A Model for Multi-party Negotiations with Majority Rule *

Sheng Zhang[1], Fillia Makedon[1], James Ford[1], Lin Ai[2]

[1] Department of Computer Science
Dartmouth College
Hanover, NH 03755, U.S.A.
{clap,makedon,jford}@cs.dartmouth.edu
[2] erin_ai@hotmail.com

Abstract. Our model of multi-party negotiations is a many parties, many issues model. The whole multi-party negotiation consists of a set of mutually influencing bilateral negotiations that are focused on different bilateral issues. We propose to use majority rule to help parties reach group agreements. When a party is not satisfied with another party's negotiation progress, he can send a primitive *oppose* to the other. Those negotiation parties who get a sufficient number of *oppose* primitives from others or those negotiation parties who lack support in opposing others will be warned to make satisfactory concessions in the following negotiation round. So the will of majority affects each party's negotiation behavior and leads to the final group agreement.

1 Introduction

Suppose there are a group of participants (which are referred to as players) in the model, each pair of them has a bilateral negotiation issue. The final group agreement is based on a set of bilateral agreements. In other words, every two players conduct a bilateral negotiation to reach an agreement on their own bilateral issue, then all these bilateral agreements are merged to form a multilateral agreement. All the bilateral negotiations are conducted synchronously. A player will be involved in the final group agreement only if he has bilateral agreements with all the other players in the group. One difficulty in reaching a group agreement is that parties in each bilateral negotiation may make slower progress or even reach a stalemate due to a disagreement. To solve this problem, we propose to use majority rule by taking into account opinions from all the players. More intuitively, when two players have a disagreement, the player who gets more support from the other players should have an advantage over the player who has less support. For example, in a three players game, if both player A and player B are opposing player C in their respective bilateral negotiations with player C (assuming the negotiation between player A and B goes well), then player C should be warned to make concessions. Using majority rule also seems to bring players more fairness. We can imagine that in traditional bilateral negotiations,

* This work was supported in part by NSF ITR 0312629.

if player C is a hard-bargaining player, player A and player B may be required to make unnecessary or larger concessions in their respective bilateral negotiations to reach agreements with player C.

The rest of this paper is organized as follows. Section 2 analyzes the related work on multi-party negotiations. Section 3 presents our negotiation model. The first part describes the negotiation objective, negotiation primitives, negotiation protocols, and players negotiation behaviors. In the second part, two approaches are introduced to resolve disagreements in negotiations according to majority rule. Finally, Section 4 concludes the paper.

2 Related Work

Kraus [1] presents a strategic negotiation model which is based on Rubinstein's model of alternating offers [2]. In the strategic model there are N agents, and they need to reach an agreement on a given issue. In each period t of the negotiation, if the negotiation has not terminated earlier, an agent whose turn it is to make an offer and each of the other agents choose to either accept of offer (choose *Yes*), reject it (choose *No*), or opt out of the negotiation (choose *Opt*). If an offer is accepted by all the agents, then the negotiation ends with an agreement. If at least one of the agents opts out of the negotiation, then the negotiation ends and a conflictual outcome results. If no agent has chosen *Opt* but at least one of the agents has rejected the offer, the negotiation proceeds to period $t+1$. Sycara [4] presents a model that combines case-based reasoning and optimization of multi-attribute utilities. The model uses persuasive argumentation as a means of guiding the negotiation process to a settlement. Sierra et al. [3] present a model for autonomous agents to reach agreements about the provision of service by one agent to another in multi-agent environments.

Our work presents a model of multi-party negotiations with multiple bilateral issues, in which the group agreement is based on a set of bilateral agreements. The main contribution of our work is that we introduce a new negotiation primitive, *oppose*, that allows the negotiation process to take into account opinions from all parties. In this approach, we use majority rule to resolve the possible disagreement in each bilateral negotiation.

3 Negotiation Model

This section describes all aspects of the proposed negotiation model.

