

Bilateral negotiation of a meeting point in a maze

Maxime Morge, Jean-Christophe Routier, Fabien Delecroix

► **To cite this version:**

Maxime Morge, Jean-Christophe Routier, Fabien Delecroix. Bilateral negotiation of a meeting point in a maze. International Conference on Practical Applications of Agents and Multi-Agent Systems (PAAMS), Jun 2014, Salamanca, Spain. Springer, 8473, pp.86-97, 2014. <hal-01058676>

HAL Id: hal-01058676

<https://hal.inria.fr/hal-01058676>

Submitted on 27 Aug 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Bilateral Negotiation of a Meeting Point in a Maze

Fabien Delecroix, Maxime Morge, and Jean-Christophe Routier

Laboratoire d'Informatique Fondamentale de Lille
Université Lille 1

Cité Scientifique- F-59655 Villeneuve d'Ascq

{fabien.delecroix,maxime.morge,jean-christophe.routier}@lifl.fr

Abstract. Negotiation between agents aims at reaching an agreement in which the conflicting interests of agents are accommodated. In this paper, we present a concrete negotiation scenario where two agents are situated in a maze and the negotiation outcome is a cell where they will meet. Based on their individual preferences (a minimal distance from their location computed from their partial knowledge of the environment), we propose a negotiation protocol which allows agents to submit more than two proposals at the same time and a conciliatory strategy. Formally, we prove that the agreement reached by such a negotiation process is Pareto-optimal and a compromise, i.e. a solution which minimizes the maximum effort for one agent. Moreover, the path between the two agents emerges from the repeated negotiations in our experiments.

1 Introduction

Negotiation is a form of interaction in which a group of agents with conflicting interests try to reach a mutually acceptable agreement over some outcomes [1]. The outcome is typically a tasks/resources allocation, a matching between agents or a joint decision. Agents' interests are conflicting in the sense that they cannot be simultaneously fully satisfied. In this perspective, negotiation can be seen as a distributed search through a space of potential agreements [2].

In this paper, we present a concrete negotiation scenario where two agents are situated in a maze and the negotiation outcome is a cell where they will meet. Based on their individual preferences (a minimal distance from their location computed from their partial knowledge of the environment), we propose a negotiation protocol which allows agents to submit more than two proposals at the same time and a strategy which consists in starting from the deal that is best for the agent and then concedes. A concession of an agent means that she proposes a new deal such there is no other preferred alternatives. Formally, we prove that the agreement reached by such a negotiation process is Pareto-optimal and a compromise, i.e. an alternative which minimizes the maximum effort for one agent to reach it. Moreover, the path between the two agents emerges from the repeated negotiations in our experiments.

Paper Overview. Section 2 describes the addressed concrete problem and why the negotiation frameworks in the literature are not sufficient for it. In section 3, we introduce the basic notions in the background of our work. Section 4 proposes a rule for multi-agents decision. Then, we present our negotiation game (Sec. 5). We describe our experiments in section 6. Finally, section 7 concludes.

2 Problem

We consider here two agents which are paratroopers landed at the two opposite corners of a maze. They aim at meeting as soon as possible, i.e minimizing the maximum number of steps for one agent to reach the meeting point. Both of them have a local perception of the environment. Each agent can perceive the walls of her current cell. Moreover, she knows her own location. Contrary to the classical rendezvous problem 3, the agents are allowed to communicate in order to negotiate the meeting point. Moreover, the meeting point can be re-negotiated during the exploration of the maze. The optimal solution for finding a meeting point requires the knowledge of the whole maze. Under this assumption, the agents can compute the shortest path between them and set the meeting point in the middle of it. By contrast, a solution which does not need any prior knowledge consists of pseudo-randomly selecting a meeting point in the first diagonal.

