

Distributed Negotiation for Collective Decision-making

Ndeye Arame Diago*, Sarvapali Ramchurn †, Samir Aknine*, Onn Shehory‡ and Mbaye Sene§

*University of Lyon 1, Email: ndeye-aramediago@etu.univ-lyon1.fr

†University of Southampton, Email: sdr1@soton.ac.uk

‡Bar Ilan University, Email: onn.shehory@gmail.com

§University of Dakar, Email: mbaye.sene@ucad.sn

Abstract—Collective decision making is a process in which participants make a collective choice from several alternatives. In this paper, we focus on collective decision contexts in which more than two selfish agents negotiate over multiple issues. We specifically consider a case of joint household energy purchase where the concerned households have to define a collective energy contract. The households involved may each be interested only in a subset of the issues at stake. We devise an effective protocol to regulate the interactions among the (household) agents and reduce their reasoning complexity. The mechanism we introduce is fully decentralized, it facilitates multi-lateral negotiation, and it reduces the complexity of the solution despite the inherent complexity of the problem.

I. INTRODUCTION

Negotiation is a process that allows participants to reach agreements. Self-interested negotiating agents may have conflicting interests. Hence, each may try to direct the negotiation towards its preferred solutions. Negotiation complexity intensifies when the number of the negotiating agents is large, the number of negotiated issues is large, and utility functions are private information. In such contexts, reaching an acceptable agreement is challenging. This calls for novel interaction models that facilitate negotiation convergence and provides an efficient and fair solution concept. In this paper, we address these challenges in a context where agents' interaction is fully decentralized (with no mediator) and simultaneous. We focus on organizational aspects of the multi-agent system, on agent interaction and decision making, and on the efficiency and fairness of the solution. We allocate much attention to the multi-agent organization, as it may significantly affect complexity, flexibility and reactivity, and impose computation and communication overheads [17] [16].

Work in game theory has looked at modelling negotiation processes for analyzing agents' behaviors within, and the result of their interaction. In [21], the authors develop a framework for Boolean games with incomplete information, using Generalized Possibilistic Logic to compactly represent agents' knowledge about the preferences of others. In the negotiation protocol proposed there agents interact sequentially, the ordering of the sequence is important, and constraints are placed on timing and targeting of proposals. In contrast, in our protocol, agents interact simultaneously

and may send proposals to any agent at any time. Another approach to tackling multilateral negotiation problems is based on heuristic search in the solution space, assuming complete information [18]. In our work we assume incomplete information and we do not use search as a solution approach. Further, we address multi-issue negotiation, in difference from prior art. Work in [8] addresses negotiation over multiple issues. They analyze computational complexity of the negotiation when issues are discussed independently, jointly and sequentially. However, they focus on bilateral and not on multilateral negotiation as we do. Another approach to multilateral negotiation is to designate a mediator. Research in [3] and [6] proposes protocols for complex contracts where there is a mediator who generates the contracts, while we seek a fully decentralized solution. Other studies concern multilateral negotiation protocols with argumentation and persuasion. In [2], the authors propose a protocol for multi-party argumentation in which agents are equipped with argumentation systems and focus on a single issue. Work in [1] focuses on the target set (of moves) that specifies minimal changes in the current state of the debate that would allow agents to satisfy their goals. It experimentally investigates how well strategies based on target sets behave. In [23], the authors model persuasion in the context of social choice based on Borda voting rule. Their approach is centralized; they focus on one agent that influences the group's decision by convincing some agents to change their vote, while in our approach interaction is fully decentralized and each agent tries to convince the others to accept its proposal.

Typically, state of the art negotiation mechanisms do not focus on the organizational aspects of the negotiating agents. Commonly, agents are free to send and receive offers from all others, unless infrastructural constraints prohibit that. When many issues are negotiated and agents have different preferences over these issues, free (typically broadcast) communication leads to increased complexity in agent reasoning and in computational costs. This could negatively affect negotiation time, convergence and the quality of the solution. To address the problems stated above, we propose a novel negotiation model based on an iterative and incremental scheme that not only fully distributes the negotiation but additionally reduces complexity. We propose a multi-agent

organization that gradually evolves during the negotiation process. The underlying approach is based on the divide and conquer concept. This consists of dividing the agents into several groups within which they debate. The groups are formed based on the similarity of agents' choices over the issues they wish to negotiate. Negotiation takes place in each group by focusing on a specific subset of issues. Thus, we seek reduction in the computational complexity of agents' reasoning. To increase protocol flexibility, we allow communication between agents belonging to different groups via "overhearing". This allows agents to obtain information beyond that acquired during their internal group exchanges. Nevertheless, agents within a group submit proposals relevant to the issues of interest to that group. However, a final outcome of the negotiation must address all issues of interest to the agents across groups and satisfy some social choice rules detailed further. Hence, our protocol allows the agents to reach agreements that address these issues. We denote such proposals as complete proposals. In our protocol, proposals submitted in different groups are refined and merged, eventually evolving into complete proposals.

The contributions of this paper include the following: (1) a distributed multilateral negotiation protocol with multiple issues; (2) strategies to be used by agents for decision making and to defend their proposals within the protocol. As part of our study, we utilize an argumentation framework that enables agents to support their proposals by arguments.

II. NEGOTIATION MODEL

Scenario: Assume a set of household agents that seek a joint energy supply contract to reduce their energy expenditure. They must negotiate different aspects of the energy contract such as the type of the product (electricity or gas), provider, duration of the contract, tariff (peak and off-peak hours, fixed price, variable price), services included, payment method, termination conditions, etc. In practice, energy consumers focus only on a subset of the issues above, depending on their needs and preferences.