3.1 Negotiation Objectives

The group agreement involves a number of acceding players, in which every pair has a bilateral agreement. There are two kinds of final group agreements after the whole negotiation process ends. The first kind is a *global group agreement* in which all players in the game accede. The second kind is a *local group agreement.*

This agreement is a group consensus of a set of players, where this set is the subset of all the players in the game. The objective of our model is to reach a global group agreement. In cases where a global group agreement is not reached, it is possible that multiple local group agreements may have been agreed to, and that a single player may be involved in more than one local group agreement. If no global group agreement is reached, the objective of the model becomes twofold reaching more local group agreements and involving as many players as possible in each local group agreement.

3.2 Notations

P_i	ith player (N players in total)
N_i^n	number of players still negotiating with P_i in the nth round
O_i^n	number of *oppose* primitives P_i receives in the nth round
UL_i^n	upper limit for P_i in the nth round
LL_i^n	lower limit for P_i in the nth round
T	time limit
$Pr(w)$	probability of a player getting a warning
$Pr(w_1)$	probability of a player getting a warning because of the case 1 (see 3.3)
$Pr(w_2)$	probability of a player getting a warning because of the case 2 (see 3.3)
λ	probability of a player opposing another player in a round

3.3 Negotiation Primitives

This section describes six negotiation primitives used in the proposed model. The first four primitives are widely used in bilateral negotiations and the last two primitives are newly introduced to help us achieve group agreements using majority rule in multi-party negotiations.

- *Call-For-Proposal (CFP)*: At the start of negotiations, each player sends one initial proposal to each of the other players, or $N-1$ proposals in total. The CFP helps each player form conjectures about what the others would like to commit and what they would like to require in the final agreement.
- *proposal*: After players get CFPs from the other players, they begin formal negotiations by sending proposals or counterproposals to each other.
- *Accept*: P_i accepts the proposal sent by P_j when he is satisfied, which means that P_i and P_j have finished the negotiation successfully with a bilateral agreement. However, P_i and P_j may still have ongoing negotiations with other players.
- *Reject*: When P_i rejects continuing the negotiation with P_j, P_i sends *reject* to P_j to end their negotiation. Moreover, this rejection also announces that P_i has quit the final group agreement, and thus all bilateral negotiations that P_i is conducting are also aborted. However, P_j can still continue those negotiations that he is conducting with other players. Therefore, a player should never rashly use *reject*. To let P_i register dissatisfaction with proposals

from P_j while still keeping himself in the negotiation process, P_i can use the primitive *Oppose*.

- *Oppose*: *Oppose* gives P_i an opportunity to oppose or give notice to P_j if P_i thinks that P_j has not made satisfactory progress in their negotiation. However, *oppose* will not terminate the ongoing negotiation.
- *Warn*: This primitive is used by either the system or a chairman in a multi-party negotiation. It will be sent to P_i in the nth round only in one of the following two cases:
  1. *Warning Case* 1: In all of the last T (*time limit*) rounds, the number of *oppose* primitives that P_i receives has always exceeded P_i's *upper limit*;
  2. *Warning Case* 2: If P_i currently gives an *oppose* to P_j, and in the past T rounds (including the current round) P_i gave *oppose* primitives to P_j, the number of *oppose* primitives that P_j receives has never exceeded P_j's *lower limit*.

When receiving a warning, P_i is also notified of the case for which he gets this warning. If P_i receives a warning for the same case in two consecutive rounds, he will be driven out of the negotiation automatically (all those negotiations that P_i is currently conducting will be aborted), and he cannot join any final group agreement although he might have some bilateral agreements with other players already. So after receiving a warning, P_i needs to take action immediately. If he was warned because of the second case, then sending a new proposal instead of using *oppose* will (by definition) ensure P_i that he will not get another warning for the second case in the next round. However, if P_i was warned because of the first case, then even if he were to make concessions immediately, other players may still give him *oppose* primitives. For example, they might not be satisfied with his negotiation performance from the view of a whole process. Different players in different rounds may have different upper limits and lower limits. We will discuss how to compute them in subsection 3.7.

3.4 Negotiation Behaviors

In the following, we describe the main negotiation behaviors of a certain player P_i in his bilateral negotiation with another player P_j. (We do not list all possible cases as that would be too complicated.)

The negotiation starts in State 1 (see Figure 1). After P_i and P_j get CFPs from each other, they start a formal negotiation (State 2) and continuously send proposals and counterproposals to each other. If either P_i or P_j accepts the proposal given by the other, the negotiation between them is finished successfully (State 3).

