals of the agents. We therefore restrict ourselves to new goals. Additional constraints may be added, for example excluding goals an agent can see to itself. However, since such additional restrictions may be unrealistic in some applications (e.g., one may delegate some tasks to a secretary even when one has the power to see to these tasks oneself), we do not consider such additional constraints.

Definition 5 (Deals in goal negotiation). In the goal negotiation protocol, a deal $d \in deals$ is a set of goals satisfying the following restrictions:

1. $d \in power(A)$
2. for all $a \in A$ there exists some $g \in d$ such that
 - (a) $g \in goals(a)$
 - (b) there does not exist a norm n in N such that $g \in \cup_{a \in A} O(n, a)$

The following example illustrates social goal negotiation in the running example.

Example 6. The agents proposing a deal select a couple of individual communicated goals from Table 1 by clicking on them. The system then send them to the other agents,

# goal proposal	speaker	goals	accepted	rejected	reason
1	a	2, 5	a	b	...
2	b	2, 4, 5	a, b, c		
...	

Table 2. Social Goal Negotiation

which can either accept them, or reject them. When they reject a communicated goal, they can add a reason to it, which is communicated to the other agents.

Let $MAS = \langle \{a, b, c\}, \{g_1, g_2, g_3, g_4, g_5\}, goals, power, \geq \rangle$ be a multiagent system representation of Table 1, where:

power: $power(a) = \{\{g_1\}, \{g_2\}, \{g_3\}\}$, $power(b) = \{\{g_2\}, \{g_3\}, \{g_4\}\}$, $power(c) = \{\{g_3\}, \{g_4\}, \{g_5\}\}$, if $G_1 \in power(A)$ and $G_2 \in power(B)$ then $G_1 \cup G_2 \in power(A \cup B)$. Agent a has the power to achieve goals g_1, g_2, g_3 , agent b has the power to achieve goals g_2, g_3, g_4 , and agent c can achieve goals g_3, g_4, g_5 . There are no conflicts among goals.

goals: $goals(a) = \{g_4, g_5\}$, $goals(b) = \{g_1, g_5\}$, $goals(c) = \{g_1, g_2\}$. Each agent desires the tasks it cannot perform itself.

Moreover, let $NMAS = \langle MAS, N, O, V \rangle$ be a normative multiagent system with $N = \{n\}$, $O(n, a_1) = \{g_1\}$. Since there has to be some benefit for agent b and c , the goals g_5 and g_2 have to be part of the social goal. Therefore, social goals (i.e., possible deals) are $\{g_2, g_5\}$ and $\{g_2, g_4, g_5\}$.

Finally, consider the negotiation. Assuming agent a is first in the order \succ , he may propose $\{g_2, g_5\}$. The other agents may accept this, or reject it and agent b will then propose $\{g_2, g_4, g_5\}$. The latter would be accepted by all agents, as they know that according to the protocol no other proposals can be made.

The example illustrates that the negotiation does not determine the outcome, in the sense that there are multiple outcomes possible. Additional constraints may be added to the negotiation strategy to further delimit the set of possible outcomes.

8. Social norm negotiation

We formalize the allowed deals during norm negotiation as obligations for each agent to see to some goals, such that all goals of the social goal are included. Again, to determine whether the obligations imply the social goal, we have to take the existing normative system into account. We assume that the normative system only creates obligations that can be fulfilled together with the already existing obligations.

Definition 6 (Fulfillable NMAS). A normative multiagent system $\langle MAS, N, O, V \rangle$ extending a multiagent system $MAS = \langle A, G, goals, power, \geq \rangle$ can be fulfilled if there exists a $G' \in power(A)$ such that all obligations are fulfilled $\cup_{n \in N, a \in A} O(n, a) \subseteq G'$.

Creating a norm entails adding obligations and violations for the norm.

Definition 7 (Add norm). Let $NMAS$ be a normative multiagent system $\langle MAS, N, O, V \rangle$ extending a multiagent system $MAS = \langle A, G, goals, power, \geq \rangle$. Adding a norm $n \notin N$ with a pair of functions $\langle o_1, o_2 \rangle$ for obligation $o_1 : A \rightarrow 2^G$ and for sanction $o_2 : A \rightarrow 2^G$ leads to the new normative multiagent system $\langle MAS, N \cup \{n\}, O \cup o_1(n), V \cup o_2(n) \rangle$.