Model: In our study, agents negotiate to arrive at a collective agreement of the sort exemplified above. Our negotiation mechanism entails dividing up the agents into groups according to their choice of issues to be discussed. Agents in each group negotiate to reach an agreement over shared issues and form an alliance. An alliance allows agents to join forces to support their joint proposal over time. In our protocol, a proposal is considered complete if it addresses all issues at stake; otherwise it is a partial proposal. A complete proposal that is supported by the majority is an acceptable collective agreement. As negotiation progresses, alliances are formed. Alliances are then gradually combined to generate acceptable collective agreements. Such combination may result from simple merges of proposals when alliances have no issues in common. Otherwise, the combination requires negotiation over conflicting attributes. This combination and

refinement process proceeds until at least one alliance arrives at a proposal that addresses all of the issues, and is accepted by a majority of the agents, or until the deadline of the negotiation is reached. The motivation for forming alliances and gradually combining them is to structure the negotiation, reduce its complexity, and facilitate its convergence. The negotiation may consist of several rounds separated by periodic, pre-defined checkpoints, and it has a deadline denoted d_l .

Our negotiation mechanism consists of three major phases. In phase 1 – assignment – the agents form groups according to similarities in issues they wish to discuss. In phase 2 – negotiation – the agents formulate proposals, send and receive proposals, and listen to other exchanges. Agents reason about and react to proposals. When beneficial they may migrate to other groups. An agent that migrates to another group may need to change its choices of issues. A set of agents that agree on the same proposal form an alliance and they encounter other alliances to merge or refine their proposals. In phase 3 – checkpoints – the agents check whether there are proposals that can be valid solutions of the negotiation. Here, the goal is to identify a winner alliance. In case of valid solutions (winner alliance), a selection process is applied and the negotiation terminates with the selected solution. Otherwise, another negotiation round begins.

III. FORMALIZATION

Agents and groups of agents: let $\mathcal{A}=\{a_1,\dots,a_n\}$ be the set of agents which negotiate over multiple issues $E=(e_1,\dots,e_k)$ with k the number of attributes. Note that here, and across this study, issues and attributes are not different. We use the term issue when referring to what agents care about, and the term attribute to refer to issues a proposal, or a solution, address. It is necessary to specify the set of attributes ahead of the negotiation. In our case, this set is defined by a protocol administrator ahead of the negotiation. The rules of the protocol dictate that agents can negotiate only over pre-specified attributes. Since the number of attributes may be large (as in our example scenario), agents can have diverse choices over attributes they wish to discuss. We denote by E_{a_i} the subset of issues selected by agent a_i . Let $\mathcal{G}=\{g_1,\dots,g_m\}$ be the set of groups formed at the protocol start. Each group g_k is a tuple $(\mathcal{A}_{g_k}, E_{g_k}, \mathcal{P}_{g_k})$ with \mathcal{A}_{g_k} the set of agents in group g_k , E_{g_k} is the set of attributes to negotiate by the agents $a_i \in \mathcal{A}_{g_k}$, \mathcal{P}_{g_k} the set of proposals submitted within group g_k .

Proposals and utility function: a proposal p_α is a tuple $(v(e_1), \dots, v(e_k))$ whose elements are the values of the attributes (as suggested by the proposing agent). When generating a proposal, an agent associates values only with attributes discussed within its group. Agents evaluate proposals according to a set of criteria $\mathcal{X}=\{x_{(1)}, \dots, x_{(r)}\}$. Each agent a_i selects, according to its objectives, a subset of the criteria set $\mathcal{X}_i \subseteq \mathcal{X}$ on which it defines its own utility function. Each evaluation criterion $x_{(j)}$ involves a subset of attributes $E_{x_{(j)}}$.

For example, an energy provider could be evaluated by its reliability, its services or on whether it provides green energy; the tariff could be assessed by its variability, etc.

Let $u_{a_i}: \mathcal{P} \rightarrow [0,1]$ be agent a_i 's utility function, where \mathcal{P} is the set of submitted proposals. This is used to establish a preference order over the set of alternatives. The utility functions of the agents are private information. Before negotiation, each agent a_i sets its desired utility range: $U_{min a_i}$ is the minimum utility value the agent may agree to, and $U_{max a_i}$ is its desired utility value, above which any proposal is acceptable. Utility ranges differ across agents.

Let $\mathcal{O} = \{\text{propose}, \text{accept}, \text{refuse}, \text{reinforce}, \text{attack}\}$ be the set of speech acts agents use. Let p_α^i be a proposal p_α submitted by agent a_i . Formally:

Propose($a_i, Ag_k, p_\alpha^i, argp_\alpha^{i+}$): $a_i \in Ag_k$ submits p_α^i with a set of positive arguments $argp_\alpha^{i+} \subseteq Argp_\alpha^{i+}$ to $\forall a_j \in Ag_k$.

Accept(a_i, Ag_k, p_α^r): $a_i \in \mathcal{A}$ sends an acceptance message for proposal p_α^r submitted by agent $a_r \in Ag_k$ to $\forall a_j \in Ag_k$.

Refuse(a_i, Ag_k, p_α^r): $a_i \in \mathcal{A}$ sends a refusal message for proposal p_α^r submitted by agent $a_r \in Ag_k$ to $\forall a_j \in Ag_k$.

Attack($a_i, Ag_k, p_\alpha^r, argp_\alpha^{r-}$): $a_i \in \mathcal{A}$ sends an attack message for proposal p_α^r submitted by agent $a_r \in Ag_k$ to $\forall a_j \in Ag_k$ with a set of negative arguments $argp_\alpha^{r-} \subseteq Argp_\alpha^{r-}$.

Reinforce($a_i, Ag_k, p_\alpha^r, argp_\alpha^{r+}$): $a_i \in \mathcal{A}$ sends a reinforcing message for proposal p_α^r submitted by $a_r \in Ag_k$ to $\forall a_j \in Ag_k$, with a set of positive arguments $argp_\alpha^{r+} \subseteq Argp_\alpha^{r+}$.