If P_i sends *oppose* to P_j in State 2, then P_i moves to the state (State 4) of opposing the other. If P_j makes a satisfactory concession, then the parties will go back to the ongoing negotiation state (State 2). However, if P_j does not make an effective concession and P_i chooses to abort the negotiation, the negotiation will fail in state 7. If P_i chooses not to terminate the negotiation in this case, then if P_i is the only or one of a small number of players currently opposing P_j,

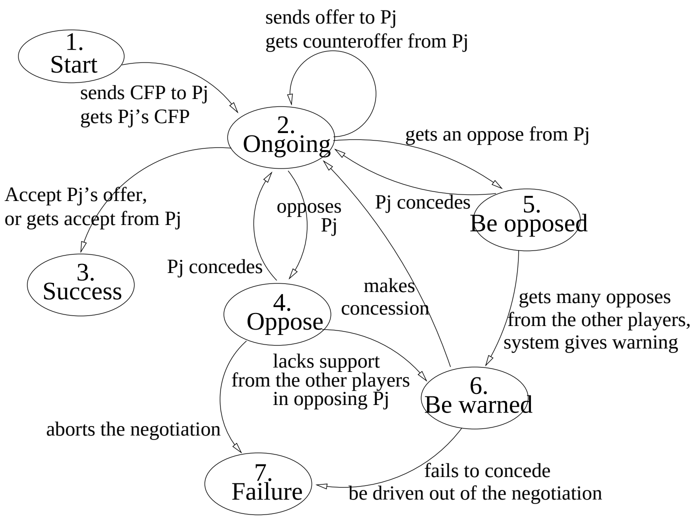

Fig. 1. The main behaviors of P_i in his negotiation with P_j

P_i will get a warning (State 6) later to urge him to make a concession (*i.e.*, to give a new proposal instead of merely opposing P_j). If P_i fails to do so, he will be driven out of the negotiation, so the negotiation between P_i and P_j also fails (State 7).

In State 2, P_i may also get an *oppose* from P_j and would then move to the state of being opposed (State 5). If P_i then produces a new proposal or counterproposal, they will return to State 2. Otherwise, if P_i also receives many *oppose* primitives in the other negotiations he is joining, then he will be warned to make concessions in these negotiations (State 6). Failing to do so will lead to State 7.

3.5 Negotiation Protocols

As all the bilateral negotiations are synchronized, only one player is allowed to send proposals (or other primitives) out at any time. So we need to coordinate all the players to ensure proposals are sent serially, one at a time. This can be done by using a ring structure like that of the strategic negotiation model [1]. Take each player as a node in the ring, and pick a node as the start node, which is the first node to send primitives to all the other nodes in the ring. Then, the next node (either clockwise or counter-clockwise) sends his primitives, and so on. When there is a player who aborts the negotiations or is driven out of the negotiations, or finishes negotiations successfully with all the other nodes in the ring, his corresponding node is removed and two neighbor nodes are connected.

A *round* for P_i is the duration from the last turn of P_i to the current turn of P_i. According to the protocol, P_i has received all the primitives sent to him from the other nodes in each round. Thus, in P_i's turn, he can apply his negotiation strategy (whatever it is) to produce counterproposals for all ongoing bilateral negotiations. P_i may also need to coordinate those counterproposals if there is a dependence constraint among two or more negotiations.

Moreover in P_i's turn, the system or chairman will compute the current upper limit and lower limit for P_i and check whether the conditions for any of the warning cases are satisfied. If the conditions for warning case 1 are satisfied, P_i will get a warning immediately. If the conditions for warning case 2 are satisfied, other corresponding players will get warnings in their next turns.

3.6 Abusing Oppose

The *oppose* primitive helps to resolve a disagreement between two players in their bilateral negotiation. In other words, a player can consider those *oppose* primitives going to the other party as support. Although this support is not direct (because other players do not know about negotiations that they don't join), it still represents important feedback summarizing the opinions of a player's negotiation partners. Practically speaking, a player who gets less support is more likely to get a warning later. However, regulating the use of *oppose* requires some care. Having only an upper limit would be problematic because a player could abuse *oppose* without ever triggering any negative effects on himself. Therefore we also set a lower limit, which is used to control this kind of abuse. Since a player only knows those primitives that he sends out or receives (he does not know how many *oppose* primitives another player receives currently), it is not possible for a player to predict whether his use of oppose can trigger a Case 1 warning for another player, and it always runs the risk of triggering a Case 2 warning for himself.