Moreover, if every agent fulfills its obligation, then the social goal is achieved.

# goal	speaker	agents and priorities	goal description	powers	obligation
2	b	b :intermediate, c :important	hit him with stones	a, b	a
4	a	a	show him a picture of M.	b, c	b
5	a	a, b	scare him to death	c	c

Table 3. Negotiated norm

Definition 8 (Deals in norm negotiation). In the norm negotiation protocol, a deal $d \in \text{deals}$ for social goal S is a pair of functions $\langle o_1, o_2 \rangle$ for obligation $o_1 : A \rightarrow 2^G$ and for sanction $o_2 : A \rightarrow 2^G$ satisfying the following conditions:

1. Adding $\langle o_1, o_2 \rangle$ to $NMAS$ for a fresh variable n (i.e., not occurring in N) leads again to a normative multiagent system $NMAS'$;
2. $NMAS'$ achieves the social goal, $\cup_{a \in A} o_1(a) = S$.
3. If $NMAS$ is fulfillable, then $NMAS'$ is too.

The running example illustrates the norm negotiation protocol.

Example 7. For the social goal $\{g_2, g_5\}$, there will always be one of the agents who does not have to see to any goal. Consider norm for the social goal $\{g_2, g_4, g_5\}$ in Table 3. A possible solution here is that each agent sees to one of the goals.

Sanctions can be added in the obvious way. In the norm negotiation as defined thus far, the need for sanctions has not been formalized yet. For this need, we have to consider the acceptance of norms.

9. Norm acceptance

An agent accepts a norm when it believes that the other agents will fulfill their obligations, and the obligation implies the goals the cycle started with. For the former we use the quasi-stability of the norm (e.g., if the norm is a convention, then we require that the norm is a Nash equilibrium). Each agent b fulfills the norm *given that all other agents fulfill the norm*. Again we have to take the existing normative system into account, so we add the condition that all other norms are fulfilled. In general, it may mean that an agent does something which it does not like to do, but it fears the sanction more than this dislike. We use the trade-off among goals \geq .

Definition 9 (Stability). A choice c of agent $b \in A$ in $NMAS$ with new norm n is $c \in \text{power}(b)$ such that $\cup_{m \in N \setminus \{n\}} O(m, b) \subseteq d$. The choices of the other agents are $oc = \cup_{a \in A \setminus \{b\}, m \in N} O(m, a) \cup V(n, a)$. The effect of choice c is $c \cup oc \cup V(n, a)$ if $O(n) \subseteq c$, $c \cup oc$ otherwise. $NMAS$ is stable if $\forall b \in A$, there is a choice c such that $O(n, b) \subseteq c$, and there is no choice $c' \geq (b)c$ with $O(n, b) \not\subseteq c'$.

Finally, we have to test whether the new situation is better than the old one for all agents. For example, we may request that the outcome in both the original multiagent system as in the new multiagent system is a Nash equilibrium, and we demand that each Nash outcome in the new system is better than each Nash outcome in the original normative multiagent system. The formalization of these concepts is along the same lines as the definition of acceptance in Definition 9.

We ask that each agent communicates to the other agents why it accepts the norm. By making a public statement, it thus commits to the norm and gives an explanation of its acceptance. Agent a accepts the norm n for goal g :

# goal	speaker	agents and priorities	goal description	powers	obligation	norm acceptance
2	<i>b</i>	<i>b</i> :intermediate, <i>c</i> :important	hit him with stones	<i>a, b</i>	<i>a</i>	<i>c</i>
4	<i>a</i>	<i>a</i>	show him a picture of M.	<i>b, c</i>	<i>b</i>	<i>a</i>
5	<i>a</i>	<i>a, b</i>	scare him to death	<i>c</i>	<i>c</i>	<i>b</i>

Table 4. Norm Acceptance

$accept\text{-}norm(a,n,g)$

Norm acceptance is illustrated in our running example.

Example 8. Let n stand for the norm accepted by the agents. Assume that every agent can select the goal for which he accepts the norm from the list of all goals of the social goal. This can be communicated to the other agents using the following three speech acts.

$accept\text{-}norm(a,n,4)$

$accept\text{-}norm(b,n,5)$

$accept\text{-}norm(c,n,2)$

The resulting table is given in Table 4.

10. Related work

In this section we discuss the existing three formalizations of the social delegation cycle.

