

Argument-Based Negotiation in a Social Context^{*}

Nishan C. Karunatilake¹, Nicholas R. Jennings¹, Iyad Rahwan², and Timothy J. Norman³

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

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

³ Department of Computing Science, University of Aberdeen, Aberdeen, UK.
tnorman@csd.abdn.ac.uk

Abstract. *Argumentation-based negotiation* (ABN) provides agents with an effective means to resolve conflicts within a multi-agent society. However, to engage in such argumentative encounters the agents require the ability to generate arguments, which, in turn, demands four fundamental capabilities: a schema to reason in a social context, a mechanism to identify a suitable set of arguments, a language and a protocol to exchange these arguments, and a decision making functionality to generate such dialogues. This paper focuses on the first two issues and formulates models to capture them. Specifically, we propose a coherent schema, based on social commitments, to capture social influences emanating from the roles and relationships of a multi-agent society. After explaining how agents can use this schema to reason within a society, we then use it to identify two major ways of exploiting social influence within ABN to resolve conflicts. The first of these allows agents to argue about the validity of each other's social reasoning, whereas the second enables agents to exploit social influences by incorporating them as parameters within their negotiation. For each of these, we use our schema to systematically capture a comprehensive set of social arguments that can be used within a multi-agent society.

Key words: Argumentation-based Negotiation, Conflict Resolution.

1 Introduction

Multi-agent systems consist of a collection of autonomous agents that inter-operate within a shared social context and that perform actions to achieve their individual and collective objectives. In such situations, the actions of these individual agents are influenced via two broad forms of motivations. First, the *internal influences* reflect the intrinsic motivations that drive the individual agent to achieve its own internal objectives. Second, as agents reside and operate within a social community, the social context itself influences their actions. Here, we categorise these latter forms as *social influences*. Now, in many cases, both forms of influence may be present and they may give conflicting motivations to the individual agent. For instance, an agent may be internally motivated to perform a specific action, whereas at the same time it may also be subject to an external social influence not to perform it. Also an agent may face situations where different social influences motivate it in a contradictory fashion (one to perform a specific action and the other not to). Moreover, in many cases agents have to carry out

^{*} The first author is a full time PhD student funded by EPSRC under the project Information Exchange (GR/S03706/01). The authors also extend their gratitude to Pietro Panzarasa, Chris Reed, and Xudong Luo for their thoughts, contributions and discussions.

their actions in environments with incomplete information. Thus, for instance, they may not be aware of the existence of all the social influences that could or indeed should affect their actions and they may also lack the knowledge of certain specific internal influences that drive other agents' behaviours.

In such complex and uncertain environments the need for the agents to interact efficiently and effectively becomes paramount. Given this, *Argumentation-Based Negotiation* (ABN) has been advocated as a promising form of interaction that allows agents to resolve their conflicts within such a society [1, 2]. In more detail, ABN allows agents to exchange additional meta-information such as justifications, critiques, and other forms of persuasive locutions within their interactions. These, in turn, allow agents to gain a wider understanding of the internal and social influences affecting their counterparts, thereby making it easier to resolve certain conflicts that arise due to incomplete knowledge. Furthermore, the negotiation element within ABN also provides a means for the agents to achieve mutually acceptable agreements to the conflicts of interests that they may have in relation to their different influences.

Against this background, the main long term objective of our work is to formulate an agent society that can use such argumentative dialogues to resolve their conflicts. Now, one of the central features required by an agent to engage in such arguments within a society is the capability to generate valid arguments during the discourse of the dialogue. We believe this demands four fundamental capabilities: (i) a schema to reason in social settings; (ii) a mechanism to identify a suitable set of arguments; (iii) a language and a protocol to exchange these arguments; and (iv) a decision making functionality to generate such dialogues. This paper focuses on the first two issues and formulates models to capture them. In so doing, we make three main contributions to the state of the art. *First*, we develop a coherent schema of social influence that provides agents with a means to reason about their actions within a society. Here we use the notion of social commitments as the basic building block for our schema and extend this notion to capture social influences emanating from the roles and relationships of a multi-agent society (see Section 2). *Second*, we illustrate how agents can use our social influence schema to *systematically* derive arguments to negotiate and resolve conflicts within a social context. In so doing, we highlight two major ways that agents can use this schema. The first allows them to affect each others' decisions indirectly by arguing about the social influences that determine their decisions. The second allows the agents to impact each others' decisions by exploiting social influences as parameters within their negotiations (see Section 3). *Third*, we perform a detailed analysis on how agents can use both these forms of social arguments to resolve conflicts with respect to existing social influences and to negotiate new social influences within a multi-agent society (see Section 4).

2 Capturing Social Influence

As mentioned in Section 1, many different forms of external influence affect the actions that an agent performs within a society. Moreover, these social influences emanate from different elements of the society. In particular, many researchers now perceive a society as a collection of *roles* inter-connected via a web of *relationships* [3, 4]. These roles and relationships represent two important aspects of social influence within a society. Specifically, when an agent operates within such a social context, it may assume certain specific *roles*, which will, in turn, guide the actions it performs. In a similar manner, the *relationships* connecting the agents enacting their respective roles also influence the actions they perform. To date, an array of existing research, both in social science and in multi-agent systems, attempts to capture the influences of these social factors on the behaviour of the individual. Nevertheless, there is little in the way of

consensus at an overarching level [5]. Given this, we progressively introduce what we believe are a minimal set of key notions and explain how we adapt them to build a coherent schema of social influence.

The notion of *social commitment* acts as our basic building block for capturing social influence. First introduced by Castelfranchi [6], it remains simple, yet expressive, and is arguably one of the fundamental approaches for modelling social behaviour among agents in multi-agent systems. In essence a social commitment (SC) is a commitment by one agent to another to perform a stipulated action. More specifically, it is defined as a four tuple relation:

$$SC = (x, y, \theta, w)$$

where x identifies the agent who is socially committed to carry out the action (termed the *debtor*), y the agent to whom the commitment is made (termed the *creditor*), θ the associated action, and w the witness of this social commitment.¹ Having defined social commitment, Castelfranchi further explains its consequences for both the agents involved. In detail, a social commitment results in the debtor attaining an *obligation* toward the creditor, to perform the stipulated action. The creditor, in turn, attains certain rights. These include the right to demand or require the performance of the action, the right to question the non-performance of the action, and, in certain instances, the right to make good any losses suffered due to its non-performance. We refer to these rights the creditor gains as the *rights to exert influence*.