3.7 Setting Warning Limits

We have introduced three kinds of limits: a lower limit for player i in the nth round (LL_i^n), an upper limit for player i in the nth round (UL_i^n) and the time limit (T) described in the primitive *warn*. Adjusting these limits directly affects the probabilities of players getting warnings. If the lower limit is too low, we cannot prevent players from abusing *oppose* because they will probably not get warnings. On the other hand, if the lower limit is too high, we may inadvertently punish innocent players. Likewise for the upper limit, settings that too high (so that neither player in a negotiation gets a warning) or too low (so that both players get warnings) will not help in resolving disagreements. A similar tradeoff also applies to the time limit: a higher time limit may make the whole multi-party negotiation last longer, but a lower time limit may cause warnings to become increasingly arbitrary. Here, we assume that the time limit is set already and discuss how to set lower limits and upper limits based on this assumption.

We assume that the probabilities of a player opposing another player (λ) in different rounds are independent (which is based on the assumption that a player will make a positive response after getting an *oppose*), so λ is always the same during the whole negotiation process. We assume that in the nth round, the number of *oppose* primitives that player i gets in the nth round (O_i^n) is binomially distributed. So the expected number is λN_i^n and the standard deviation is $\sqrt{N_i^n \lambda(1-\lambda)}$.

The probabilities of a certain player getting a warning ($Pr(w)$), getting a warning for Case 1 ($Pr(w_1)$), and for Case 2 ($Pr(w_2)$) satisfy $(1 - Pr(w_1)) \times (1 - Pr(w_2)) = 1 - Pr(w)$. If we hope $Pr(w_1) = Pr(w_2)$, and $Pr(w) \leq \epsilon$, we need to satisfy:

$$Pr(w_1) = Pr(w_2) \leq 1 - \sqrt{1-\epsilon}.$$

P_i will get a warning in the kth round because of warning Case 1 if $\forall n$: $k-T+1 \leq n \leq k$, we have $O_i^n \geq UL_i^n$, so

$$Pr(w_1) \geq \prod_{n=k-T+1}^{k} Pr(O_i^n \geq UL_i^n).$$

A player other than P_i who intends to abuse the primitive *oppose* to P_i will get a warning if $\forall n$: nth round is one of those past T rounds (including the current round) in which he sends an *oppose* to P_i, we have $O_i^n \leq LL_i^n$. Therefore, $Pr(w_2)$ is larger than or equal to the probability that this player gets a warning in this case. So we have:

$$Pr(w_2) \geq \prod_{n\text{th round} \in \text{ those T rounds}} Pr(O_i^n \leq LL_i^n).$$

To simplify the computation, we just let

$$Pr(O_i^n \geq UL_i^n) \leq \sqrt[T]{Pr(w_1)} \leq \sqrt[T]{1 - \sqrt{1-\epsilon}},$$

$$Pr(O_i^n \leq LL_i^n) \leq \sqrt[T]{Pr(w_2)} \leq \sqrt[T]{1 - \sqrt{1-\epsilon}}.$$

Using Chebyshev's Inequality, we can get:

$$UL_i^n \geq \lambda N_i^n + \frac{\sqrt{N_i^n \lambda(1-\lambda)}}{\sqrt[2T]{1-\sqrt{1-\epsilon}}},$$

$$LL_i^n \leq \lambda N_i^n - \frac{\sqrt{N_i^n \lambda(1-\lambda)}}{\sqrt[2T]{1-\sqrt{1-\epsilon}}}.$$

So UL_i^n and LL_i^n are decided by N_i^n, ϵ, λ, and T. Of these four terms, N_i^n can be computed directly, ϵ and T can be decided from previous experience, and λ can be computed from past multi-party negotiations.