10.1. Norms are a class of constraints deserving special analysis

Meyer and Wieringa define normative systems as “systems in the behavior of which norms play a role and which need normative concepts in order to be described or specified” [38, preface]. Alchourrón and Bulygin [1] define a normative system inspired by Tarskian deductive systems:

“When a deductive correlation is such that the first sentence of the ordered pair is a case and the second is a solution, it will be called normative. If among the deductive correlations of the set α there is at least one normative correlation, we shall say that the set α has normative consequences. A system of sentences which has some normative consequences will be called a normative system.” [1, p.55].

Jones and Carmo [29] introduce agents in the definition of a normative system by defining it as “sets of agents whose interactions are norm-governed; the norms prescribe how the agents ideally should and should not behave. [...] Importantly, the norms allow for the possibility that actual behavior may at times deviate from the ideal, i.e., that violations of obligations, or of agents’ rights, may occur.” Since the agents’ control over the norms is not explicit here, we use the following definition.

A normative multi-agent system is a multi-agent system together with normative systems in which agents can decide whether to follow the explicitly represented norms or not, and the normative systems specify how and in which extent the agents can modify the norms. [13]

Note that this definition makes no presumptions about the internal architecture of an agent or of the way norms find their expression in agent’s behavior.

10.1.1. Representation of norms

Since norms are explicitly represented, according to our definition of a normative multi-agent system, the question should be raised how norms are represented. Norms can be interpreted as a special kind of constraint, and represented depending on the domain in which they occur. However, the representation of norms by domain dependent constraints runs into the question what happens when norms are violated. Not all agents behave according to the norm, and the system has to deal with it. In other words, norms are not hard constraints, but soft constraints. For example, the system may sanction violations or reward good behavior. Thus, the normative system has to monitor the behavior of agents and enforce the sanctions. Also, when norms are represented as domain dependent constraints, the question will be raised how to represent permissive norms, and how they relate to obligations. Whereas obligations and prohibitions can be represented as constraints, this does not seem to hold for permissions. For example, how to represent the permission to access a resource under an access control system? Finally, when norms are represented as domain dependent constraints, the question can be raised how norms evolve.

We therefore believe that norms should be represented as a domain independent theory. For example, deontic logic [51, 49, 48, 35, 36, 37] studies logical relations among obligations and permissions, and more in particular violations and contrary-to-duty obligations, permissions and their relation to obligations, and the dynamics of obligations over time. Therefore, insights from deontic logic can be used to represent and reason with norms in multi-agent systems. Deontic logic also offers representations of norms as rules or conditionals. However, there are several aspects of norms which are not covered by constraints nor by deontic logic, such as the relation among the cognitive abilities of agents and the global properties of norms. Meyer and Wieringa explain why normative systems are intimately related with deontic logic.

“Until recently in specifications of systems in computational environments the distinction between normative behavior (as it *should be*) and actual behavior (as it *is*) has been disregarded: mostly it is not possible to specify that some system behavior is non-normative (illegal) but nevertheless possible. Often illegal behavior is just ruled out by specification, although it is very important to be able to specify what should happen if such illegal but possible behaviors occurs! Deontic logic provides a means to do just this by using special modal operators that indicate the status of behavior: that is whether it is legal (normative) or not” [38, preface].

10.1.2. Norms and agents

Conte *et al.* [20] distinguish two distinct sets of problems in normative multi-agent systems research. On the one hand, they claim that legal theory and deontic logic supply a theory of norm-governed interaction of autonomous agents while at the same time lacking a model that integrates the different social and normative concepts of this theory. On the other hand, they claim that three other problems are of interest in multi-agent systems research on norms: how agents can acquire norms, how agents can violate norms, and how an agent can be autonomous. Agent decision making in normative systems and the relation between desires and obligations has been studied in agent architectures [15], which thus explain how norms and obligations influence agent behavior.

An important question in normative multi-agent systems is where norms come from. Norms are not necessarily created by legislators, but they can also be negotiated among agents, or they can emerge spontaneously, making the agents norm autonomous [50].

In electronic commerce research, for example, cognitive foundations of social norms and contracts are studied [8]. Protocols and social mechanisms are now being developed to support such creations of norms in multi-agent systems. Moreover, agents like legislators playing a role in the normative system have to be regulated themselves by procedural norms [12], raising the question how these new kind of norms are related to the other kinds of norms.