This notion of social commitment resulting in an obligation and rights to exert influence, allows us a means to capture social influences between two agents. Thus, when a certain agent is socially committed to another to perform a specific action, it subjects itself to the social influences of the other to perform that action. The ensuing obligation, on one hand, allows us to capture how an agent gets subjected to the social influence of another, whereas, the rights to exert influence, on the other hand, model how an agent gains the ability to exert such social influence upon another. Thereby, the notion of social commitment gives an elegant mechanism to capture social influence resulting between two agents.

Given this basic building block for modelling social influence between specific pairs of agents, we now proceed to explain how this notion is extended 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. To illustrate this, consider the relationship between the two roles supervisor and student. For instance, assume the relationship socially influences the student to produce and hand over his thesis to the supervisor in a timely manner. This influence we can perceive as a social commitment that exists between the roles supervisor and student (the student is socially committed to the supervisor to perform the stipulated action). As a consequence of this social commitment, the student attains an obligation toward the supervisor to carry out this related action. On the other hand, the supervisor gains the right to exert influence on the student by either demanding that he does so or through questioning his non-performance. In a similar manner, the supervisor may be influenced to review and comment on the thesis. This again is another social commitment associated with the relationship. In this instance, it subjects the supervisor to an obligation to review the thesis while the student gains the right to demand its performance. In this manner, social commitment again provides

¹ In the desire to maintain simplicity within our schema, we avoid incorporating the witness in our future discussions (as Castelfranchi did in his subsequent expositions).

an effective means to capture the social influences emanating through roles and relationships of the society (independently of the specific agents who take on the roles).

This extension to the basic definition of social commitment is inspired primarily by the work of Cavedon and Sonenberg [3]. Their work investigates how different social influences emanating via roles and relationships affect the agent's prioritising of goals. However, we refrain from going into the level of modalities of agents (such as goals, beliefs and intentions), but rather stay at the level of actions.² The motivation for doing so is twofold. First, our primary interest in this work is to use our model to capture arguments that our agents can use to argue about their actions in an agent society. We also aim to implement this argumentation system and test its performance under various arguing strategies (see Section 6). To this end, we believe a model that focuses on the level of actions, as opposed to goals, beliefs and intentions, will reduce the complexity of our future work. Second, an agent adopting a goal, a belief or an intention can be perceived as an act that it performs. Therefore, focusing on the level of actions loses little in terms of expressiveness.

Our extension also adds certain modifications to the original definition of social commitment. In more detail, we allow a social commitment to exist between roles and not only between agents. The rationale for doing so is to relax the highly constraining requirement present within the Cavedon and Sonenberg model that forces all known roles in a relationship to be filled if any one is occupied. To explain this, consider the previous example relationship between the roles student and supervisor. If we define the social commitment between these two roles it captures the general influence within the relationship. Thus, if some particular person (e.g. Andy) assumes the role of student, he would still be obligated to produce the thesis to his supervisor even though, at the moment, the school has not appointed a specific supervisor to him. Therefore, this subtle deviation allows the agents to maintain a social commitment even though the other party of the relationship is not instantiated.

It is also important to clarify our notion of obligation. Here, we do not strictly adhere to the analysis of Castelfranchi that an honest agent will always gain an internal commitment (resulting in an intention to perform that action) for all its social commitments [6]. On the contrary, in accordance with the works of Cavedon and Sonenberg [3] and Dignum *et al.* [5, 7], we believe that all ensuing obligations resulting due to social commitments exert their own degree of influence upon the individual. Thus, certain social commitments may cause a stronger social influence than others. This is, we believe, an important characteristic in realistic multi-agent societies, where autonomous agents are subjected to contradicting external influences (which may also conflict with their internal influences). Therefore, if an agent is subjected to obligations that either contradict or hinder each other's performance, the agent will make a choice about which obligation to honour. To facilitate this choice, we allow the agents to associate each obligation (resulting due to a social commitment) with its own specific degree of influence. We believe this degree of influence is dependent on two main factors. First, is the relationship that the social commitment is a part of. In more detail, two different social commitments related with the same action, but part of different relationships, can cause obligations with different degrees of influence to the agent. Second, it is also dependent on the associated action. Thus even in the same relationship, certain obligations associated with certain actions may cause a stronger influence than others. Given this descriptive definition and the underlying motivations of our model of social influence, we now formulate these notions to develop a notational representation of the schema.³

² For an extended logical formalism that captures how both the beliefs and intentions, in addition to the goals, of an agent are affected via social influences refer to [4].

³ However, in the following it is not our objective to provide a formal logical definition to the problem of resolving conflicts among obligations. Such a task is non-trivial and some progress to this end is

Definition 1: Let:

- $A = \{a_1, \dots, a_n\}$ denote a finite set of agents,
- $R = \{r_1, \dots, r_n\}$ denote a finite set of roles,
- $P = \{p_1, \dots, p_n\}$ denote a finite set of relationships,
- $\Theta = \{\theta_1, \dots, \theta_n\}$ denote a finite set of actions,
- $\text{Act} : A \times R$ denote the fact that an agent is acting a role,
- $\text{RoleOf} : R \times P$ denote the fact that a role is related to a relationship, and
- $\text{In} : A \times R \times P$ denote the fact that an agent acting a role is part of a relationship.

If an agent acts a certain role and that role is related to a specific relationship, then that agent acting that role is said to be part of that relationship (as per Cavedon and Sonenberg [3]):

$$\text{Act}(a, r) \wedge \text{RoleOf}(r, p) \rightarrow \text{In}(a, r, p) \quad (\text{Rel. Rule})$$

Definition 2: Let SC denote a finite set of social commitments and $SC_\theta^{x \rightarrow y} \in SC$ denote a social commitment where x is the entity (agent or role) acting as the debtor, y is the entity acting as the creditor, and θ is the related action.