Each submitted proposal p_α is associated with a tuple $(\mathcal{T}_{p_\alpha}^{ac}, \mathcal{T}_{p_\alpha}^{re}, \mathcal{T}_{p_\alpha}^{rf}, \mathcal{T}_{p_\alpha}^{at}, \nu_{p_\alpha}, ws(p_\alpha), \epsilon_{p_\alpha})$ whose elements are, respectively, its acceptance, reinforcing, refusal and attack rates, its support ratio value with $\nu_{p_\alpha} = \frac{\mathcal{T}_{p_\alpha}^{ac} + \mathcal{T}_{p_\alpha}^{re}}{\mathcal{T}_{p_\alpha}^{rf} + \mathcal{T}_{p_\alpha}^{at} + 1}$, its social satisfaction value and its time stamp. The rates are computed w.r.t. the number of agents in the system. For example $\mathcal{T}_{p_\alpha}^{ac}$ is the ratio between the number of acceptances and the number of agents n . Other rates are calculated similarly.

IV. NEGOTIATION PROTOCOL

A. Assignment phase

We assume that each agent has its own preferences over the negotiated issues. However, some agents could make common choices over some issues. In our example, let $\mathcal{A} = \{a_1, a_2, a_3, a_4, a_5\}$ be the households that aim to negotiate an energy contract involving attributes e_1, e_2, e_3, e_4, e_5 which are, respectively, *providers, products, tariff, duration of the contract and payment method*. Consider agents' issues of interest as follows: $E_{a_1} = (e_1, e_2, e_3), E_{a_2} = (e_2, e_3), E_{a_3} = (e_1, e_5), E_{a_4} = (e_1, e_2), E_{a_5} = (e_1, e_4, e_5)$. Agents form negotiation groups based on preference similarity. That is, they join groups in which their preferences over attributes are similar to those of other agents. This is done in a decentralized manner, based on individual agent decisions. Note that agents do not share their exact preferences, but their "issues of interest" are

disclosed. Our protocol entails that an agent cannot belong simultaneously to more one group.

Let $\gamma_i = |E_{a_i}|$ be the number of preferred attributes for a_i . Agent a_i sorts the agents according to the number of attributes they share. So it generates the following subsets: $O_i = O_{\gamma_i}, O_{\gamma_i-1}, O_{\gamma_i-2}, \dots, O_0$ which, respectively, represent the set of agents with which it shares the same attributes, the set of agents with which it shares γ_i-1 attributes, the set of agents with which it shares γ_i-2 attributes and so on. As an exemple, agent a_1 holds $\gamma_1=3$ attributes and $O_1 = \{O_3 = \emptyset, O_2 = \{a_2, a_4\}, O_1 = \{a_3, a_5\}, O_0 = \emptyset\}$.

Agent groups are formed in rounds. Firstly, each agent a_i requests the agents with which it shares the highest number of attributes to form a group. a_j accepts a_i 's request if it also shares with a_i the highest number of attributes. Once an agent accepts a request, it cannot accept again another request. But it can proceed and request other agents to join its group. After this phase, some agents may have not joined any group. For these agents, the process proceeds to another round in which they make group requests based one their second choice. This iterates until all agents join groups.

B. Negotiation phase

Negotiation takes place within each group, across groups according to the inter-group communication policy, and eventually between alliances. It occurs in several phases separated by checkpoints. The latter consist of verifying whether there is a proposal that meets the requirements of a solution (see Section V). Negotiation ends when an agreement is found during a checkpoint phase or when the deadline is reached.

1) *Inside a group*: Each agent submits its proposals within its group following diverse strategies: it can send all of its proposals at once, or incrementally. The incremental approach consists of submitting the best proposal first. It requires a strategy for switching to the next proposal in the agent's proposal list. Firstly, the utility value of this proposal needs to be higher than the expected minimum utility U_{min} of the agent. Additionally, if within a certain period of time δ_t the proposal does not meet a certain rate of acceptance (by other group members), the agent can move to its second-best proposal, and so on. While sending proposals to others and negotiating them, an agent also receives proposals from other agents and assesses them. It can accept or refuse those proposals. The motivation for partitioning the agents into groups is to simplify and reduce complexity of the interaction among the agents. To facilitate that, we introduce some interaction rules \mathbf{R}_i that limit the communication complexity, as follows:

\mathbf{R}_{i1} : An agent can submit a limited number of proposals δ_p and it can make a limited number of refusals δ_r .

\mathbf{R}_{i2} : An agent can submit its proposals only to the members of the same group at any time of the negotiation phase.

\mathbf{R}_{i3} : An agent can accept a proposal made by another agent

at any time of the negotiation phase even if it had refused or attacked this proposal before.

R_{i4}: An agent can accept and reinforce a proposal made by another agent at any time of the negotiation phase even if it had not refused or attacked this proposal before.

R_{i5}: An agent can refuse and attack a proposal made by another agent at any time of the negotiation phase if it had not accepted it before.

R_{i6}: A proposal can not be accepted, refused, attacked or reinforced several times by the same agent. These rules result in 5 possible decision sequences:

Accept, Refuse, Accept → *Reinforce, Refuse* → *Attack, Refuse* → *Attack* → *Accept*. The interaction rules streamline the negotiation and limit complexity. They also enable an agent to make concessions by accepting, e.g., a proposal that it refused before. An agent's concession consists of scaling down its desired utility range.

2) *Inter-group communication*: An inter-group communication policy defines: when the agents become aware of the exchanges carried out outside their group; when during the negotiation they can express their viewpoints about the perceived proposals; and when during the negotiation they can move to another group. There are multiple ways to define an inter-group communication policy. In this study, we focus on a policy that allows communication across groups named Free Inter-group Communication (FIC). With FIC, agents have a visibility over proposals submitted outside their group, and they can react upon these proposals subject to the interaction rules described below.