When norms are created, the question can be raised how they are enforced. For example, when a contract is violated, the violator may have to pay a penalty. But then there has to be a monitoring and sanctioning system, for example police agents in an electronic institution. Such protocols or roles in a multi-agent system are part of the construction of social reality, and Searle [41] has argued that such social realities are constructed by constitutive norms. This raises the question how to represent such constitutive or counts-as norms, and how they are related to regulative norms like obligations and permissions [9].

10.1.3. Norms and other concepts

Not only the relation between norms and agents must be studied, but also the relation between norms and other social and legal concepts. How do norms structure organizations? How do norms coordinate groups and societies? How about the contract frames in which contracts live? How about the relation between legal courts? Though in some normative multi-agent systems there is only a single normative system, there can also be several of them, raising the question how normative systems interact. For example, in a virtual community of resource providers each provider may have its own normative system, which raises the question how one system can authorize access in another system, or how global policies can be defined to regulate these local policies [9].

10.2. Social delegation cycle using artificial social systems

The problem studied in artificial social systems is the design, emergence or more generally the creation of social laws. Shoham and Tennenholtz [42] introduce social laws in a setting without utilities, and they define *rational* social laws as social laws that improve a social game variable [43]. We follow Tennenholtz' presentation for stable social laws [47]. In Tennenholtz' game-theoretic artificial social systems, the goals or desires of agents are represented by their utilities. A game or multi-agent encounter is a set of agents with for each agent a set of strategies and a utility function defined on each possible combination of strategies. The social goal is represented by a minimal value for the social game variable. Tennenholtz [47] uses as game variable the maximin value. This represents safety level decisions, in the sense that the agent optimizes its worst outcome assuming the other agents may follow any of their possible behaviors.

The social norm is represented by a social law, characterized as a restriction to the strategies available to the agents. It is *useful* with respect to an efficiency parameter e if each agent can choose a strategy that guarantees it a payoff of at least e . A social law is *quasi-stable* if an agent does not profit from violating the law, as long as the other agent conforms to the social law (i.e., selects strategies allowed by the law). Quasi-stable conventions correspond to Nash equilibria. The efficiency parameter can be seen as a social kind of *utility aspiration level*, as studied by Simon [44]. Such aspiration levels have been studied to deal with limited or resource-bounded reasoning, and have led to the development of goals and planning in artificial intelligence; we therefore use a goal based ontology in this paper. The three steps of the social delegation cycle in this classi-

cal game-theoretic setting can be represented as follows. Goal negotiation implies that the efficiency parameter is higher than the utility the agents expect without the norm, for example represented by the Nash equilibria of the game. Norm negotiation implies that the social law is useful (with respect to the efficiency parameter). The acceptance relation implies that the social law is quasi-stable.

The game-theoretical model has several drawbacks as a basis of our communication model. Due to the uniform description of agents in the game-theoretic model, it is less clear how to distinguish among kinds of agents. For example, the unique utility aspiration level does not distinguish the powers of agents to negotiate a better deal for themselves than for the other agents. Moreover, the formalization of the social delegation cycle does neither give a clue on how the efficiency parameter is negotiated, nor on how the social law is negotiated. For example, the goals or desires of the agents as well as other mental attitudes may play a role in this negotiation. There is no sanction or control system in the model (adding a normative system to encode enforceable social laws to the artificial social system complicates the model [6]). Finally, an additional drawback is that the three ingredients of the model (agent goals, social goals, and social laws) are formalized in three completely different ways.

10.3. Power and dependencies

The second formalization of the social delegation cycle is the highly detailed model we have proposed within normative multiagent systems [10]. The challenge to define social mechanisms, as we see it, is to go beyond the classical game theoretic model by introducing social and cognitive concepts and a negotiation model, but doing so in a minimal way. In the model proposed in this paper we therefore keep goals and obligations abstract and we do not describe them by first-order (or propositional) logic or their rule structure, we do not introduce decisions, actions, tasks, and so on. Similar concerns are also mentioned by Wooldridge and Dunne in their qualitative game theory [53].

In [11] we introduce a norm negotiation model based on power and dependence structures. Power may affect norm negotiation in various ways, and we therefore propose to analyze the norm negotiation problem in terms of social concepts like power and dependence. Power has been identified as a central concept for modeling social phenomena in multi-agent systems by various authors [14, 16, 22, 33], as Castelfranchi observes both to enrich agent theory and to develop experimental, conceptual and theoretical new instruments for the social sciences [17].

