

Managing Social Influences through Argumentation-Based Negotiation

Nishan C. Karunatilake¹, Nicholas R. Jennings¹, Iyad Rahwan², Sarvapali D. Ramchurn¹

¹ School of Electronics and Computer Science
University of Southampton
Southampton, SO17 1BJ, UK.
{nnc02r,nrj,sdr}@ecs.soton.ac.uk

²Institute of Informatics
The British University in Dubai
P.O.Box 502216 Dubai, UAE.
irahwan@acm.org

ABSTRACT

Social influences play an important part in the actions that an individual agent may perform within a multi-agent society. However, the incomplete knowledge and the diverse and conflicting influences present within such societies, may stop an agent from abiding by all its social influences. This may, in turn, lead to conflicts that the agents need to identify, manage, and resolve in order for the society to behave in a coherent manner. To this end, we present an empirical study of an argumentation-based negotiation (ABN) approach that allows the agents to detect such conflicts, and then manage and resolve them through the use of argumentative dialogues. To test our theory, we map our ABN model to a multi-agent task allocation scenario. Our results show that using an argumentation approach allows agents to both efficiently and effectively manage their social influences even under high degrees of incompleteness. Finally, we show that allowing agents to argue and resolve such conflicts early in the negotiation encounter increases their efficiency in managing social influences.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]

General Terms

Theory, Algorithms, Experimentation

Keywords

Argumentation-based Negotiation, Conflict Resolution.

1. SOCIAL ARGUMENTATION MODEL

In abstract, our framework consists of four main elements: (i) a *schema* for reasoning about social influence, (ii) a set of *social arguments* that make use of this schema, (iii) a *language and protocol* for facilitating dialogue about social influence, and (iv) a set of *decision functions* that agents may use to generate dialogues within the protocol (for a comprehensive formal representation of the framework refer to [3, 4]).

1.1 Social Influence Schema

The notion of *social commitment* [1] acts as our basic building block for capturing social influence. In essence, a social commitment ($SC_g^{x \rightarrow y}$) is a commitment by one agent x (termed the *debtor*)

to another y (termed the *creditor*) to perform a stipulated action. As a result of such a social commitment, the debtor is said to attain an *obligation* toward the creditor, to perform the stipulated action. The creditor, in turn, attains certain rights (i.e., the right to demand or require the stipulated action, the right to question its non-performance, and the right to demand compensation to make good any losses suffered due to its non-performance). We refer to these as *rights to exert influence*. We extend this notion to capture social influences resulting due to factors such as roles and relationships within a wider multi-agent society (i.e., those that rely on the structure of the society, rather than the specific individuals who happen to be committed to one another). Specifically, since most relationships involve the related parties carrying out certain actions for each other, we can view a relationship as an encapsulation of social commitments between the associated roles. Figure 1 formulates these notions as a schema of social influence.

1.2 Social Arguments

The social influence schema can be used to systematically identify social arguments that can be used to negotiate in the presence of social influences. Specifically, we identify two major ways:

Socially Influencing Decisions: In a social context, an agent can affect another agent's decisions by arguing about the validity of the other's social reasoning. Specifically, agents can argue about the premises within the social influence schema to undercut and rebut the justifications to perform certain actions within a society.

Negotiating Social Influence: Agents can also use social influences (i.e., obligations and rights) as additional parameters within their negotiations. Specifically, here the agents use negotiation as a tool for trading social influences.

1.3 Language and Protocol

To enable agents to express their arguments, we define two complementary languages: the *domain language* and the *communication language*. The former allows the agents to express premises about their social context and the latter to construct arguments to engage in their discourse to resolve conflicts. The protocol indicates the legal ordering of communication utterances and has six main stages: (i) *opening*, (ii) *conflict recognition*, (iii) *conflict diagnosis*, (iv) *conflict management*, (v) *agreement*, and (vi) *closing*. In operation, it is defined as a dialogue game protocol which gives locutions rules, commitment rules, and structural rules for each of the locutions in the communication language.