R_{i7}: the members of a group cannot submit a proposal to agents in another group until they move into that group.

The partition of agents into separate groups aims for the interaction within each group to take place among agents with similar attributes they want to discuss. We do not prohibit interaction with other agents as discussed above, however we assume that such interaction will be limited as agents will have limited interest in other groups.

R_{i8}: a_i may accept a proposal which addresses some of its issues of interest and submitted outside its group while staying within its own group.

When an agent is not satisfied by its group's proposals, it may look for another group with which it shares some issues. It may either accept proposals submitted in that group, or migrate to that group to submit its proposals there. Once an agent migrates to a new group, it is immediately granted all interaction rights within that group, as dictated by the protocol. Agent migration is regulated with mobility rules (**R_m**) as follows:

R_{m1}: a_i cannot visit the same group more than δ_v times.

R_{m2}: a_i cannot migrate more than δ_d times. These rules channel the interactions of agents by forcing them to stay a certain time within a group. Subject to these rules, an agent strategically chooses the group it wishes to join.

The FIC interaction policy increases the flexibility of

inter-group agent communication. Its advantage is the ability to arrive at better agreements, and the disadvantage is an increased communication cost. In broadcast communication, agents receive every proposal submitted even if they do not focus on its issues of interest. This involves unnecessary messages processings and unnecessary communications. With FIC, the agent decides by itself to acquire information from other groups when this is a beneficial. We assert that agent decisions to join or change groups, and to change attributes, are not more complex than having all the agents negotiate in one big group.

3) *Decision making*: Agents' decisions are made based on their utility from the proposals at hand. Ahead of the negotiation process, an agent a_i (or more likely its owner or administrator) sets its desired utility range: $Umin_{a_i}$ is the minimum utility value the agent may agree to, and $Umax_{a_i}$ is its desired utility value, above which any proposal is acceptable. Our negotiation problem assumes different utility functions. Hence, we introduce an alternative way for representing the agent's satisfaction level with a proposal. The utility thresholds of each agent split its utility line into three ranges. We associate these ranges with agent satisfaction. The ranges $[0, Umin]$, $[Umin, Umax]$, $[Umax, \infty]$ are associated with satisfaction levels $\sigma_0, \sigma_1, \sigma_2$, respectively. Note that agents may have different utility thresholds. Consequently, two agents whose utility from a proposal is the same may have different satisfaction levels from the same proposal. From the above, given a proposal p^j_β from a_j to a_i and the utility from it $u_{a_i}(p^j_\beta)$, if $u_{a_i}(p^j_\beta) < Umin_{a_i}$ then it is in the best interest for a_i to refuse and attack p^j_β , as it has a low, σ_0 satisfaction value. Similar strategy interpretations can be made for σ_1 and σ_2 . That is, σ_1 is associated with acceptance at a moderate satisfaction level, and σ_2 is associated with a high satisfaction level. In the rest of this paper, we set the value of $\sigma_0=0$.

4) *Alliances*: An alliance is formed by at least two agents when they agree on one proposal. The goal is for agents to join forces in order to better support their proposal. In this work, joining forces entails that agents within an alliance define and implement common decision-making strategies for negotiating with alliances and specifically the utility range of the alliance. This is computed based on the members' expected utilities. Alternative strategy definitions are possible too but are not discussed in this study [24]. An alliance c_x is represented by a tuple $(A_c, Umin_c, Umax_c, p_c)$. It sets its desired utility range such that $Umin_c = \min\{Umin_{a_i}\}_{a_i \in A_c}$ is the minimum utility value each agent in the alliance may agree to, and $Umax_c = \min\{Umax_{a_i}\}_{a_i \in A_c}$ is its desired utility value, above which any proposal is acceptable. Thus, we aim to preserve the expected utility of each agent. Each agent evaluates a proposal according to its utility function and makes decisions according to the desired utility range of its alliance. This means that an agent refuses and attacks a proposal if its utility is lower than $Umin_c$, otherwise it

accepts and reinforces it. In this way, the agents are led to make concessions by accepting proposals which would have been refused if they considered their own desired utility range. This concerns every agent a_i in the alliance whose minimum expected utility $U_{min_{a_i}}$ is more than U_{min_c} . However, an agent's individual satisfaction value from a proposal depends on its own desired utility range. Thus, when a_i accepts a proposal p_α such that $U_{min_c} < u_{a_i}(p_\alpha) < U_{min_{a_i}}$ or $U_{max_c} < u_{a_i}(p_\alpha) < U_{max_{a_i}}$, or $U_{min_{a_i}} < u_{a_i}(p_\alpha) < U_{max_{a_i}}$ its satisfaction value is σ_1 , otherwise its satisfaction value is σ_2 . The protocol motivates the agents to incrementally make concessions. This formation mechanism facilitates convergence of the negotiation.

During the negotiation, different alliances can form. Each aims to get a maximal number of agents with a proposal addressing all of the issues. To this end, it has to convince other agents to accept the proposal it defends, or convince other alliances to join forces, either by merging proposals or by negotiating over conflicting attributes to refine and generate a joint proposal. An agent can join an alliance at any time by accepting the proposal defended by this alliance (rule \mathbf{R}_{i5}). An agent cannot leave its alliance because it cannot remove proposals it has already accepted. But it can join another alliance by accepting another proposal (rule \mathbf{R}_{i3}).