The model in [11] is based on a distinction between social goal negotiation and the negotiation of the obligations with their control system. Roughly, the social goals are the benefits of the new norm for the agents, and the obligations are the costs of the new norm for the agents in the sense that agents risk being sanctioned. Moreover, in particular, when representatives of the agents negotiate the social goals and norms, the agents still have to accept the negotiated norms. The norm is accepted when the norm is quasi-stable in the sense that agents will act according to the norm, and effective in the sense that fulfilment of the norm leads to achievement of the agents' desires – i.e., when the benefits outweigh the costs.

Our new model is based on a minimal extension of Tennenholtz' game theoretic model of the social delegation cycle. We add a negotiation protocol, sanction and control, and besides acceptance also effectiveness. It is minimal in the sense that, compared to our earlier model [10] in normative multiagent systems, we do not represent the rule structure of norms, we do not use decision variables, and so on. Also, as discussed in

this paper, we do not add goal hierarchy, definitions, etc. The model therefore focusses on various uses of power: the power as ability to achieve goals and in negotiation.

11. Summary

Massively multiplayer online games have now become among the most complex and sophisticated online social spaces in existence [4, 34]. As the lives people lead in virtual worlds become more and more connected with their offline lives such online spaces are shedding significant light on questions of liberty and responsibility, and on how an online existence reverberates in what we think of as a gamers real life [34, 27]. A restricted version of norm creation is built in some of the games, such as for example A Tale in the Desert, but we envision in this paper a gaming environment in which the players can create arbitrary norms to coordinate their behaviors.

In this paper we introduce an agent communication protocol and speech acts for norm negotiation. It creates individual or contractual obligations fulfilling goals of the agents based on the social delegation cycle. The first problem is that the social delegation cycle refers to the individual goals of the agents, which are not accessible to other agents and which are not represented in the agent environment. This problem is solved by the first step of our communication protocol.

1. First, agents communicate their individual goals and powers.
2. Second, they propose social goals which can be accepted or rejected by other agents.
3. Third, they propose obligations and sanctions to achieve the social goal, which can again be accepted or rejected.
4. Finally, the agents accept the new norm by indicating which of their communicated individual goals the norm achieves.

Consequently, the social delegation cycle does not refer to the individual private goals of the agents, but to the communicated public goals.

The second problem, to define norm acceptance, is solved by the final step of our norm negotiation protocol, in which each agent has to commit itself to the social norm by indicating which of its publicly communicated goals the norm achieves.

Finally, the solution of the third problem, that is, how to use the norm negotiation protocol in online multi-player games and social spaces like The Sims Online or Second Life, is to use an abstract power and dependence based representation of the social delegation cycle. Clearly the game players cannot detail their individual probabilities and utilities as required by Tennenholtz' game-theoretic model, and they cannot detail the decision variables of the model in [10]. However, they can detail their goals and powers, as they can use natural language to describe their goals, and detail their powers by selecting the relevant goals. This solution is illustrated by our running example of norm negotiation in multi-player online gaming, where the goals are lines of text entered in a table, powers are cross-tables, social goals are sets of individual goals, and norms are distributions of the goals over the agents.

An important subtask of the social delegation cycle is to break down goals in sub-goals, using a goal hierarchy. However, to focus on the main social mechanism, in this paper we do not detail the sub-protocol for subgoaling. Therefore, we assume that the individual goals communicated by the agents are such that they can be seen to by other agents. The extension with a protocol for subgoaling is left for further research.