A social commitment $SC_\theta^{x \rightarrow y}$ will result in the debtor attaining an obligation toward the creditor to perform a stipulated action and the creditor, in turn, attaining the right to influence the performance of that action (as per Castelfranchi [6]):

$$SC_\theta^{x \rightarrow y} \rightarrow [O_\theta^{x \rightarrow y}]_x^f \wedge [R_\theta^{y \rightarrow x}]_y, \quad (\text{S-Com Rule})$$

where:

- $[O_\theta^{x \rightarrow y}]_x^f$ represents the obligation that x attains that subjects it to an influence of a degree f (see the previous description in Section 2) toward y to perform θ and
- $[R_\theta^{y \rightarrow x}]_y$ represents the right that y attains which gives it the ability to demand, question, and require x regarding the performance of θ .

Definition 3: Let:

- $\text{DebtorOf} : (R \cup A) \times SC$ denote that a role (or an agent) is the debtor in a social commitment,
- $\text{CreditorOf} : (R \cup A) \times SC$ denote that a role (or an agent) is the creditor in a social commitment,
- $\text{ActionOf} : \Theta \times SC$ denote that an act is associated with a social commitment, and
- $\text{AssocWith} : SC \times P$ denote that a social commitment is associated with a relationship.

If the roles associated with the relationship are both the creditor and the debtor of a particular social commitment, then we declare that social commitment is associated with the relationship (as per Section 2).

Given these definitions, we can capture the social influences within multi-agent systems as a schema. To this end, Figure 1 gives a natural language representation of the schema and a notational representation is captured via formulae (1) through (6). In the following section, we will use this schema to systematically capture the social arguments that agents can use to argue in societies.

Applying Rel. Rule to a society where: $a_i, a_j \in A \wedge r_i, r_j \in R \wedge p \in P$ s.t. $\text{Act}(a_i, r_i), \text{Act}(a_j, r_j), \text{RoleOf}(r_i, p), \text{RoleOf}(r_j, p)$ hold true, we obtain:

$$\text{Act}(a_i, r_i) \wedge \text{RoleOf}(r_i, p) \rightarrow \text{In}(a_i, r_i, p) \quad (1)$$

$$\text{Act}(a_j, r_j) \wedge \text{RoleOf}(r_j, p) \rightarrow \text{In}(a_j, r_j, p). \quad (2)$$

reported in the works of Torre & Tan [8] and Ross [9]. However, a detailed discussion is not within the scope of this paper.

An agent a_i acting the role r_i
 Leads it to be part of the relationship p
 With another agent a_j acting the role r_j

A social commitment $SC_{\theta}^{r_i \rightarrow r_j}$ associated with p

- Leads to a_i attaining an obligation O toward r_j ,
 Which subjects it to an influence of degree f
 To perform the action θ
- And, in turn, leads to a_j attaining the right R toward r_i
 To demand, question and require the performance of action θ

Fig. 1. Natural Language Representation of the Schema of Social Influence.

Now, consider a social commitment $SC_{\theta}^{r_i \rightarrow r_j}$ associated with the relationship p in this society. Applying this to Definition 3 we obtain:

$$\begin{aligned} & (\text{DebtorOf}(r_i, SC) \wedge \text{RoleOf}(r_i, p)) \wedge (\text{CreditorOf}(r_j, SC) \wedge \text{RoleOf}(r_j, p)) \\ & \wedge \text{ActionOf}(\theta, SC) \rightarrow \text{AssocWith}(SC_{\theta}^{r_i \rightarrow r_j}, p). \end{aligned} \quad (3)$$

Applying the S-Comm rule to $SC_{\theta}^{r_i \rightarrow r_j}$ we obtain:

$$SC_{\theta}^{r_i \rightarrow r_j} \rightarrow [O_{\theta}^{r_i \rightarrow r_j}]_{r_i}^f \wedge [R_{\theta}^{r_j \rightarrow r_i}]_{r_j}. \quad (4)$$

Combining (4), (1) and (3) we obtain:

$$\text{In}(a_i, r_i, p) \wedge \text{AssocWith}(SC_{\theta}^{r_i \rightarrow r_j}, p) \rightarrow [O_{\theta}^{a_i \rightarrow r_j}]_{a_i}^f. \quad (5)$$

Combining (4), (2) and (3) we obtain:

$$\text{In}(a_j, r_j, p) \wedge \text{AssocWith}(SC_{\theta}^{r_i \rightarrow r_j}, p) \rightarrow [R_{\theta}^{a_j \rightarrow r_i}]_{a_j}. \quad (6)$$

3 Capturing Social Arguments

When agents operate within a society of incomplete information with diverse and conflicting influences, they may, in certain instances, lack the knowledge, the motivation and/or the capacity to enact all their social commitments. However, to function as a coherent society it is important for these agents to have a means to resolve such conflicts and come to a mutual understanding about their actions. As argued in Section 1, ABN provides one such means. However, to argue in such a society, the agents need to have the capability to first identify the arguments to use. To this end, here we present how agents can use our social influence schema to systematically identify arguments to negotiate within a society. We term these arguments *social arguments*, not only to emphasise their ability to resolve conflicts within a society, but also to highlight the fact that they use the social influence present within the system as a core means in changing decisions and outcomes within the society. More specifically, we have identified two major ways in which social influence can be used to change decisions and outcomes and thereby resolve conflicts between agents. These are depicted in Figure 2 and are described in more detail in the following.




Fig. 2. Interplay of Social Influence and Argumentation-Based Negotiation.

3.1 Socially Influencing Decisions

One way to affect an agent’s decisions is by arguing about the validity of that agent’s practical reasoning [10, 11]. Similarly, in a social context, an agent can affect another agent’s decisions by arguing about the validity of the latter’s social reasoning. In more detail, agents’ decisions to perform (or not to perform) actions are based on their internal and/or social influences. Thus, these influences formulate the justification (or the reason) behind their decisions. Therefore, agents can affect each other’s decisions indirectly by affecting the social influences that determine their decisions (see Figure 2(a)). Specifically, in the case of actions motivated via social influences through the roles and relationships of a structured society, this justification to act (or not to act) flows from the social influence schema (see Section 2). Given this, we can further classify the ways that agents can socially influence each other’s decisions into two broad categories:

1. Undercut the opponent’s existing justification to perform (or not) an action by disputing certain premises within the schema which motivates its opposing decision.
2. Rebut the opposing decision to act (or not) by,
 - (a) Pointing out information about an alternative schema that justifies the decision not to act (or act as the case may be).
 - (b) Pointing out information about conflicts that could or should prevent the opponent from executing its opposing decision.