In order to illustrate this problem, we consider a 3×3 maze (cf Fig 1). At the second step of the resolution, Alice is in the cell d while Bob is in the cell f . Each agent computes the distance to reach all the other cells based on its knowledge. For this purpose, an agent takes into consideration the perceived walls and she assumes that there is no wall between the cells it did not visit yet. In other words, the computation is performed by an A-star algorithm where the future path-cost function is the Manhattan distance. For instance, Bob supposes that 3 steps are required to reach the cell c since she is aware there is a wall between c and f , and so it plans to go through e and b . However, this path cannot be followed since there is a wall between e and b the agent is not aware of.

Since we want to minimize the maximum number of steps for an agent to reach the meeting point, the cell e is a good candidate even if d , e and f are Pareto-optimal (see Def. 1). In order to solve this distributed solving problem, we need a negotiation protocol and a strategy which allow to reach a fair solution. It is worth noticing that the communication of their position is not enough to reach a rendezvous with a pure strategy 3. Moreover, the communication of the wall is useless since the agents explore different parts of the maze.

Related Works. Many negotiation frameworks have been proposed in the literature (see 1 for a survey) depending on the object of negotiation, the agents' preferences (2 or n), the protocol and the strategy. First of all, we consider here 2 agents (the paratroopers) negotiating a single-issue with discrete values (the meeting point).

Model for the Agents' Preferences. Most of the literature assume that the preferences are represented by utility functions in order to negotiate a payoff,




Fig. 1. The maze (at middle) and its internal representation for Alice (at left) and Bob (at right) at time $t = 2$. In the latter, the visited cells are in grey, some walls may be still unknown and each cell is labelled with an estimation of the shortest path length.

i.e. a point on the curve of these functions. In our particular example, we can remark that two points are at the same distance of Alice : a and e . Then, we must propose a negotiation strategy based on large preferences for agents. More generally, we will only assume a preference relation which is incomplete in order to consider not only the indifference but also the incomparability between alternatives. [4] allows agents to have a qualitative preference model (i.e. a CP-net) for negotiating. Similarly, we will compute the rank of an alternative using its depth in the preference graph.

Protocol. The negotiation protocol is the set of rules which regulates the exchanges of proposals. It can be symmetric (e.g. [5]) or asymmetric (e.g. [6]). Contrary to [4], our protocol does not give the priority to one agent and each agent can submit more than one alternative at each round: it is required for reaching a fair agreement with large preferences.

Strategy. The negotiation strategy must be set up according to the model for the agents' preferences and the protocol. Its features are the availability of information about opponents and its efficiency. While some negotiation strategies make the assumption of perfect information (e.g. [5]) we think we cannot make such assumption in our context. Most of the existing strategies leads to a Pareto-optimal solution (e.g. [5]). Additionally, we aim at proposing a negotiation process for distributed problem solving which minimizes the maximum effort for one agent. [7] considers social choice theory to allow agents to choose among alternatives based on their social value since it knows the preferences of others. We do not assume here any knowledge about the preferences of the opponents.

In summary, we aim at proposing:

1. a negotiation strategy based on large (and eventually incomplete) preferences which does not assume that agents know the preferences of each other;
2. a protocol which allows more than two offers per round;
3. a fair negotiation process which does not give priority to one agent and which minimize the maximum effort for one agent.

3 Background

In order to represent the taste of the decision maker and to compare the alternatives to each other, we assume here a preference relation on a non-empty finite set of alternatives \mathcal{X} , i.e. a preorder relation (reflexive and transitive) denoted \succsim . By contrast, the corresponding strict preference relation (denoted \succ) is a strict order, i.e. transitive and asymmetric. The indifference relation captures the indifference of the decision maker between alternatives. It means that the decision maker believes that, according to its preferences, there is no real difference between x and y . Moreover, We remark that the preference relation can be incomplete. Two alternatives can be incomparable if it is impossible for the decision maker to compare them. It can be interpreted as a way for the decision maker to refuse to commit due to an uncertain judgment. Contrary to the indifference relation, the incomparability relation is not transitive.

The notion of non-dominance allows to distinguish the alternatives for which there is no preferred alternatives. The set of **non-dominated alternatives** over \mathcal{X} wrt \succsim is the set: $ND(\mathcal{X}, \succsim) = \{x \in \mathcal{X} \mid \forall y \in \mathcal{X}, \neg(y \succ x)\}$. It is worth noticing that there is always at least one non-dominated alternative.