References

- [1] C.E. Alchourrón and E. Bulygin. *Normative Systems*. Springer, Wien, 1971.
- [2] J. F. Allen and C. R. Perrault. Analyzing intention in utterances. *Artificial Intelligence*, 15:143–178, 1980.
- [3] J.L. Austin. *How to do things with words*. Harvard University Press, Cambridge Mass., 1962.
- [4] F. Beau. *Culture d'Univers*. FYP editions, 2007.
- [5] G. Boella, L. Sauro, and L. van der Torre. Social viewpoints on multiagent systems. In *Procs. of the 3rd International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS'04)*, pages 1358–1359, 2004.
- [6] G. Boella and L. van der Torre. Enforceable social laws. In *Procs. of the 4th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'05)*, pages 682–689, 2005.
- [7] G. Boella and L. van der Torre. Fair distribution of collective obligations. In *Procs. of the 17th European Conference on Artificial Intelligence (ECAI'06)*, pages 721–722, 2006.
- [8] G. Boella and L. van der Torre. A game theoretic approach to contracts in multiagent systems. *IEEE Transactions on Systems, Man and Cybernetics - Part C*, 36(1):68–79, 2006.
- [9] G. Boella and L. van der Torre. Security policies for sharing knowledge in virtual communities. *IEEE Transactions on Systems, Man and Cybernetics - Part A*, 36(3):439–450, 2006.
- [10] G. Boella and L. van der Torre. Norm negotiation in multiagent systems. *International Journal of cooperative Information Systems (IJCIS)*, 16(1), 2007.
- [11] G. Boella and L. van der Torre. Power in norm negotiation. In *Agent and Multi-Agent Systems: Technologies and Applications, First KES International Symposium, KES-AMSTA 2007*, volume 4496 of LNCS, pages 436–446, Berlin, 2007. Springer.
- [12] G. Boella and L. van der Torre. Substantive and procedural norms in normative multiagent systems. *Journal of Applied Logic*, in press.
- [13] G. Boella, L. van der Torre, and H. Verhagen. Introduction to normative multiagent systems. *Computational and Mathematical Organizational Theory, Special issue on Normative Multiagent Systems*, 12(2-3):71–79, 2006.
- [14] S. Brainov and T. Sandholm. Power, dependence and stability in multi-agent plans. In *Procs. of the 21st National Conference on Artificial Intelligence (AAAI'99)*, pages 11–16, 1999.
- [15] J. Broersen, M. Dastani, J. Hulstijn, and L. van der Torre. Goal generation in the BOID architecture. *Cognitive Science Quarterly*, 2(3-4):428–447, 2002.
- [16] C. Castelfranchi. Modeling social action for AI agents. *Artificial Intelligence*, 103(1-2):157–182, 1998.
- [17] C. Castelfranchi. Micro-macro constitution of power. *ProtoSociology*, 18-19, 2003.
- [18] M. Cavazza. AI in computer games: Survey and perspectives. *Virtual Reality*, 5:223–235, 2000.
- [19] P. R. Cohen and C. R. Perrault. Elements of a plan-based theory of speech acts. *Cognitive Science*, 3:177–212, 1979.
- [20] R. Conte, C. Castelfranchi, and F. Dignum. Autonomous norm-acceptance. In *Intelligent Agents V (ATAL'98)*, LNAI 1555, pages 99–112. Springer, 1999.
- [21] R. Conte, R. Falcone, and G. Sartor. Agents and norms: How to fill the gap? *Artificial Intelligence and Law*, 7(1):1–15, 1999.
- [22] R. Conte and J. Sichman. Multi-agent dependence by dependence graphs. In *Procs. of the 1st International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'02)*, pages 483–490, 2002.
- [23] FIPA. FIPA ACL communicative act library specification. Technical Report SC00037J, Foundation for Intelligent Physical Agents, 2002.
- [24] N. Fornara and M. Colombetti. A commitment-based approach to agent communication. *Applied Artificial Intelligence*, 18(9-10):853–866, 2004.
- [25] J. Hulstijn G. Boella, R. Damiano and L. van der Torre. A common ontology of agent communication languages: Modeling mental attitudes and social commitments using roles. *Applied Ontology*, 2(3-4):217 – 265, 2007.
- [26] B. Gaudou, A. Herzig, and D. Longin. A logical framework for grounding-based dialogue analysis. *Electronic Notes in Theoretical Computer Science*, 157(4):117–137, 2006.
- [27] James Grimmelmann. Virtual borders: The interdependence of real and virtual worlds. *First Monday*, 11(2), 2006.
- [28] C. L. Hamblin. *Fallacies*. Methuen & Co, London, 1970.
- [29] A. Jones and J. Carmo. Deontic logic and contrary-to-duties. In D. Gabbay and F. Guentner, editors, *Handbook of Philosophical Logic*, volume 3, pages 203–279. Kluwer, Dordrecht (NL), 2001.
- [30] J. Laird and M. van Lent. Human-level AI's killer application: Interactive computer games. *AI Magazine*, 22(2):15–25, 2001.
- [31] David Lewis. *Convention: A Philosophical Study*. Harvard University Press, Cambridge, 1969.