Given this, in the following we highlight how agents can systematically use the social influence schema to identify these possible types of arguments to socially influence each other’s decisions. For clarity, we present our arguments both in natural language and using notation. The domain language of our notational representation flows naturally from our schema while the communication language (see Table 1) is inspired from the works of [12], [13], and [14].⁴ To denote the arguments we define three additional predicates (to the ones defined in Section 2); namely (i) $\text{InfluenceOf}(O, f)$ denotes that f is the degree of influence associated with the obligation O ; (ii) $\text{do}(a, \phi)$ (or $\text{do}(\phi)$ in the abbreviated form) indicating the agent a to perform ϕ (where ϕ can be an action, an obligation, a right to influence, *adopt* a new obligation, or *stop* an existing relationship); (iii) $\text{Conflict}(\text{do}(\phi_i), \text{do}(\phi_j))$ denotes a conflict between $\text{do}(\phi_i)$ and $\text{do}(\phi_j)$. Finally, in order to illustrate how agents can exploit third party social influences within their arguments, we denote two additional relationships (apart from p defined in Section 2) as p' and p'' ; the former between the roles r_i and r_k that the agents a_i and a_k hold, and the latter between the roles r_j and r_k that the agents a_j and a_k hold respectively.

⁴ Due to space limitations we intentionally avoid a detailed discussion on the language, the protocol, and the agents’ decision making functions of our ABN system. See [15] for further details.

Table 1. High-level description of the communication language

Locution	Description
<i>OPEN-DIALOGUE</i>	Open the dialogue.
<i>CLOSE-DIALOGUE</i>	Close the dialogue.
<i>ASSERT(l)</i>	Assert a certain proposition <i>l</i> .
<i>CHALLENGE(l)</i>	Challenge the justification for the proposition <i>l</i> .
<i>PROPOSE(do(l) ⇒ do(m))</i>	Propose the performance of <i>l</i> in return for the performance of <i>m</i> .
<i>ACCEPT(do(l) ⇒ do(m))</i>	Accept the performance of <i>l</i> in return for the performance of <i>m</i> .
<i>REJECT(do(l) ⇒ do(m))</i>	Reject the performance of <i>l</i> in return for the performance of <i>m</i> .

1. Dispute (Dsp.) existing premises to undercut the opponent's existing justification.

- i. Dsp. a_i is acting role r_i : $ASSERT(\neg Act(a_i, r_i))$.
- ii. Dsp. a_j is acting role r_j : $ASSERT(\neg Act(a_j, r_j))$.
- iii. Dsp. r_i is related to the relationship p : $ASSERT(\neg RoleOf(r_i, p))$.
- iv. Dsp. r_j is related to the relationship p : $ASSERT(\neg RoleOf(r_j, p))$.
- v. Dsp. SC is associated with the relationship p : $ASSERT(\neg AssocWith(SC_{\theta}^{r_i \rightarrow r_j}, p))$.
- vi. Dsp. f is the degree of influence associated with O: $ASSERT(\neg InfluenceOf(O, f))$.
- vii. Dsp. θ is the action associated with O: $ASSERT(\neg ActionOf(O, \theta))$.
- viii. Dsp. θ is the action associated with R: $ASSERT(\neg ActionOf(R, \theta))$.

2. Point out (P-o) new premises about an alternative schema to rebut the opposing decision.

- i. P-o a_i is acting the role r_i : $ASSERT(Act(a_i, r_i))$.
- ii. P-o a_j is acting the role r_j : $ASSERT(Act(a_j, r_j))$.
- iii. P-o r_i is related to the relationship p : $ASSERT(RoleOf(r_i, p))$.
- iv. P-o r_j is related to the relationship p : $ASSERT(RoleOf(r_j, p))$.
- v. P-o SC is a social commitment associated with the relationship p :
 $ASSERT(AssocWith(SC_{\theta}^{r_i \rightarrow r_j}, p))$
- vi. P-o f is the degree of influence associated with the obligation O: $ASSERT(InfluenceOf(O, f))$.
- vii. P-o θ is the action associated with the obligation O: $ASSERT(ActionOf(O, \theta))$.
- viii. P-o θ is the action associated with the right R: $ASSERT(ActionOf(R, \theta))$.
- ix. P-o a_i 's obligation O to perform: $ASSERT(O_{\theta}^{a_i \rightarrow r_j})$
- x. P-o a_j 's right to demand, question and require the action θ : $ASSERT(R_{\theta}^{a_j \rightarrow r_i})$

3. Point out conflicts that prevent executing the decision to rebut the opposing decision.

- (a) Conflicts with respect to O.
 - i. P-o a conflict between two different obligations due toward the same role:
 $ASSERT(Conflict(do(O_{\theta}^{a_i \rightarrow r_j}), do(O_{\theta'}^{a_i \rightarrow r_j})))$.
 - ii. P-o a conflict between two different obligations due toward different roles:
 $ASSERT(Conflict(do(O_{\theta}^{a_i \rightarrow r_j}), do(O_{\theta'}^{a_i \rightarrow r_k})))$
- (b) Conflicts with respect to R.
 - i. P-o a conflict between two different rights to exert influence upon the same role:
 $ASSERT(Conflict(do(R_{\theta}^{a_j \rightarrow r_i}), do(R_{\theta'}^{a_j \rightarrow r_i})))$
 - ii. P-o a conflict between two different rights to exert influence upon different roles:
 $ASSERT(Conflict(do(R_{\theta}^{a_j \rightarrow r_i}), do(R_{\theta'}^{a_j \rightarrow r_k})))$
- (c) Conflicts with respect to θ and another action θ' such that (i) θ' is an alternative to the same effect as θ ; (ii) θ' either hinders, obstructs, or has negative side effects to θ (refer to [10]).
 $ASSERT(Conflict(do(\theta), do(\theta')))$

3.2 Negotiating Social Influence

In the previous section, we explored various forms of arguments that agents can use to alter the social influences and thereby change each other's decisions. In this section, we explore a different way in which agents can use social reasoning in negotiation. Here, instead of using

social argumentation as a tool to affect decisions, agents use negotiation as a tool for “trading influences”. In other words, the social influences are incorporated as additional parameters of the negotiation object itself [16] (see Figure 2(b)). The following presents a list of what we believe to be the most important social arguments that would allow the agents to exploit social influences in such a manner.