Example 1. We consider here the set of alternatives $\mathcal{X} = \{a, b, c, d, e, f, g, h, i\}$ and the two preference relations \succsim_1 and \succsim_2 over \mathcal{X} corresponding to our previous example. In our case, a cell is preferred to another if the estimated distance towards the first cell is at least as good as the second cell. The preference graph of \succsim_1 (resp. \succsim_2) is represented at left (resp. at right) of Fig. 2 as a directed graph where a node represents an alternative and there is a edge from x to y when x is at least as good as y . We can remark that $ND(\mathcal{X}, \succsim_1) = \{d\}$ and $ND(\mathcal{X}, \succsim_2) = \{f\}$.




Fig. 2. Preference relation \succsim_1 (at left) and \succsim_2 (at right)

We focus now on a group of agents (two or more) taking a joint decision. We consider here a non-empty finite set of alternatives \mathcal{X} , a set of agents Ω and for each agent $i \in \Omega$, \succsim_i is the preference relation for i over \mathcal{X} .

The Pareto rule is used by a group of agents to compare two alternatives. An alternative x Pareto-dominates an alternative y for a group of agents if x is at

least as good as y for all the agents and at least one agent strictly prefers x to y . An alternative is Pareto-optimal if it is not Pareto-dominated.

Definition 1 (Pareto). *The **Pareto-dominance** relation $\succ_{\Omega} \subseteq \mathcal{X} \times \mathcal{X}$ is defined such that $\forall (x, y) \in \mathcal{X}^2, x \succ_{\Omega} y \Leftrightarrow (\forall i \in \Omega, x \succ_i y) \wedge (\exists j \in \Omega, x \succ_j y)$. The set of **Pareto-optimal** alternatives for Ω over \mathcal{X} is the set $ND(\mathcal{X}, \succ_{\Omega}) = \{x \in \mathcal{X} \mid \forall y \in \mathcal{X} \neg(y \succ_{\Omega} x)\}$*

It is worth noticing that the Pareto-dominance is a strict order. The Pareto-optimality captures the notion of rationality for multi-agents. Indeed, the alternatives are Pareto-dominated since, from the viewpoint of the group of agents, better alternatives are available. Moreover, there is always a Pareto-optimal alternative. The Pareto-optimality is not a sufficiently discriminatory rule. For instance, in our previous example (cf. Fig. 2), the set of Pareto-optimal alternatives is $\{e, d, f\}$. However, e seems to be more “fair” than d and f .

4 Multi-agents Decision

We aim at setting a rule for multi-agents decision which is Pareto-inclusive and we want to warrant the existence of a solution. We call them compromises.

Like the utility functions allow to evaluate the individual satisfaction of agents, we introduce a rank function for evaluating the effort performed by an agent to accept an alternative.

Definition 2 (Individual rank function). *Let \succ_i the preference relation of the agent i over \mathcal{X} . Our rank function is defined such that :*

$$r(x, \mathcal{X}, \succ_i) = \begin{cases} 1 & \text{if } x \in ND(\mathcal{X}, \succ_i) \\ 1 + r(x, \mathcal{X} \setminus ND(\mathcal{X}, \succ_i), \succ_i) & \text{otherwise} \end{cases}$$

The rank of an alternative is its level on the preference graph. By taking into account the rank of an alternative, we make the assumption that any concession - the fact to withdraw an alternative for a worst one - of any agent represents the same effort.

In order to obtain Pareto-optimal and fair alternatives, we define the leximin rule on the alternative ranks.

Definition 3 (Leximin Preference). *Let $x, y \in \mathcal{X}$ be two alternatives. We denote $\mathbf{x}_r(\Omega) = (x_1, \dots, x_n)$, the vector of alternative ranks in decreasing order. We say that x is strictly **leximin-preferred** than y (denoted $x \succ_{lex} y$) iff $\exists k \leq n, \forall i < k, x_i = y_i$ and $x_k < y_k$. The **leximin-optimal** set over \mathcal{X} is $ND(\mathcal{X}, \succ_{lex})$.*

The leximin relation is a partial strict order.