1.4 Decision Making Functionality

The protocol described above gives agents a number of different options, at various stages, as to what utterances to make. For instance, after a proposal the receiving agent could either accept or reject it. After a rejection, the agent may choose to challenge this rejection, end the dialogue, or forward an alternative proposal. An

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'06 May 8–12 2006, Hakodate, Hokkaido, Japan.
Copyright 2006 ACM 1-59593-303-4/06/0005 ...\$5.00.




Figure 1: Schema of Social Influence.

	r_0	r_1	r_2
r_0	0	1	0
r_1	1	0	1
r_2	0	1	0

(a) Rol-Rel map

	r_0	r_1	r_2
r_0	[0:0]	[200:0]	[0:0]
r_1	[400:100]	[0:0]	[200:600]
r_2	[0:0]	[700:200]	[0:0]

(b) Social commitment map

	r_0	r_1	r_2
a_0	1	0	0
a_1	0	1	1
a_2	0	1	0

(c) Ag-Rol map

Figure 2: Social Influence Model.

agent, therefore, still requires a decision mechanism for selecting a particular utterance among the available legal options. To this end, for each of the possible dialogue moves, we specify general decision making algorithms to give the agents that capability (see [3]).

2. ARGUMENTATION CONTEXT

To evaluate how our argumentation model can be used to manage social influences, we require a computational context, in which a number of agents interact in the presence of social influences and conflicts arise as a natural consequence of these interactions. In abstract, the context is based on a multi-agent task allocation scenario where a collection of self-interested agents interact to obtain services to achieve a given set of actions (similar to that of [2]).

2.1 Modelling Social Influences

To map social influences to this context we *first* embody a role-relationship structure into the multi-agent society by randomly linking a defined number of roles to create relationships (see Figure 2(a) where 1 indicates that a relationship exists between the two related roles). *Second*, to this structure, we randomly specify a set of social commitments. We do so by defining a decommitment penalty cost for each capability in each of the active edges (see Figure 2(b)). For instance, in Figure 2(b) the entry [400:100] in row 1, column 2 indicates that the role r_0 is committed to provide capabilities c_0 and c_1 to a holder of the role r_1 . If the agent holding the role r_0 chooses to violate this commitment it will have to pay 400 and 100 (respectively for c_0 and c_1) if asked. *Finally*, we assign these roles to the actual agents using the mapping shown in Figure 2(c).

Since one of the aims in our experiments is to test how agents use argumentation to manage and resolve conflicts created due to incomplete knowledge about their social influences, we generate a number of settings by varying the level of knowledge seeded to the agents. Specifically, we give only a subset of the agent-role mapping. Thus, a certain agent may not know all the roles that it or another agent may act. This lack of knowledge may, in turn, lead to conflicts within the society, since certain agents may know certain facts about the society that others are unaware of. By controlling the level of missing knowledge, we generate an array of settings ranging from perfect knowledge (0% missing knowledge) to the case where agents are completely unaware of their social influences (100% missing knowledge) within the society.

2.2 Agent Interaction

Agents within the system argue and negotiate with each other to find willing and capable partners to accomplish their actions. In essence, an agent that requires a certain capability will generate and forward *proposals* to another selected agent within the community requesting it to sell its services in exchange for a certain reward. If

Algorithm 1 The *negotiate()* method.

```

1:  $[p_0, p_1, \dots, p_{max}] \leftarrow generateProposals()$ 
2:  $p \leftarrow p_0$ 
3:  $isAccepted \leftarrow false$ 
4: {Loop till either the agent agrees or the last proposal fails.}
5: while ( $isAccepted \neq true \parallel p \leq p_{max}$ ) do
6:    $response \leftarrow PROPOSE(p)$ 
7:   if ( $response = "accept"$ ) then
8:      $isAccepted \leftarrow true$ 
9:   else
10:    if ( $p \neq p_{max}$ ) then
11:       $p \leftarrow getNextViableProposal()$ 
12: return  $isAccepted$ 