We introduced upper limits to resolve disagreements in bilateral negotiations. However, whatever upper limits are set to, we can always meet a problem like the following that we cannot handle by our current approach. Consider two players opposing to each other in a bilateral negotiation, who meanwhile both get a number of *oppose* primitives (more than or equal to their respective upper limits). If this situation is maintained, they will both ultimately get warnings. This is probably not fair to one of them. We propose to solve this problem by introducing another approach which also uses the idea of majority rule.

3.8 Potential Value

In the following, we will use a simple example which is based on coalitions to present this approach. In our example (see Figure 2(a)), five players are conducting a multi-party negotiation. After several rounds of negotiation that resolve the competition in each coalition, P_1,P_2, and P_3 form one coalition while P_4 and P_5 form another. Suppose each bilateral negotiation between two players in the same coalition has succeeded. However, suppose that in those bilateral negotiations across coalitions, players give *oppose* to each other. Because to each P_i, $O_i^n = N_i^n$ (assuming all players just finished the nth round), all players will get warnings soon.

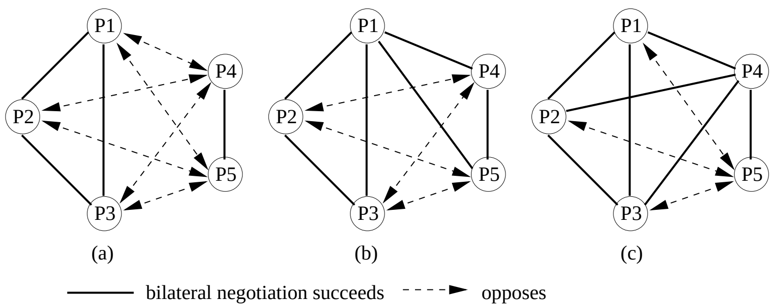

Fig. 2. An example of Potential Value. (a) is the original situation. (b) is the situation when P_1 finishes all negotiations successfully with those players currently opposing him. (c) is the situation when P_4 finishes all negotiations successfully with those players currently opposing him.

If each player is considered as a vertex in an undirected graph and an edge is added between P_i and P_j if the negotiation between them has succeeded, then each connected component in the graph exactly represents a local group agreement. For any connected component c in this graph, we define $Value(c)$ (the value of component c) as the number of edges in this component. Furthermore, we define $Value(P_i)$ (the value of a player P_i) as the sum of values of those connected components that include P_i. In other words,

$$Value(P_i) = \sum_{\forall c, P_i \in c} Value(c).$$

In our example (see Figure 2(a)), the connected component including P_1 is $\{(P_1, P_2, P_3)\}$, so the current value of player P_1 is 3. Similarly, the connected component including P_4 is $\{(P_4, P_5)\}$, so the current value of P_4 is 1.

The potential value of a player P_i (defined as $PV(P_i)$) is the difference between the future value of P_i (when all the players who currently *oppose* P_i finally finish negotiations successfully with P_i) and the current value of P_i. When all the players who currently *oppose* P_1 (*i.e.*,P_4 and P_5) finally finish negotiations successfully with P_1, the case is like Figure 2(b). In that situation, connected components including P_1 are $\{(P_1, P_2, P_3), (P_1, P_4, P_5)\}$, so the future contribution of P_1 is $3+3=6$. Therefore, $PV(P_1) = 6-3 = 3$. Similarly in figure 2(a), if all the players who currently *oppose* P_4 (*i.e.*, P_1,P_2 and P_3) finish negotiations successfully with P_4, then the situation is like Figure 2(c). So the future value of P_4 is $6+1=7$, and $PV(P_4) = 7-1 = 6$.

Those players who have larger potential values have the potential to contribute to more and larger connected components (more and larger local group agreements) when getting warnings compared with those players having smaller potential values. Idea is therefore that if two players *oppose* each other and meanwhile both of them receive *oppose* primitives in excess of their upper limits, we will do the following modification: we remove the *oppose* that is sent from the player having a higher potential value to the player having a lower potential value. Algorithm 1 is used to compute the updated number of *oppose* ($O_i'^n$) after this modification. This value replaces the original value O_i^n in checking warning Case 1 for P_i in the nth round. (O_i^n is still used in checking warning Case 2.) This algorithm ensures that as long as those players who will possibly be warned do not have the same potential value, there is at least one player who will be warned so that the whole multi-party negotiation will not become a stalemate.