All the compromises are Pareto-optimal and there is always one.

Property 1 (Compromises)

1. $Compromises(\mathcal{X}, \Omega) \subseteq ND(\mathcal{X}, \succ_{\Omega})$.
2. There always exists $x \in \mathcal{X}$ such that $x \in Compromises(\mathcal{X}, \Omega)$.

Proof 1 (Compromises)

1. *Proof by contradiction. Let x be a compromise over \mathcal{X} for Ω . We assume that there is $y \in \mathcal{X}$ which Pareto-dominates x . From Def. 1 we deduce (1) $\forall i \in \Omega, y \succ_i x$ and (2) there exists an agent $j \in \Omega$ such that $y \succ_j x$. From (1) and Def. 2, we deduce that $\forall i \in \Omega, r(x, \mathcal{X}, \succ_i) \geq r(y, \mathcal{X}, \succ_i)$. From (2) and Def. 2, we deduce there exists an agent $j \in \Omega$ such that $r(x, \mathcal{X}, \succ_j) > r(y, \mathcal{X}, \succ_j)$. From Def. 3, we deduce that $y \succ_{lex} x$ and so x is not a compromise, which is a contradiction.*
2. *Since the leximin is an order over a non-empty finite set, this set contains at least one minimal element.*

We can remark that some Pareto-optimal may not be compromises. In our example (cf. Fig 2), e is the only compromise.

In summary, a compromise captures the fairness of a solution. This rule is Pareto-inclusive and a compromise always exists.

5 Bilateral Negotiation

We propose a bilateral negotiation game in order to reach an agreement. Firstly, we introduce the protocol. Secondly, we present a negotiation strategy. Finally, we evaluate the agreements.

Protocol. Since we do not want to give priority to the agent which speaks first, we consider here a simultaneous game made of several rounds. For this purpose, we introduce an arbitration mechanism¹.

At each round, the arbitrator collects the proposals - each of them is a set of alternatives. The protocol forbids the empty proposals and the repetition of an alternative in two proposals of the same participant. When an alternative is proposed by both participants, the arbitrator closes the game, detects the agreements and informs the participants. Otherwise, it forwards these proposals to the other participant and a new round starts (see Fig. 3).

In the rest of this paper, we denote $p_1(k)$ the proposal of the agent 1 at round $k \geq 1$, i.e. a set of alternatives. At each round t , we denote $H_1(t)$ the history of the proposals of the agent 1 such that $H_1(t) = \bigcup_{k < t} p_1(k)$ and $H(t)$ is the game history $H(t) = H_1(t) \cup H_2(t)$. Moreover, the set of alternatives sent by the agent 1 during the history $H_1(t)$ is denoted $\mathcal{P}_1(t) = \{x \in \mathcal{X} \mid \exists p \in H_1(t), x \in p\}$.

At each round t , if the arbitrator identifies at least one alternative proposed by both participants, the game ends. The arbitrator gives priority to the alternatives which appear earlier in the history (cf. Algo. 1) and it returns this set

¹ For clarity, we choose to implement the synchronization with a mediator. However, decentralized solutions are easy to establish.




Fig. 3. Protocol for the participants (at top) and for the arbitrator (at bottom). The initial states are represented by double circles and the final ones are in black.

of alternatives, called **agreements**, to the participants. An agreement is not necessarily unique. In order to reach a collective decision, the arbitrator may select one agreement by using a (pseudo-)randomized function.

Algorithm 1. Arbitration

Data: $H(t)$: history, $p_1(t)$ and $p_2(t)$: proposals

```

1 for  $k = 0$  to  $t$  do
2    $A = (p_1(k) \cap p_2(t)) \cup (p_2(k) \cap p_1(t))$ ;
3   if  $A \neq \emptyset$  then
4      $send(\{a_1, a_2\}, agreement(A))$ ;
5     return;
6  $send(a_2, proposal(p_1(t)))$ ;
7  $send(a_1, proposal(p_2(t)))$ ;
```

Strategy. The strategy of a participant interfaces with the protocol through the condition mechanism of utterance and interpretation of the proposals. Obviously, when an agent receives a proposal, she updates her representation of the history. The content of the proposals is determined by the strategy of the agent.