5) *Interaction across alliances*: According to the structure of the initial groups of agents and the interaction rules, there may exist different relations between the alliances:

link_A: since we allow agents to accept several proposals (\mathbf{R}_{i3}), it could happen that two alliances c_{x,g_k}, c_{y,g_k} share some agents, i.e., $c_{x,g_k} \cap c_{y,g_k} \neq \emptyset$ (see Fig. 1).

link_E: according to the group creation policy we use, it could happen that two alliances c_{x,g_k}, c_{y,g_q} share some attributes, i.e., $E_{g_k} \cap E_{g_q} \neq \emptyset$ (see Fig. 1).

link_{FIC}: according to \mathbf{R}_{i8} rule, it could happen that two alliances c_{x,g_k}, c_{y,g_q} from different groups g_k, g_q share some agents, i.e., $c_{x,g_k} \cap c_{y,g_q} \neq \emptyset$ (see Fig. 1).

An alliance uses several criteria to choose another alliance




Figure 1. Structure of the groups and alliances

outside its group in order to merge or refine their proposals. These criteria are the number of agents in the other alliance, the number of attributes it shares with that alliance and the number of agents it shares with that alliance. Alliances that have formed are publicly known. Each alliance decides by itself the alliances with which it would like to join forces.

We prevent an alliance from choosing another from the same group since their proposals have already been discussed.

Merging and refinement operations are denoted, respectively, by f_m and f_r . To perform these, at least two proposals from two alliances c_x, c_y are needed. Alliances decide upon merging or refining proposals based on the gain they would achieve and the cost of the operation. The cost and the gain are given, respectively by f_{cost} and f_{gain} functions defined as follows: $f_{cost}(c_x, c_y) = \frac{|E_{c_x} \cap E_{c_y}|}{|E_{c_x} \cup E_{c_y}|} \times \frac{|A_{c_x} \cup A_{c_y}| - |A_{c_x} \cap A_{c_y}|}{|A_{c_x} \cup A_{c_y}|}$. $f_{gain}(c_x, c_y) = \frac{|E_{c_x} \cup E_{c_y}|}{|A_{c_x} \cup A_{c_y}|} \times \frac{|A_{c_x} \cup A_{c_y}|}{|A_{c_x} \cup A_{c_y}|}$, n, k represent, respectively, the number of agents in the system and the number of all issues at stake. There are other approaches to define these criteria [25]. Here we define cost and gain functions tailored to our context. The functions can be generalised to many alliances.

Property 1: The cost of a merging process is equal to zero.

This makes sense because there are no attributes over which agents are in conflict, i.e., $E_{c_x} \cap E_{c_y} = \emptyset \Leftrightarrow |E_{c_x} \cap E_{c_y}| = 0 \Rightarrow f_{cost}(c_x, c_y) = 0$. We define the expected utility function of an alliance c_x from joining forces with another c_y by $u(c_x, c_y) = f_{gain}(c_x, c_y) - f_{cost}(c_x, c_y)$. Each alliance chooses the alliances that maximise its utility.

6) *Merging and refinement process*: Two proposals can be merged when there are no common attributes. Merging entails creating a new proposal with the attributes of the merged proposals. In Figure 1, we suppose that alliances c_{1,g_1} and c_{4,g_3} are going to support, respectively, the proposals $p_1 = (v(e_1), v(e_2))$ and $p_2 = (v(e_3), v(e_4))$, where $v(e_i)$ is the value of attribute e_i . To merge their proposals, they should define a proposal $p_x = f_m(p_1, p_2) = (v(e_1), v(e_2), v(e_3), v(e_4))$ with f_m the merging function. The merging process is simple because we assume that attributes are independent and each alliance is not concerned with attributes from the other alliance. It is beneficial for the alliances to merge their proposals because it allows them to better defend their proposal by becoming an alliance with more agents. However, to become the winner alliance, the defended proposals must be complete. A complete proposal includes all attributes discussed during the negotiation. It should additionally meet the requirements of a solution, as in Section V. Hence, each alliance strategically chooses the alliances with which it forms a group.

The refinement process of the proposals supported by the alliances consists of negotiating over common attributes whose values are different. This negotiation concerns only the values of the attributes in conflict and consists of finding a compromise. Each agent in these alliances can submit a proposal by modifying only the values of these attributes. For instance, when two alliances having, respectively, the proposals $p_1 = (v(e_1), v(e_3), v(e_4))$, $p_2 = (v(e_2), v(e_3), v(e_5))$ form a group of alliances g_1^c , a new proposal in g_1^c is a proposal $p_x = f_r(p_1, p_2) = (v(e_1), v(e_2), set(v(e_3)))$,

$v(e_4), v(e_5)$) with f_r the refinement function and $set(v(e_3))$ is a new e_3 value. There are many ways to resolve the conflict. This is out of scope of this work.

C. Checkpoint phase

In our mechanism, negotiation takes place in rounds, separated by checkpoints. Periodical checkpoints allow checking the progress of the negotiation by analyzing the results of the exchanges performed in all groups. This consists of verifying whether there is a complete proposal that is supported by the majority, i.e., a winner alliance. Checkpoints eventually allow to conclude whether the negotiation has arrived at a solution, or whether it failed. Two results are possible at the end of a checkpoint: either there is a solution, in which case the negotiation ends in success, or there is not, in which case it continuous if the deadline is not reached. If at the last checkpoint no solution is found the negotiation ends in failure. Before the negotiation, we set the number of checkpoints q , and their schedule. This information is shared with all agents. Let $\mathcal{C}_{pt} = \{C_{pt_1}, \dots, C_{pt_q}\}$ be the set of checkpoints. Each checkpoint C_{pt_r} is a tuple (t_r, S_r) , t_r is the time of the checkpoint and S_r the set of proposals acceptable as a solution. $\mathcal{C}_{pt} = \{C_{pt_1} = (t_1, S_1), \dots, C_{pt_q} = (t_q, S_q)\}$ is ordered chronologically ($t_1 < \dots, t_q < d_t$). We arbitrarily set the majority threshold $\phi_{maj} = 50\%$. Any proposal p_α with acceptance rate $\mathcal{T}_{p_\alpha}^{ac} > \phi_{maj}$ is an acceptable solution. When checkpoint C_{pt_r} occurs, all $a_i \in \mathcal{A}$ add their acceptable proposals to S_r . If $S_r = \emptyset$ then there is no solution at that checkpoint. If $r \neq q$, negotiation proceeds to next checkpoint $C_{pt_{r+1}}$. If $|S_r| = 1$ then there is one solution and the negotiation ends in success. If $|S_r| > 1$ the negotiation ends in success and the solution is selected according to the social decision process detailed below. If $r = q$ and $S_r = \emptyset$ the negotiation ends in failure.