```

Algorithm 2 The *argue()* method.

```

1:  $H_o \leftarrow challengeJustification()$  {Challenge for the opponent's justification}
2:  $H_p \leftarrow generateJustification()$  {Generate personal justification}
3: if ( $isValid(H_o) = false$ ) then
4:   {Assert invalid premises of  $H_o$ }
5: else
6:   {Adopt premises of  $H_o$  into personal knowledge}
7: if ( $isValid(H_p) = false$ ) then
8:   {Correct invalid premises of  $H_p$  within personal knowledge}
9: else
10:  {Assert  $H_p$ }

```

the receiving agent perceives this proposal to be viable and believes it is capable of performing it, then will *accept* it. Otherwise it will *reject* the proposal. In case of a reject, the original proposing agent will attempt to forward a modified proposal (Algorithm 1).

Agents detect conflicts (such as the one highlighted in Section 2.1) by analysing the decommitment penalties paid by their counterparts for violating their social commitments. Once detected, agents attempt to resolve them by exchanging their respective justifications. These would take the form of the social influence schema and are then analysed to diagnose the cause of the conflict. If there are inconsistencies between them, social arguments are used to highlight these. If they are both valid, then each agent would point-out alternative justifications via asserting missing knowledge (Algorithm 2).

3. MANAGING SOCIAL INFLUENCES

To empirically evaluate how our ABN model allows agents to manage their social influences, we define a number of different interaction strategies to allow the agents to manage conflicts related to their social influences and experiment their relative performance benefits. The underlying motivation for these strategies is our social influence schema, which gives the agents different rights; namely the right to demand compensation and the right to challenge non-performance of social commitments.

3.1 Demanding Compensation

If an agent violates a social commitment, one of the ways its counterpart can react is by exercising its right to demand compensation (see Section 1.1). This formulates our baseline strategy which extends our negotiation algorithm by allowing the agents to demand compensation in cases where negotiation fails (Algorithm 3). Once requested, the agent that violated its social commitment will pay the related the penalty. However, in imperfect information settings, a particular agent may violate a social commitment simply because it was not aware of it (i.e., due to the lack of knowledge of its roles or those of its counterparts). In such situations, an agent may pay a decommitment penalty different to what the other believes it should get, which may, in turn, lead to conflicts. In such situations, our second strategy allows agents to use social arguments to argue about their social influences and, thereby, manage their conflicts (Algorithm 4). Our hypothesis here is that by allowing agents to argue about their social influences we are providing them with a coherent mechanism to manage and resolve their conflicts and, thereby, allowing them to gain a better outcome as a society. To this end, the former strategy acts as our control experiment and the latter as the test experiment (see Figure 3).

Algorithm 3 Claim-Penalty-Non-Argue strategy.

```

1:  $isAccepted \leftarrow negotiate()$ 
2: if ( $isAccepted = \text{false}$ ) then
3:    $compensation \leftarrow demandCompensation()$ 

```

Algorithm 4 Claim-Penalty-Argue strategy.

```

1:  $isAccepted \leftarrow negotiate()$ 
2: if ( $isAccepted = \text{false}$ ) then
3:    $compensation \leftarrow demandCompensation()$ 
4:   if ( $compensation < rightToPenalty$ ) then
5:      $argue()$ 

```

Observation 1: The argumentation strategy allows agents to manage their social influences even at high uncertainty levels.

The effectiveness of the the non-argue strategy falls more rapidly than the argue one. This is because the argue method allows agents to manage and resolve conflicts of opinion that they may have about their social influences. For instance, if a certain agent is unaware of a role that another acts, it may correct this through arguing with that agent. Thus, arguing allows agents to correct such gaps in their knowledge and, thereby, resolve any conflicts that may arise as a result. In this manner, ABN allows the agents to manage their social influences even at high uncertainty levels. Thereby, as a society, the agents can accomplish more of their actions and gain a higher total earnings value. The non-arguing approach, which does not allow them to argue about their social influences, leaves such conflicts unresolved and reduces the population earnings as knowledge imperfections increase within the society.

Observation 2: In cases of perfect information and complete uncertainty both strategies perform equally.