In this section, we adopt the viewpoint of the agent 1 (since the roles are symmetric). Her strategy (see Algo. 2) is legal: a proposal is not empty and the alternatives are not repeated in its proposals. Moreover, the strategy is rational: the agent chooses among the legal alternatives (\mathcal{X}') those that may be compromises from her viewpoint. The negotiation heuristic consists in choosing among the rational alternatives.

A **conciliatory** agent selects all the rational alternatives in order to reach an agreement as soon as possible ($select(Rat, H(t)) = Rat$). By contrast, other strategies may aim at minimizing the individual rank of an agreement. Since we consider distributed problem solving, we will make the assumption that agents are collaborative and so conciliatory. An agent concedes since she proposes a new deal such there is no other preferred alternatives which are legal.

Algorithm 2. Negotiation strategy

Data: $H(t)$: the history, \mathcal{X} : the set of alternatives
1 $\mathcal{X}' = \mathcal{X} \setminus \mathcal{P}_1(t)$;
2 $Rat = ND(\mathcal{X}', \succ_1)$;
3 $E = select(Rat, H(t))$;
4 $send(arbitrator, proposal(E))$;

Example 2. *Let us consider the preferences represented in Fig. 2. We consider the negotiation games where the agent 1 (resp. 2) adopts a conciliatory strategy. In this game, the agent 1 (resp. 2) starts with the alternative d (resp. the alternative f). At the second round, the agent 1 offers both a and e while the agent 2 offers both e and i . Therefore, e is an agreement.*

Theoretical Evaluation. We identify here the properties of the agreements.

First, we can remark that the negotiation game always leads to an agreement.

Property 2 (Guaranteed success). *Our negotiation game ends successfully.*

Proof 2 (Guaranteed success). *Since the set of alternatives is finite and the protocol forbids the repetition of alternatives in the proposals of the same agent, the game ends. The set of alternatives with the precedence relation in the history is a non-empty and finite set with a total order. Therefore, this set has always at least one minimal element, i.e. an agreement.*

An agreement reached by a negotiation game between two conciliatory agents is a Pareto-optimal alternative and a compromise.

Theorem 1 (Agreements). *If both participants adopt a conciliatory strategy, then the set of agreements reached $A_{CC} \subseteq \mathcal{X}$ is such that:*

1. $A_{CC} \subseteq ND(\mathcal{X}, \succ_\Omega)$
2. $x_{CC} \in A_{CC} \Rightarrow compromise(x_{CC}, \Omega, \mathcal{X})$

First, we remark that the timing of proposals depends on their rank.

Lemma 1. *Let $x \in \mathcal{X}$ be an alternative and a negotiation game which stops at time $\theta > 0$. $r(x, \mathcal{X}, \succ_i) < \theta \Rightarrow x \in p_i(r(x, \mathcal{X}, \succ_i))$.*

Proof 3 (Agreements)

1. *Let $x \in A_{CC}$ be an agreement. We prove by contradiction that $A_{CC} \subseteq ND(\mathcal{X}, \succ_\Omega)$. We assume $x \notin ND(\mathcal{X}, \succ_\Omega)$. Therefore, $\exists y, (y \succ_1 x \wedge y \succ_2 x) \vee (y \succ_2 x \wedge y \succ_1 x)$. So, $\exists t, y \in \mathcal{P}_1(t) \wedge y \in \mathcal{P}_2(t) \wedge (x \notin \mathcal{P}_1(t) \vee x \notin \mathcal{P}_2(t))$. Therefore $x \notin A_{CC}$ which is a contradiction.*
2. *Let $x \in A_{CC}$ be an agreement and θ the time when the game stops. Proof by contradiction. We assume $x \notin Compromises(\mathcal{X}, \Omega)$. (A) From (A) and Def. 3 and Prop. 1, $\exists y \in \mathcal{X}, y \succ_{lex} x$. (B) From Def. 3 and (B), we distinguish 2 cases:*