4. Use O as a parameter of negotiation.

- i. Promise to (or threaten not to) undertake one or many future obligations if the other agent performs (or does not perform) a certain action θ .⁵

$$PROPOSE(do(a_j, \theta) \Rightarrow do(a_i, \text{adopt } O_{\theta'}^{a_i \rightarrow a_j}))$$

$$PROPOSE(do(a_j, \theta) \Rightarrow \neg do(a_i, \text{adopt } O_{\theta'}^{a_i \rightarrow a_j}))$$

$$PROPOSE(\neg do(a_j, \theta) \Rightarrow do(a_i, \text{adopt } O_{\theta'}^{a_i \rightarrow a_j}))$$

$$PROPOSE(\neg do(a_j, \theta) \Rightarrow \neg do(a_i, \text{adopt } O_{\theta'}^{a_i \rightarrow a_j}))$$
- ii. Promise to (or threaten not to) honour one or many existing obligations if the other agent performs (or does not perform) a certain action θ : $PROPOSE(\pm do(a_j, \theta) \Rightarrow \pm do(a_i, O_{\theta}^{a_i \rightarrow a_j}))$

5. Use R as a parameter of negotiation.

- i. Promise not to (or threaten to) exercise the right to influence one or many existing obligations if the other agent performs (or does not perform) a certain action θ :
$$PROPOSE(\pm do(a_j, \theta) \Rightarrow \pm do(a_i, R_{\theta'}^{a_i \rightarrow a_j}))$$

6. Use third party obligations and rights as a parameter of negotiation.

- i. Third party obligations
 - i. Promise to (or threaten not to) undertake one or more future obligations toward a_k to perform θ' , if a_j would (or would not) exercise its right to influence a certain agent a_l to perform θ :
$$PROPOSE(\pm do(a_j, R_{\theta}^{a_j \rightarrow a_l}) \Rightarrow \pm do(a_i, \text{adopt } O_{\theta'}^{a_i \rightarrow a_k}))$$
 - ii. Promise to (or threaten not to) honour one or more existing obligations toward a_k to perform θ' , if a_j would (or would not) exercise its right to influence a certain agent a_l to perform θ :
$$PROPOSE(\pm do(a_j, R_{\theta}^{a_j \rightarrow a_l}) \Rightarrow \pm do(a_i, O_{\theta'}^{a_i \rightarrow a_k}))$$
- ii. Third party rights
 - i. Promise to (or threaten not to) exercise the right to influence one or many existing obligations toward a_k to perform θ' , if a_j would honour its existing obligation to perform θ :
$$PROPOSE(do(a_j, O_{\theta}^{a_i \rightarrow a_j}) \Rightarrow \neg do(a_i, R_{\theta'}^{a_i \rightarrow a_k}))$$

$$PROPOSE(\neg do(a_j, O_{\theta}^{a_i \rightarrow a_j}) \Rightarrow do(a_i, R_{\theta'}^{a_i \rightarrow a_k}))$$

7. Use P as a parameter of negotiation.

- i. Threaten to terminate p (its own relationship with a_j) or p' (a third party relationship that a_i has with a_k), if the agent a_j performs (or does not perform) a certain action θ :
$$PROPOSE(\pm do(a_j, \theta) \Rightarrow do(a_i, \text{stop } p))$$

$$PROPOSE(\pm do(a_j, \theta) \Rightarrow do(a_i, \text{stop } p'))$$
- ii. Threaten to influence another agent (a_k) to terminate its relationship p'' with a_j , if a_j performs (or does not perform) a certain action θ .
$$PROPOSE(\pm do(a_j, \theta) \Rightarrow do(a_i, R_{\text{stop } p''}^{a_i \rightarrow a_k}))$$

In summary, these social arguments allow agents to resolve conflicts in two main ways. The first set of arguments facilitate critical discussion about the social influence schema; thus, these allow the agents to critically question and understand the underlying reasons for each others' action. This form of engagement not only allows the agents to extend their incomplete knowledge of the society, but also provides a means to convince their counterparts to change decisions based on such incomplete information, thereby, resolving conflicts within a society. The second set of arguments allows the agents to exploit social influences constructively within their negotiations. Thus, providing agents with additional parameters to influence their counterpart to reach agreements and thereby resolve conflicts through a negotiation encounter. Having systematically captured such social arguments through our schema of social influence, in the following section we present an illustrative case study to analyse their use within ABN to overcome conflicts in multi-agent systems.

⁵ To save space, we will denote these four variations as $PROPOSE(\pm do(a_j, \theta) \Rightarrow \pm do(a_i, \theta'))$.

4 Arguing with Social Influence

To illustrate how the above set of social arguments can be used to resolve conflicts and change outcomes within a social context, we consider the following case study. The scenario is based on a small community comprising three individual agents (referred to as Andy, Ben, and Carl). Each of these agents has certain specific roles in the scenario and they are inter-connected via a set of relationships. Andy, for instance, has two roles (one as a PhD student and another as a project partner), while Ben has the role of supervisor and Carl has the role project manager. The community has two defined relationships; one between Ben and Andy of type supervisor-student and the second between Carl and Andy of type project manager-project partner. To highlight the use of social arguments we consider the following initial setting where Andy is obliged to perform the following three actions:

θ_1 : Obligated toward Ben to *write his thesis*.

θ_2 : Obligated toward Ben to *write a journal paper*.

θ_3 : Obligated toward Carl to *write a software component*.

Apart from these three obligatory tasks, we also assume that Carl wants Andy to undertake an additional obligation to *integrate the software* (referred to as θ_4) to which Andy has expressed his dissent. In this context, we assume that Andy is only capable of performing one of the above tasks due to time restrictions. Therefore, after prioritising them, Andy chooses to perform action θ_1 . In the following we analyse two cases where both Ben and Carl attempt to use the social arguments captured in Section 3 to convince Andy to change his decision. In more detail, the first case analyses Ben's attempt to convince Andy to prioritise writing the journal paper to his thesis. The second case analyses Carl's attempt to negotiate with Andy to undertake the additional obligation of integrating the software. Both of these cases illustrate the use of social arguments to change decisions within a society; the former emphasises their effect in resolving conflicts with respect to existing social influences, while the latter uses them to negotiate new social influences.