- [32] Karen E. Lochbaum. A collaborative planning model of intentional structure. *Computational Linguistics*, 24(4):525–572, 1998.
- [33] F. Lopez y Lopez. *Social Power and Norms: Impact on agent behaviour*. PhD thesis, University of Southampton, 2003.
- [34] P. Ludlow and M. Wallace. *The Second Life Herald*. MIT Press, Cambridge (MA), 2007.
- [35] D. Makinson and L. van der Torre. Input-output logics. *Journal of Philosophical Logic*, 29:383–408, 2000.
- [36] D. Makinson and L. van der Torre. Constraints for input-output logics. *Journal of Philosophical Logic*, 30:155–185, 2001.
- [37] D. Makinson and L. van der Torre. Permissions from an input-output perspective. *Journal of Philosophical Logic*, 32(4):391–416, 2003.
- [38] J. Meyer and R. Wieringa. *Deontic Logic in Computer Science: Normative System Specification*. John Wiley and Sons, Chichester, England, 1993.
- [39] M. Nickles, F. Fischer, and G. Weiss. Communication attitudes: A formal approach to ostensible intentions, and individual and group opinions. *Electronic Notes in Theoretical Computer Science*, 157(4):95–115, 2006.
- [40] J.R. Searle. *Speech Acts: an Essay in the Philosophy of Language*. Cambridge University Press, Cambridge, UK, 1969.
- [41] J.R. Searle. *The Construction of Social Reality*. The Free Press, New York, 1995.
- [42] Y. Shoham and M. Tennenholtz. On social laws for artificial agent societies: off-line design. *Artificial Intelligence*, 73 (1-2):231 – 252, 1995.
- [43] Y. Shoham and M. Tennenholtz. On the emergence of social conventions: Modeling, analysis and simulations. *Artificial Intelligence*, 94(1–2):139–166, 1997.
- [44] H.A. Simon. A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69:99–118, 1955.
- [45] M.P. Singh. A social semantics for agent communication languages. In *Issues in Agent Communication (AC'00)*, LNCS 1916, pages 31 – 45. Springer, Berlin, 2000.
- [46] R.C. Stalnaker. Common ground. *Linguistics and Philosophy*, 25:701–719, 2002.
- [47] M. Tennenholtz. On stable social laws and qualitative equilibria. *Artificial Intelligence*, 102(1):1–20, 1998.
- [48] L. van der Torre. Contextual deontic logic: Normative agents, violations and independence. *Annals of Mathematics and Artificial Intelligence*, 37(1-2):33–63, 2003.
- [49] L. van der Torre and Y. Tan. Contrary-to-duty reasoning with preference-based dyadic obligations. *Annals of Mathematics and Artificial Intelligence*, 27(1-4):49–78, 1999.
- [50] H. Verhagen. On the learning of norms. In *Procs. of Modelling Autonomous Agents in a Multi-Agent World (MAAMAW'99)*, 1999.
- [51] G. H. von Wright. Deontic logic. *Mind*, 60:1–15, 1951.
- [52] Douglas N. Walton and Erik C. Krabbe. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. State University of New York Press, New York, 1995.
- [53] M. Wooldridge and P.E. Dunne. On the computational complexity of qualitative coalitional games. *Artificial Intelligence*, 158(1):27–73, 2004.

Author Biographies

boella.jpg

Guido Boella is received the PhD degree at the University of Torino in 2000. He is currently professor at the Department of Computer Science of the University of Torino. His research interests include multi-agent systems, in particular, normative systems, institutions and roles using qualitative decision theory. He organized the first workshops on normative multi-agent systems NorMAS, on coordination and organization CoOrg, and the AAAI Fall Symposium and ECOOP Roles.

Parice Caire is currently PhD student at the University of Luxembourg. She is interested in digital cities, virtual communities and tools for conviviality.

insert photo

insert photo

Leendert van der Torre is currently full professor in intelligent systems at the University of Luxembourg. He received the PhD degree from Erasmus University Rotterdam in 1997. He has developed the so-called input/output logics and the BOID agent architecture. His current research interests include deontic logic, qualitative game theory and coordination in normative multiagent systems. He is responsible for priority P1 on security and reliability of University of Luxembourg.

Correspondence and offprint requests to: Leendert van der Torre, Computer Science and Communication, University of Luxembourg, Luxembourg. Email: leendert@vandertorre.com