- (a) $\max_{a \in \Omega} r(y, \mathcal{X}, \succ_a) < \max_{a \in \Omega} r(x, \mathcal{X}, \succ_a)$ (C)
 From (C) and Lemma 4, we deduce there exists $t \geq 1$ when both participants play y but only one have played x . Therefore, from Algo. 4, $x \notin A_{CC}$, which is a contradiction.
- (b) $\max_{a \in \Omega} r(y, \mathcal{X}, \succ_a) = \max_{a \in \Omega} r(x, \mathcal{X}, \succ_a)$ (D) and
 $\min_{a \in \Omega} r(y, \mathcal{X}, \succ_a) < \min_{a \in \Omega} r(x, \mathcal{X}, \succ_a)$ (E).
 From (D), (E) and Lemma 4, we deduce that x is played first before y while x and y are played a second time during the same round. From Algo. 4, we deduce that $x \notin A_{CC}$, which is a contradiction.

Actually, the game aims at splitting the effort between the participants.

6 Experiments

Our experiments aim at comparing the improvement of the distributed solving with the negotiation and its communication costs.

The pseudo-random generation of a maze of size $n \times n$ is performed by:

- a pseudo-randomized version of Prim’s algorithm. This algorithm results in mazes with many short dead ends and the solution is usually pretty direct as well;
- a depth-first search algorithm. This algorithm results in mazes with fewer but longer dead ends, and the solution is usually very long and twisty.

We pseudo-randomly generate a new $n \times n$ maze, then we try to solve the corresponding problem by:

1. the pseudo-randomly selection of a meeting point in the fist diagonal;
2. the negotiation of the meeting point at each step;
3. the negotiation of the meeting point every 2 steps;
4. the negotiation of the meeting point every 4 steps;
5. the negotiation of the meeting point each time the path length towards the meeting point increases for one agent due to the discovery of a wall.

We consider $n \in [2; 14]$ and for each n , we generate 100 experiments.

In Fig. 4, we consider the maximum number of steps for one agent and the total number for the two agents. Negotiating rather than pseudo-randomly choosing a meeting point improves the time to reach it. More surprisingly, the global satisfaction of the agents is improved by the negotiation in simple mazes. In other words, negotiations may help to find shortest paths.

In Fig. 5, we compare the communication cost of our different strategies in terms of number of negotiations, total number of rounds and total number of proposed alternatives. The results are presented for mazes generated with the randomize version of Prim’s algorithm but they are very similar when the depth-first search algorithm is used. Negotiating at each step increases the communication cost with similar results for the number of steps. It seems that the communication cost of the strategies # 3 and # 5 are similar.