V. NEGOTIATION OUTCOME

In classical score voting, the score of each candidate is the sum of obtained scores. The winner is the candidate with the highest score. But it could happen that this candidate is not the winner when we simply consider the majority rule. Let A, B, C a set which must decide between two candidates x, y with score 0,1,2. They get, respectively, the results (1,1,1) and (2,2,0). With majority rule, x is the winner by considering score 0 means a missing vote. With a score voting, y is the winner with a score equal to 4. Nevertheless, each of these methods does not always ensure fairness [26]. To overcome this, we propose a new solution concept that combines effectiveness and fairness [11][10][26]. The concept of fairness we use here draws on [11]. It consists of deciding an outcome which satisfies at the same time the majority rule and the score voting. By doing so we introduce a new method to evaluate the social satisfaction of an outcome which is not limited to the sum of agents' satisfactions as done for x in the example stated above.

On a concrete level, we use a mixed voting system based on approval voting [15] and score voting in order to combine the advantages and to limit the drawbacks of those methods. According to the resulting method, each agent may accept any number of proposals it prefers as an outcome of the negotiation (and refuse the others). An agent allocates a score to each proposal. A proposal score is the agent's satisfaction from it, i.e., σ_0, σ_1 or σ_2 , (cf Section IV-B3). To determine the outcome of the negotiation, we propose a social decision process based on a set of hierarchical decision criteria in order to enable selection of a single preferred proposal when there are several acceptable proposals. It consists of establishing a social preference order [10] [11] over the set of acceptable proposals built by the agents during the negotiation. These criteria are the social satisfaction, the support value (section III) and the time stamp of the proposal. Note that using the acceptance rate as the only decision metric is insufficient. This is in particularly the case when there are several proposals with the same acceptance rate. We introduce a new score aggregation function named social satisfaction w_s (defined below) which satisfies the rules of majority and score voting.

Social satisfaction function: We represent the evolution of a submitted proposal during the negotiation process by a set of successive states $S_{p_\alpha} = \{s_0, s_1, \dots, s_q\}$. A proposal moves from one state to another whenever an agent sends an acceptance, a reinforcing, an attack or a refusal message. The first state s_0 corresponds to the submission stage of the proposal. For each proposal p_α , its state set S_{p_α} corresponds to a satisfaction profile $\sigma_{p_\alpha} = (\sigma_{s_0}, \sigma_{s_1}, \dots, \sigma_{s_q})$, $\sigma_{s_i} \in \{\sigma_0, \sigma_1, \sigma_2\}$ which represents the satisfaction level of the agent of which action moved the proposal from state s_{i-1} to state s_i , after the action was performed (i.e., the proposal moved to state s_i). The social satisfaction function $w_s : \mathcal{P} \rightarrow \mathbb{R}_+$ assigns a positive real value to each submitted proposal, designating its social satisfaction score. To facilitate fairness, w_s draws on Rawls' difference principle (or maximin) [11], which regulates inequalities as follows: it only permits inequalities that work to the advantage of the worst-off. The value of w_s depends on the way in which it evolves throughout the negotiation. Given the above, the social satisfaction

function is defined as follows: $w_s(p_\alpha) = \frac{\sum_{i=0}^q \sigma_{s_i}}{n - n_{ap_\alpha}(s_q)}$, where n denotes the number of agents, $n_{ap_\alpha}(s_q)$ denotes the number of agents that accepted the proposal p_α in state s_q , $n_{ap_\alpha}(s_q) \leq n - 1$.

Proposition 1: $w_s(p_\alpha, s_i)$ is positive and increasing in S_{p_α} .

Proof: In $s_0, w_s(p_\alpha, s_0) = \frac{\sigma_{s_0}}{n-1} > 0$, σ_{s_0} is the satisfaction of the agent that submitted p_α . The submission of a proposal is an implicit acceptance (by the submitter). Each $s_i \neq s_0$ matches an acceptance, a reinforcing, a refusal or, an attack. $w_s(p_\alpha, s_i) = \frac{\sum_{k=0}^i \sigma_{s_k}}{n - n_{ap_\alpha}(s_i)}$. When p_α switches from s_i to s_j : if s_j matches a re-

fusal, attack, or a reinforcing then $\sigma_{s_j}=0, n_{ap\alpha}(s_j)=n_{ap\alpha}(s_i)$
 $\sum_{k=0}^i \sigma_{s_k} = \sum_{k=0}^j \sigma_{s_k} r \Leftrightarrow w_s(p_\alpha, s_j) = w_s(p_\alpha, s_i)$; if s_j matches an ac-
 ceptance, $\sigma_{s_j} > 0 \frac{\sum_{k=0}^j \sigma_{s_k}}{n - n_{ap\alpha}(s_j)} = \frac{\sum_{k=0}^i \sigma_{s_k} + \sigma_{s_j}}{n - (n_{ap\alpha}(s_i) + 1)} > \frac{\sum_{k=0}^i \sigma_{s_k}}{n - n_{ap\alpha}(s_i)} \Leftrightarrow$
 $w_s(p_\alpha, s_j) > w_s(p_\alpha, s_i)$. w_s increases according to the accep-
 tances. It is maximized when p_α is accepted by all. ■

Proposition 2: The proposal with the highest value of so-
 cial satisfaction has also the highest acceptance rate (keeping
 σ_1 and σ_2 within some range limitation).