4.1 Resolving Conflicts between Existing Social Influences

The following dialogue sets the scene for the first case by highlighting the constraint and the conflict of interest between the different priorities of both Andy and Ben (i.e., Andy wanting to do θ_1 while Ben wanting Andy to do θ_2):⁶

L1 - Ben: *Can you finish the journal paper?*

L2 - Andy: *No, I can't.*

L3 - Ben: *Why not?*

While L2 revealed the conflict, L3 is an attempt by Ben to identify the reason behind Andy's refusal. This is an important junction of the dialogue for Andy. Here, he can either choose to explain the reasons (thereby answer Ben's question) or challenge Ben's right to question him. First, we will consider the later option where Andy chooses to challenge Ben. This can arise due to a number of reasons. For example, due to incomplete information in the system, Andy may not be aware that θ_2 is an obligatory action. On the other hand, he may be aware that θ_2 is an obligation, but may not be aware that the obligation is toward Ben (thus, he is not convinced of Ben's right to question). Due to any of these reasons Andy could challenge

⁶ We choose to denote all argumentative dialogues using natural language, since we believe it allows the reader to easily understand their conceptual differences without worrying about the notational syntax. However, these can be easily encoded in the language highlighted in Table 1 (e.g., see Figure 3).

Ben's right to question, which would shift⁷ the dialogue toward a critical discussion. The social arguments identified in Section 3 allow Ben to respond appropriately to this form of critical questioning, thus, justify his right to question. The following illustrates an example case where Ben combines social arguments 2.i, 2.ix, and 2.x to retort back to Andy's question:

L4 - Andy: Why do you ask?

L5 - Ben: I am your supervisor and you have an obligation toward me to finish the paper, which gives me the right to question its non-performance.

This form of questioning not only allows Andy to expand his incomplete information of the society, but also provides him with a means to filter out the individuals to which he does not have to justify his non-performance. For instance, if George, a PhD colleague of Andy, played the role of the questioner and asked L3, he will not be able to answer Andy's critical question (L4).

Having analysed how Ben can use the social arguments to respond to Andy's critical questioning, we will now proceed to the next important step in this dialogue. Once questioned and convinced of Ben's right to question, Andy is obliged to give reasons for his non-performance. To illustrate this, consider the following dialogue where Andy uses the social argument 3.a.i within L4 as his reason:

L1 - Ben: Can you finish the journal paper?

L2 - Andy: No, I can't.

L3 - Ben: Why not?

L4 - Andy: I have to finish the thesis, and I can't do two things together.

Having established the conflict and the reason behind it, it is now up to Ben to convince Andy to change his priorities (i.e., finish the paper before writing the thesis). To achieve this goal, Ben can use a number of different approaches. In the following we will analyse some of these and illustrate how our captured set of social arguments help Ben to achieve his goal:

1. *Socially Influencing Decisions.*

One option available to Ben is to focus on the validity of Andy's social reasoning. In other words, Ben can attempt to change Andy's decision (to write the thesis over finishing the paper), by indirectly arguing about the social influences that determined this. In argumentative dialogue terms, this reflects a shift toward a more persuasive form of a dialogue [17]. Specifically, the social arguments captured in Section 3.1 provide Ben with a number of means to achieve this. For instance, Ben can attempt to identify the motives that prompted Andy to prioritise the thesis over the paper and argue and persuade him to change his motives. The following dialogue illustrates how Ben uses his expert opinion⁸ to change Andy's perception of the relative importance of these actions and thereby reach an agreement (see Figure 3 for a notational representation of this using our domain and communication language):

L5 - Ben: But you are obliged to finish the paper.

L6 - Andy: Yes, but I am also obliged to write the thesis and I believe it influences me more than the obligation to finish the journal paper.

⁷ A change from one dialogue type to another within the same discussion is generally referred to as a *shift* in argumentation theory. However, due to space restrictions, we avoid an expansive discussion about different dialogue types, shifts between them, and fallacies involved when performing such shifts. For a detailed discussion refer to [17, 18].

⁸ Argument from expert opinion is a specialised type of argument scheme discussed in depth by Walton [11]. One of the main strengths of our social arguments is to provide a means for the agents to exploit such schemes within a social context.

<p>Let:</p> <ul style="list-style-type: none"> - O_1 denote the obligation to perform θ_1 (finishing his thesis) and f_1 its associated degree of influence, - O_2 denote the obligation to perform θ_2 (write a journal paper) and f_2 its associated degree of influence. <p>Ben: <i>OPEN-DIALOGUE</i> Andy: <i>OPEN-DIALOGUE</i> Ben: <i>PROPOSE</i>($do(\theta_2)$) Andy: <i>REJECT</i>($do(\theta_2)$) Ben: <i>CHALLENGE</i>($\neg do(\theta_2)$) Andy: <i>ASSERT</i>($Conflict(do(\theta_2), do(\theta_1))$) Ben: <i>ASSERT</i>($O_2$) Andy: <i>ASSERT</i>($O_1 \wedge (f_1 > f_2)$) Ben: <i>ASSERT</i>($\neg(f_1 > f_2) \wedge (f_2 > f_1)$) Andy: <i>ACCEPT</i>($(f_2 > f_1) \wedge do(\theta_2) \wedge \neg do(\theta_1)$) Ben: <i>CLOSE-DIALOGUE</i> Andy: <i>CLOSE-DIALOGUE</i></p>

Fig. 3. Notational representation of the sample ABN dialogue.