Algorithm 1: Computation of $O_i'^n$ in checking warning Case 1

(1) compute O_i^n and $PV(P_i)$
(2) $O_i'^n \leftarrow O_i^n$
(3) **if** $O_i^n \geq UL_i^n$
(4)   **foreach** P_j currently sending *oppose* to P_i
(5)   **if** P_i is also currently sending *oppose* to P_j and ($O_j^n \geq UL_j^n$ or $O_j^{n-1} \geq UL_j^{n-1}$ if O_j^n is not computed yet)
(6)   **if** $PV(P_i) < PV(P_j)$
(7) $O_i'^n \leftarrow O_i'^n - 1$

In our example, the original computation of O_i^n in warning Case 1 gives us the Figure 3(a). After applying Algorithm 1, we have Figure 3(b), in which $O_1'^n$ is now equal to 0 because $PV(P_1) < PV(P_4)$ and $PV(P_1) < PV(P_5)$. Similarly, $O_2'^n = 0$, $O_3'^n = 0$, $O_4'^n = 3$ and $O_5'^n = 3$. If this situation continues for a series of rounds, P_4 and P_5 will eventually be warned to make concessions. From another point of view, we can see that a larger coalition (*i.e.*,$\{P_1, P_2, P_3\}$) will eventually dominate over a smaller coalition (*i.e.*,$\{P_4, P_5\}$).

However, it is still possible that disagreements cannot be resolved when all the players who receive more *oppose* primitives than their upper limits have the same potential value. If at that time players continue to *oppose* each other in all ongoing bilateral negotiations, which means there is no real progress in any bilateral negotiation, then the system can terminate the whole multi-party negotiation. It is also necessary to terminate the whole multi-party negotiation when it has undergone a longer time than a predetermined limit. This limit can be decided by all the players to prevent negotiations going to an infinite loop.

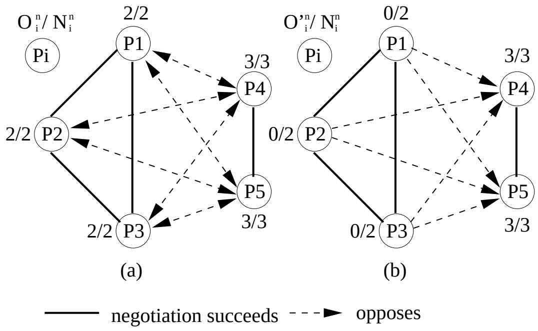

Fig. 3. (a) is the original computation of O_i^n. (b) is the computation of $O_i'^n$ after using potential values. Those *oppose* primitives from players having higher potential values to player having lower potential values are removed in (b).

Thus, in this approach several resolutions are possible. The multi-party negotiation ends (either finishes naturally or is terminated by the system) as one of the following three cases: 1) all the remaining players reach a local group agreement; 2) those players not finishing all their negotiations have a same potential value; 3) the multi-party negotiation lasts too long. In cases 2) and 3), all local group agreements in the ending moment can still be considered valid.

4 Conclusion

We present a model of multi-party negotiations where group agreements are constructed using multiple bilateral agreements. The whole negotiation process is based on a set of synchronous bilateral negotiations. Two approaches using majority rule are presented to resolve possible disagreements in these bilateral negotiations to help all parties reach a final group agreement. One uses the statistical method to set upper limits and lower limits and the other uses potential values. Future work includes conducting more accurate analysis on the statistical method, and designing other approaches or incentives to induce players to reach a group agreement.

References

1. S. Kraus. Automated negotiation and decision making in multiagent environments. volume 2086 of *Lecture Notes in Computer Science*, pages 150–172, 2001.
2. A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 1982.
3. C. Sierra, P. Faratin, and N. Jennings. A service-oriented negotiation model between autonomous agents. In *Proceedings of the 8th European Workshop on Modeling Autonomous Agents in a Multi-Agent World (MAAMAW-97)*, pages 17–35, Ronneby, Sweden, 1997.
4. K. Sycara. Persuasive argumentation in negotiation. *Theory and Decision*, 28(3):203–242, May 1990.