Proof: Let $f, g, f(k) = \frac{(k+1)\sigma_1}{n-k}, g(k) = \frac{(k+1)\sigma_2}{n-k}$ two functions,
 $k \in [1, n-1]$ and \mathcal{P}_k the set of proposals with k acceptances.

$$\begin{aligned} \forall p_\alpha \in \mathcal{P}_k, \forall \sigma_1 \in]0, \sigma_2[, f(k) \leq w_s(p_\alpha) \leq g(k). \\ \forall p_\alpha \in \mathcal{P}_k, p_\beta \in \mathcal{P}_{k+1}, \exists \sigma_1 \in]0, \sigma_2[: \\ f(k) \leq w_s(p_\alpha) \leq g(k) < f(k+1) \leq w_s(p_\beta) \leq g(k+1). \quad (1) \end{aligned}$$

$f(k+1) > g(k) \Leftrightarrow \frac{(k+2)}{n-k-1}\sigma_1 > \frac{(k+1)}{n-k}\sigma_2$. Let $z(k) = \frac{(k+1)}{n-k}\sigma_2$
 $\times \frac{n-k-1}{k+2}$: $\forall \sigma_2 > 0, \sigma_1 \in]M, \sigma_2[, f(k+1) > g(k)$, M is the maxi-
 mum value of $z(k)$ with $k \in [1, n-1]$.

In summary, proposition 3 is true $\forall \sigma_2 > 0, \sigma_1 \in]M, \sigma_2[$ ■

Preference order: The social preference relationship de-
 scribed by w_s is complete, as there are no incomparable
 proposals. It is reflexive and transitive too. w_s is used
 in order to establish a social preference order when there
 are several proposals with acceptance rate greater than or
 equal to ϕ_{maj} . This threshold should be greater than 50% to
 guarantee a majority. The limitation of w_s is that σ_1 and σ_2
 can not take any value, they are interdependent. When we
 set σ_2 , σ_1 must comprise in $\in]M, \sigma_2[$ in order to guarantee
 the expected result.

We use the support ratio value to differentiate between
 proposals when there are several proposals that have the
 same value of social satisfaction (i.e., social indifference).
 When there is more than one proposal with the same support
 ratio value, we use the time stamp as a differentiator. We
 assume that there is a difference ϵ between the times in
 which two proposals got a majority.

We establish the preference order process described above
 to lead the negotiation to a collective agreement which
 satisfies the majority. Electing a proposal as the outcome of
 the negotiation matches the choice of the winner alliance.

Property 2: The proposed social decision process guar-
 antees the uniqueness of the solution.

The time stamp is the last criterion used when there are
 several proposals with the same social satisfaction value and
 the same values. In this case, the proposal with the smallest
 time stamp is the solution.

VI. NEGOTIATION COMPLEXITY

We define $D_c \in [0, 1]$ the complexity factor which expresses
 the similarity of agents (belonging to different groups)
 according to their choices of attributes they want to discuss.

Definition Let m the number of formed groups and m_c the
 number of groups $g_k \in \mathcal{G}$ such that $\exists g_q \in \mathcal{G}$ and $E_{g_k} \cap E_{g_q} \neq \emptyset$. The
 complexity factor $D_c = \frac{m_c}{m}$.

The complexity factor D_c expresses the complexity of the
 negotiation process and the alliances' interaction. According
 to D_c , we identify three situations:

- $D_c = 0$ when all of the initial groups formed are disjoint.
 This means that $\forall g_k \in \mathcal{G}, \nexists g_q \in \mathcal{G} / E_{g_k} \cap E_{g_q} \neq \emptyset$ so $m_c = 0$. Hence,
 alliances from different groups have no common attributes.
 They only proceed to merge their proposals. When $D_c = 0$ the
 negotiation is less complex compared to the case of $D_c > 0$.
 Here, each alliance chooses another to merge their proposals
 according to the size of the alliances (the number of agents
 in the alliances).

- $0 < D_c < 1$ when some groups are not disjoint. This means
 that $\exists g_k, g_q \in \mathcal{G} / E_{g_k} \cap E_{g_q} \neq \emptyset$ so $m_c > 0$. The alliances proceed to
 merge or to refine (negotiate the attributes in conflict) their
 proposals.

- $D_c = 1$ when each group shares at least some
 attributes with another group. This means that
 $\forall g_k \in \mathcal{G}, \exists g_q \in \mathcal{G} / E_{g_k} \cap E_{g_q} \neq \emptyset$ so $m_c = m$. The alliances proceed
 only to refine (negotiate the attributes in conflict) their
 proposals.

VII. EXPERIMENTAL RESULTS

To analyze the performance of our protocol, we have imple-
 mented an example of negotiation scenario in the energy
 market. We have tested our protocol by varying the number
 of agents n and the number of issues k . For each experiment,
 each agent chooses randomly its interesting attributes and
 seeks to form a group with the one with which it shares the
 maximum number of attributes. We define utility function
 used by an agent a_i to evaluate an a_j in order to request it
 to form a group as follows: $u_{g_i}(a_j) = |E_{a_i} \cap E_{a_j}|, u_{g_i} \in [0, \gamma_i]$. The
 experiment has been repeated 350 times.