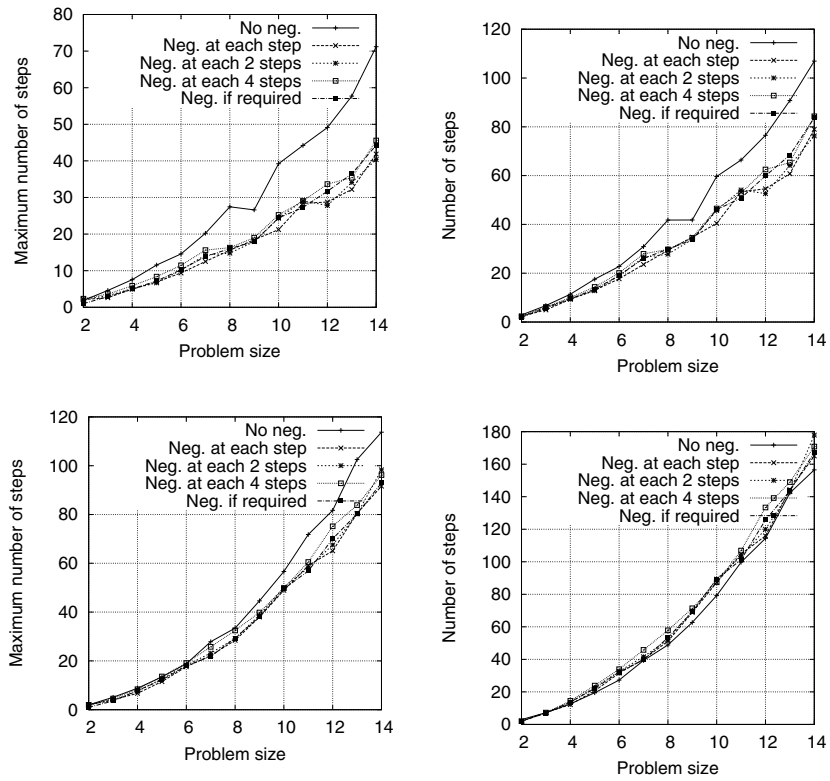


Fig. 4. Maximum number of steps by one agent and the total number of steps performed by two agents with mazes generated with a pseudo-randomized version of Prim's algorithm (top) and with a depth-first search algorithm (bottom).




Fig. 5. Number of negotiations (at top), rounds (at center) and proposals (at bottom)

We generate now a 8×8 maze with the pseudo-randomized version of Prim’s algorithm and the location of the agents. We partially destroy the walls until the density (nb of walls/nb of initial walls) is d %. Then, we try to solve the corresponding problem with the help of our 5 strategies. We consider $d \in [0; 100]$. For each d , we generate 100 experiments. In Fig. 6 we observe that negotiating rather than pseudo-randomly choosing a meeting point improves the time to reach it. Moreover, we cannot distinguish the efficiency of the different strategies. However, these strategies have different communication cost. For instance, the strategy # 5 is closed to the strategy # 3 when the density is high.




Fig. 6. Number of maximum steps (at left) and proposals (at right) depending on the density of the maze

In order to resume, the strategy #4 seems to be the best one due to the maximum number of steps by one agent and its communication cost whatever the maze is.

7 Conclusion

In this paper, we have proposed a negotiation protocol which allows agents to make more than two offers per round and a negotiation strategy based on large (and eventually incomplete) preferences which does not assume that agents know the preferences of each other. In this way, we have proposed a fair negotiation process which does not give priority to one agent and which minimizes the maximum effort for one agent. We have applied our framework for distributed problem solving. In particular, we have considered the case of two agents in a maze which aims at negotiating a meeting in order to reach it as soon as possible. In our experiments, the negotiation improves the resolution of this problem and the path between the agents emerge from the repeated negotiations.

We are currently extending our bilateral negotiation game to a multi-party negotiation game with more than two agents. Even if our definitions for concessions and compromises are suitable, we adapt the negotiation strategies and the arbitration mechanism. In this way, we will allow more than two agents to negotiate a meeting point.

References

1. Fatima, S., Rahwan, I.: Negotiation and Bargaining. In: Multiagent Systems, pp. 143–176. MIT Press (2013)
2. Jennings, N.R., Faratin, P., Lomuscio, A.R., Parsons, S., Sierra, C., Wooldridge, M.: Automated negotiation: Prospects methods and challenges. *GDN* 10(2), 199–215 (2001)
3. Alpern, S., Gal, S.: *The Theory of Search Games and Rendezvous*. Kluwer (2003)
4. Aydođan, R., Baarslag, T., Hindriks, K.V., Jonker, C.M., Yolum, P.: Heuristic-based approaches for CP-nets in negotiation. In: Ito, T., Zhang, M., Robu, V., Matsuo, T. (eds.) *Complex Automated Negotiations*. *SCI*, vol. 435, pp. 115–126. Springer, Heidelberg (2012)
5. Rosenschein, J.S., Zlotkin, G.: *Rules of Encounter - Designing Conventions for Automated Negotiation among Computers*. MIT Press (1994)
6. Rubinstein, A.: Perfect equilibrium in a bargaining model. *Econometrica* 50(1), 97–102 (1982)
7. Endriss, U.: Monotonic concession protocols for multilateral negotiation. In: *Proc. of AAMAS*, pp. 392–399. ACM (2006)