L7 - Ben: *In my expert opinion, I believe it is more important at this point to finish the paper than the thesis. You should change your opinion.*

L8 - Andy: *I adhere to your expert opinion, therefore I will finish the paper.*

Apart from focusing on obligations and rights, Ben could also socially influence Andy by focusing on the related actions. For instance, Ben could reveal additional information that was not readily available for Andy at the time of prioritising his actions. One way of doing this is to highlight the potential merits of writing the paper before the thesis as follows:

L5 - Ben: *If you finish the paper, it will help you to write the thesis since you can reuse the same material.*

Ben can do this by using the social argument 3.c.i. He can also emphasise the potential disadvantages of Andy's choice by making use of the social argument 3.c.ii:

L5 - Ben: *If you attempt to write the thesis without this paper it will negatively affect the quality of your thesis.*

2. *Negotiating social influence.*

Another option Ben has is to focus on using his existing social influences as parameters to negotiate a mutually acceptable agreement and thereby change Andy's original decision. The social arguments captured in Section 3.2 predominantly facilitate this form of an argumentative dialogue. For instance, Ben can focus on his right to influence the obligatory action as follows:

L5 - Ben: *You are obliged to me to finish the paper and I have a right to demand that you do so.*

L6 - Andy: *True, but you also have the right to demand me to write thesis. These rights are in conflict due to time restrictions.*

L7 - Ben: *Agreed. I will promise not to exercise my right to demand the thesis provided that you finish the paper.*

Social arguments captured through the schema would facilitate this discussion. Specifically, for L5 Ben uses the arguments 2.ix and 2.x, for L6 Andy uses 2.x and 3.b.i, and for L7 Ben, in turn, uses 5.i to reach a mutually acceptable agreement. Alternatively, Ben can focus on a reciprocal obligation that he has toward Andy and use it as a parameter in negotiation. The following illustrates how Ben uses social arguments 2.x and 4.ii in combination to forward a threat (L5) in his negotiation:

L5 - Ben: *I am obliged to review your thesis. However, if you do so without writing the paper I will neither read nor review your thesis.*

L6 - Andy: *In that case I will write the paper first.*

4.2 Negotiating New Social Influences

Here we analyse how social arguments provide agents with a means to negotiate new social influences within a society. To this end, we consider a case where Carl is attempting to convince Andy to adopt a new obligation to integrate the software, but Andy is not keen in doing so due to his current time constraints. The following illustrate three possibilities that allow Carl to exploit social influence within his negotiation and in each case we highlight how our social arguments facilitate such an approach:

1. *Exploit existing rights to influence*: Carl can focus his negotiation on his existing right to influence θ_3 . More specifically, Carl can point out his existing right to demand the performance of θ_3 and offer to refrain from doing so if Andy adopts the additional obligation. Since adopting an obligation is also an action (see Section 3.2), the social arguments 2.x and 5.i will facilitate this line of negotiation.
2. *Exploit Andy's current situation*: Carl can also make use of Andy's current situation in his negotiation. In more detail, if Carl knows that Andy is pressed for time to write the paper, he can offer to help him achieve this if, in exchange, Andy agrees to do the integration. In this line of argument, Andy is proposing to undertake an additional obligation (helping Andy write the paper) in exchange for Andy adopting the additional obligation. Carl can achieve this by using the social argument 4.i in his negotiation dialogue.
3. *Trade existing obligations*: Carl can offer to trade with Andy his current obligation to do θ_3 with the new obligation to perform θ_4 . This could be a more agreeable solution for both Carl and Andy especially if performing θ_4 is not an immediate requirement. This line of argument is a combination of the above two since here Carl rescinds his rights on θ_3 (using argument 4.i), whilst Andy adopts a new obligation on θ_4 (using argument 5.i).

5 Related Work

As detailed in Section 1, one of the central features required by an agent to argue and resolve its conflicts is its capability to generate valid arguments during the discourse. This area is extensively researched in current ABN literature and a wide variety of approaches have been proposed to model this capability within a computational entity [2]. In the following, we review some of these and place our model in context by highlighting similarities and distinctions with these efforts.

Our work greatly benefits from the approaches used in *argumentation schemes* [11] to systematically identify arguments. In more detail, argumentation schemes capture stereotypical patterns of reasoning upon which communication structures can be built. Increasingly, they are used in computational contexts, including multi-agent systems, since they hold potential for significant improvements in reasoning and communication abilities in such systems [19]. For instance, the recent work of Reed & Walton [19] presents a general framework for specifying such schemes in computational contexts and the work of Atkinson *et al.* [10] uses this in their schematic approach to capture a particular style of dialogue over actions. In a similar manner, we have captured how agents function and reason *within a society* as a schema of social influence and use it to systematically identify social arguments. In systematically categorising our social arguments, we also draw from the logical approaches to ABN [12, 20]. Broadly, these systems formulate an argument as a certain sequence of inferences leading to a logical conclusion and, in turn, allow the agents to construct attacks by either disqualifying one or more of these inferences (undercut) or formulating an alternative series of inference leading to the

opposite conclusion (rebut). We follow the same systematic manner in organising the different social arguments identified from our schema.

Another fundamental work in computational argument generation is that of Kraus, Sycara & Evenchik [21]. In essence, their work allows agents to use promises, threats and various forms of appeals during a negotiation encounter. Now, our social arguments, particularly those that allow agents to negotiate social influences, hold certain similarities to their forms of threats, promises and appeals. However, there are two important distinctions. First, their main focus is in formulating interactions between two agents, whereas we expressly take into account the impact of society by way of social commitments. Second, they do not take into account incomplete information between the two agents. Thus, they do not provide agents with a means to resolve conflicts due to such imperfections that are often present with a multi-agent system. In contrast, our social arguments captured in Section 3.1 allow agents to argue about their social influences and overcome such conflicts within a society.

The work of Sierra *et al.* [22] is an important initial attempt to extend the work of Kraus *et al.* to a social context. Similar to our approach (and unlike [21]) they allow agents to argue in social contexts with imperfect information. However, they only consider authority based relationships, which we believe only capture a specialised form of social context (i.e., institutions or formal organisations). Our work, on the other hand, presents a more generic way of capturing social influences of roles and relationships (i.e., using social commitment with different degrees of influence). This not only provides a simple unified mechanism to simulate social contexts with a wide array of relationships exerting different social influences upon the agents, but also allows us to experiment with our agents' ability to argue, negotiate and resolve conflicts in such disparate social systems.