The reason for both strategies performing equally when there is perfect information (0 level) is because there are no knowledge imperfections. Thus, agents do not need to argue to correct conflicts simply because such conflicts do not exist. The reason for both strategies performing equally when there is a complete lack of knowledge is because none of the agents in the society are aware of any social influences (even though they exist). Thus, they are not able to detect any conflicts or violations and, as a consequence, agents do not resort to arguing to manage such conflicts. Therefore, when there is a complete lack of knowledge, the strategy that uses the argue performs the same as the non-argue one.

Observation 3: At all knowledge levels, the argumentation strategy exchanges fewer messages than the non-arguing one.

The reason for this is that, even though agents use some messages to argue and correct their incomplete knowledge, thereafter the agents use their corrected knowledge in subsequent interactions. However, if the agents do not argue to correct such knowledge imperfections, they negotiate more frequently since they cannot use their social influence. Thus, this one-off increase of argue messages becomes insignificant when compared to the increase in the propose, accept, and reject messages due to the increased number of negotiations.

3.2 Questioning Non-Performance

When a particular social commitment is violated our social influence schema also gives the agents the right to challenge and demand a justification for this non-performance. It is generally argued in ABN theory that allowing agents to exchange meta-information gives them the capability to understand each others' reasons and, thus, provides a more efficient method of resolving conflicts under uncertainty. In a similar manner, we believe that providing the agents with the capability to challenge and demand justifications for violating social commitments allows the agents to gain a wider understanding of the internal and social influences affecting their counterparts, thereby, providing a more efficient method for managing social influences in the presence of incomplete knowledge. Using this intuition as the underlying hypothesis, we use our




Figure 3: Efficiency and Effectiveness of the Argue and Non-Argue strategies with 20 Agents and 3 Roles.




Figure 4: Efficiency and Effectiveness of the various argumentation strategies.

previous best strategy *Claim-Penalty-Argue* as the control experiment and design two other strategies (*Argue-In-First-Rejection* and *Argue-In-Last-Rejection*) to experiment with the effect of allowing the agents to challenge non-performance at different stages within the negotiation encounter. The former allows the agent to challenge after the receipt of the first rejection and the latter after the last rejection. Figure 4 shows our results.

Observation 4: The effectiveness of the various argumentation strategies are broadly similar.

This is because all three strategies argue and resolve conflicts even though they decide to argue at different points within the encounter. Thus, we do not expect to have significant differences in the number of conflicts resolved. Thus, the effectiveness stays the same.

Observation 5: Challenging earlier in the dialogue, significantly increases the efficiency of managing social influences.

The reason for this behaviour is based on how the agents use these reasons exchanged during the argue phase. In the *Claim-Penalty-Argue* strategy the main objective of arguing is to resolve the conflict regarding the penalty value that should be paid. It does not however, attempt to find out the reason for why its counterpart rejected its proposal. By challenging the reason for the rejection, the latter two strategies gain certain meta-information, which the agents constructively use in their subsequent interactions. For instance, if the counterpart rejected the proposal due to lack of capability, it can be excluded in future if the agent requires a capability which is equal or greater. Arguing in the first rejection provides this information earlier in the negotiation, which, in turn, gives the agent more capacity to exploit such information (even in the present negotiation) than getting it in the last encounter.

4. REFERENCES

- [1] C. Castelfranchi. Commitments: From individual intentions to groups and organizations. In *Proc. of ICMAS'95*, pages 41–48, San Francisco, CA, 1995.
- [2] N. C. Karunatilake and N. R. Jennings. Is it worth arguing? In *Argumentation in Multi-Agent Systems*, LNAI 3366, pages 234–250, NY, USA, 2004.
- [3] N. C. Karunatilake, N. R. Jennings, I. Rahwan, and T. J. Norman. Arguing and negotiating in the presence of social influences. In *Proc. of CEEMAS'05*, LNAI 3690, pages 223–235, Budapest, Hungary, 2005.
- [4] N. C. Karunatilake, N. R. Jennings, I. Rahwan, and T. J. Norman. Argument-based negotiation in a social context. In *Proc. of ArgMAS'05*, pages 74–88, Utrecht, The Netherlands, 2005.