Firstly, we have analyzed the number of initial groups
 m which are formed and the number of rounds r needed
 to form them. experimental results show that r increases
 when k increase, but it does not increase when n increases.
 This implies negotiation time depends on the number of
 issues but not on the number of agents. Thus our solution
 indeed removes the complexity associated with the size of
 the agent community, as intended by this study. Secondly,
 we have analyzed, the complexity factor defined in section
 VI according to the ratio between of agents and the number
 of issues. We denote by $r_i = n/k$. the curve of figure 2
 shows that for 16,6% of performed experiments, $D_c = 0$ with
 $r_i = 6/16 = 0.4$, for 54,6%, $0 < D_c < 1$ with $r_i = 34/13 = 2.6$ and for
 30%, $D_c = 1$ with $r_i = 18/17 = 1.1$. These results will be used as
 an index to guide agents to define their negotiation tactics
 according to the structure of the initial groups they will form.
 This is not detailed in this work.




Figure 2. Complexity factor

VIII. CONCLUSION

This paper has presented a new multilateral negotiation mechanism based on an incremental arrival at a solution. The mechanism is based on a divide and conquer approach. In our protocol agent interactions are fully decentralized. It respects individual preferences however maintains fairness as well. In our mechanism, each agent manages its interactions by following the rules of the protocol. Agents submit proposals, reason about them, and take action to promote their desired proposals and demote others. It leads to incremental and gradual convergence into a single solution which is both supported by the majority of agents, and fair. In future research, we intend to refine the protocol and formally prove its convergence to a solution.

REFERENCES

- [1] D. Kontarinis, E. Bonzon, N. Maudet, P. Moraitis, Empirical Evaluation of strategies for Multiparty Argumentative Debates, Workshop CLIMA XV Prague Czech Republic, (2014).
- [2] E. Bonzon, N. Maudet, On the Outcomes of Multiparty Persuasion, 8th International Workshop ArgMAS, (2011).
- [3] M. Klein, P. Faratin, H. Sayama, Y. Bar-Yam, Protocols for negotiating complex contracts, Intelligent Systems IEEE,(2003).
- [4] M. Wardeh, T. Bench-Capon, F. Coenen, Multi-Party Argument from Experience, 6th Workshop, ArgMAS (2009).
- [5] P. Frank, M. Dignum, Gerard A. W. Vreeswijk, Towards a Testbed for Multi-party Dialogues, International Workshop, ACL, pages 212-230 (2003).
- [6] R. Aydogan, K. V. Hindriks, C. M. Jonker, Multilateral Mediated Negotiation Protocols with Feedback, In Fifth ACAN Workshop, (2012).
- [7] U. Endriss, Monotonic concession protocols for multilateral negotiation, AAMAS, pages 392-399 (2006).
- [8] Shaheen S. Fatima, Michael Wooldridge, Nicholas R. Jennings, Multi-Issue Negotiation with Deadlines, JAIR, (2006)
- [9] A. Nongailard, P. Mathieu, B. Jaumard, La negociation du bien-etre social utilitaire, JFSMA'2008, pages 55-64(2008).
- [10] K. Arrow, A. Sen, k. Suzumura, Handbook of Social Choice and Welfare, Volume 2, Pages 1-962 (2011)
- [11] P. Adair, La Theorie de la justice de John Rawls. Contrat social versus utilitarisme. In: Revue francaise de science politique, pages 81-96 (1991).
- [12] B. Golden, P. Perny, Infinite order Lorenz dominance for fair multiagent optimization, AAMAS, pages 383-390 (2010).
- [13] W. Marshall, I. Olkin, Inequalities: theory of majorization and its applications, Academic Press, London (1979).
- [14] E. Posner, K. Spier, A. Vermeule, Divide and Conquer, Harvard Law School John M. Olin Center for Law, Economics and Business Discussion Paper Series. Paper 636 (2009).
- [15] S. Brams, P. Fishburn, Going from theory to practice: the mixed success of Approval Voting, memo, New York University.
- [16] A. Schultdt, J.O. Berndt, O. Herzog, The Interaction Effort in Autonomous Logistics Processes: Potential and Limitations for Cooperation In Autonomous Cooperation and Control in Logistics, pages 77-90, Springer, Berlin Heidelberg (2011).
- [17] B. Horling, V. Lesser, A Survey of Multi-Agent Organizational Paradigms, The Knowledge Engineering Review, volume 19:4 pages 281-316 (2005).
- [18] D. de Jonge, C. Sierra, NB3: a Multilateral Negotiation Algorithm for Large, Non-linear Agreement Spaces with Limited Time, AAMAS, Volume 29, Issue 5, pp 896-942 (2015).
- [19] M. A. Bercoff, L'art de negocier avec la methode Harward, Editions Eyrolles, 2009
- [20] A. Rosenfeld, I. Zuckerman, E. Segal-halevi, O. Drein, S. Kraus, NegoChat-A: a chat-based negotiation agent with bounded rationality, AAMAS, Volume 30, Issue 1,(2016).
- [21] S. De Clercq, S. Schockaert, A. Nowe, M. De Cock, Multi-lateral Negotiation in Boolean Games with Incomplete Information Using Generalized Possibilistic Logic, IJCAI,(2015).
- [22] U. Grandi, D. Grossi, P. Turrini, Equilibrium Refinement through Negotiation in Binary Voting, IJCAI, (2015).
- [23] N. Hazon, R. Lin, S. Kraus, How to Change a Group's Collective Decision?, IJCAI, (2013).
- [24] P. Faratin, C. Sierra, and N. R. Jennings, Negotiation Decision Functions for Autonomous Agents, Robotics and Autonomous Systems,(1998).
- [25] A. Farinelli, M. Bicego, F. Bistaffa and S. D. Ramchurn, A hierarchical clustering approach to large-scale near-optimal coalition formation with quality guarantees, EAAI, (2017).
- [26] Zuo B., Zheng S., Wu H. Multi-issue Agent Negotiation Based on Fairness, In International Conference Web Information Systems and Mining, (2009).