6 Conclusions and Future Work

The long term objective of our work is to formulate an agent society that can use argumentative dialogues to resolve their conflicts. As mentioned in Section 1, we believe this requires agents to have four fundamental capabilities. To this end, this paper addresses the first two issues; namely a schema to reason in a social context and a mechanism to identify a suitable set of arguments. We achieve these aims by (i) developing a coherent schema for agents to function among different social influences and (ii) designing a model that allows agents to systematically use this schema to capture social arguments to negotiate and resolve conflicts within a social context. We also highlighted the two main ways in which social influence and ABN mutually enhance one another in terms of effectively resolving conflicts and demonstrated their operation in an illustrative case study.

In addition to the above two issues, we believe that agents also require a language and a protocol to exchange these arguments, and a decision making functionality to generate such dialogues [15]. As mentioned in Section 3.1, our domain language flows naturally from our schema and the communication language is adapted from the works of Amgoud *et al.* [12] and McBurney *et al.* [14]. In abstract, our protocol has six stages; namely opening, conflict recognition, conflict diagnosis, conflict management, agreement, and closing. Apart from the opening and the closing stages, which provide synchronisation points for the agents, the remaining four comply well with the pragma-dialectics model for critical discussion proposed by van Eemeren and Grootendorst [23]. Furthermore, for each locution type we have defined their respective pre-condition and commitment rules. Finally, we have defined the decision making functions for each of these dialogue moves first, at an abstract level and then in a more domain dependent level. However, since the main objective of this paper is to set the conceptual grounding (and also due to space restrictions) we choose to exclude detailed explanations

of these issues from this paper. In future, we aim to expand upon our current implementation by designing different argument selection strategies, thus allow the agents to adopt different tactics in resolving conflicts in a multi-agent community.

References

1. Karunatillake, N.C., Jennings, N.R.: Is it worth arguing? In: *Argumentation in Multi-Agent Systems (Proc. of ArgMAS 2004)*. LNAI 3366, NY, USA, Springer-Verlag (2004) 234–250
2. Rahwan, I., Ramchurn, S.D., Jennings, N.R., McBurney, P., Parsons, S., Sonenberg, L.: Argumentation-based negotiation. *The Knowledge Engineering Review* **18** (2003) 343–375
3. Cavedon, L., Sonenberg, L.: On social commitment, roles and preferred goals. In: *Proc. of the 3rd Int. Conf. on Multi-Agent Systems (ICMAS'98)*. (1998) 80–86
4. Panzarasa, P., Jennings, N.R., Norman, T.J.: Social mental shaping: Modelling the impact of sociality on the mental states of autonomous agents. *Computational Intelligence* **17** (2001) 738–782
5. Dignum, F., Morley, D., Sonenberg, E.A., Cavedon, L.: Towards socially sophisticated BDI agents. In: *Proc. of the 4th Int. Conf. on Multi-agent Systems*, Boston, USA (2000) 111–118
6. Castelfranchi, C.: Commitments: From individual intentions to groups and organizations. In: *Proc. of the 1st Int. Conf. on Multi-agent Systems (ICMAS'95)*, San Francisco, CA (1995) 41–48
7. Dignum, V., Kinny, D., Sonenberg, L.: Motivational attitudes of agents: On desires, obligations and norms. In: *Proc. of the 2nd Int. Workshop of Central Eastern Europe on Multi-Agent Systems (CEEMAS'01)*. Volume 2296., Poland (2001) 61–70
8. van der Torre, L., Tan, Y.H.: Contrary-to-duty reasoning with preference-based dyadic obligations. *Annals of Mathematics and Artificial Intelligence* **27** (1999) 49–78
9. Ross, A.: Imperatives and logic. *Theoria* **7** (1941) 53–71
10. Atkinson, K., Bench-Capon, T., McBurney, P.: A dialogue game protocol for multi-agent argument over proposals for action. In: *Argumentation in Multi-Agent Systems (Proc. of ArgMAS 2004)*. LNAI 3366, NY, USA, Springer-Verlag (2004) 149–161
11. Walton, D.N.: *Argumentation Schemes for Presumptive Reasoning*. Erlbaum, Mahwah, NJ (1996)
12. Amgoud, L., Parson, S., Maudet, N.: Argument, dialogue and negotiation. In Horn, W., ed.: *Proc. of the 14th European Conference on Artificial Intelligence (ECAI'00)*, Berlin (2000) 338–342
13. MacKenzie, J.: Question-begging in non-cumulative systems. *Journal of philosophical logic* **8** (1979) 117–133
14. McBurney, P., van Eijk, R., Parsons, S., Amgoud, L.: A dialogue-game protocol for agent purchase negotiations. *Autonomous Agents and Multi-Agent Systems* **7** (2003) 235–273
15. Karunatillake, N.C., Jennings, N.R., Rahwan, I., Norman, T.J.: Arguing and negotiating in the presence of social influences. In: *Proc. of the 4th Int. Central and Eastern European Conference on Multi-Agent Systems (CEEMAS'05)*, Budapest, Hungary (2005) to appear
16. Faratin, P., Sierra, C., Jennings, N.R.: Using similarity criteria to make trade-offs in automated negotiations. *Artificial Intelligence* **142** (2002) 205–237
17. Reed, C.A.: Dialogue frames in agent communication. In: *Proc. of the 3rd Int. Conf. on Multi-Agent Systems*. (1998) 246–253
18. Walton, D.N., Krabbe, E.C.W.: *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. State Univ. of NY (1995)
19. Reed, C.A., Walton, D.N.: Towards a formal and implemented model of argumentation schemes in agent communication. In: *Argumentation in Multi-Agent Systems (Proc. of ArgMAS 2004)*. LNAI 3366, NY, USA, Springer-Verlag (2004) 19–30
20. Parsons, S., Sierra, C., Jennings, N.R.: Agents that reason and negotiate by arguing. *Journal of Logic and Computation* **8** (1998) 261–292
21. Kraus, S., Sycara, K., Evenchik, A.: Reaching agreements through argumentation. *Artificial Intelligence* **104** (1998) 1–69
22. Sierra, C., Jennings, N.R., Noriega, P., Parsons, S.: A framework for argumentation-based negotiation. In: *Proc. of 4th Int. Workshop on Agent Theories Architectures and Languages (ATAL'97)*, Rhode Island, USA (1998) 167–182
23. van Eemeren, F.H., Grootendorst, R.: *Argumentation, Communication, and Fallacies*. Lawrence Erlbaum Associates, Inc, Hillsdale NJ (1992)