



HAL
open science

Multi-party negotiation with preferences rather than utilities

Fabien Delecroix, Maxime Morge, Thomas Nachtergaelle, Jean-Christophe Routier

► **To cite this version:**

Fabien Delecroix, Maxime Morge, Thomas Nachtergaelle, Jean-Christophe Routier. Multi-party negotiation with preferences rather than utilities. *Multiagent and Grid Systems - An International Journal of Cloud Computing*, 2016, 12 (2), pp.27. 10.3233/MGS-160248 . hal-01327502

HAL Id: hal-01327502

<https://inria.hal.science/hal-01327502>

Submitted on 6 Jun 2016

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.

Multi-party Negotiation with Preferences Rather than Utilities

Fabien Delecroix*¹, Maxime Morge*^{† 1}, Thomas Nachtergaelle*¹, and
Jean-Christophe Routier*¹

¹*Univ. Lille, CNRS, Centrale Lille, UMR 9189 - CRISTAL - Centre de Recherche en Informatique Signal et Automatique de Lille, F-59000 Lille, France*

Abstract

We propose here a multi-party negotiation game where agents exchange offers over a single issue. Contrary to most of existing works, we capture the agents' rationality with incomplete (weak and partial) preferences rather than with utility functions. We prove that our negotiation process is always successful and reach fair agreements when the agents are cooperative. In order to illustrate how our negotiation game can tackle distributed problem solving, we introduce a toy problem: two or more agents are paratroopers landed in an unknown maze. Since they aim to meet as soon as possible, they negotiate a rendezvous during their exploration. Our experimentations show that the added value of negotiation is a fast resolution since this adaptive heuristic shares the efforts performed by the agents.

Keywords: Conflict handling, Negotiation, Distributed Problem Solving

*Prenom.Nom@univ-lille1.fr

[†]Corresponding author Phone: +333 20 33 77 12 / Fax: +333 28 77 85 37

1 Introduction

Muti-agent systems are suitable for analysing, designing and implementing systems composed of autonomous systems which interact. Conflict handling has become a subject of central interest in distributed artificial intelligence over the past years since it is required to design distributed services and systems and to support conflict resolution between humans. Conflict are first-class citizens in such situations since (human and software) agents may have epistemic and practical conflicts. Even if they are in competition, agents having conflict of interests must collaborate in order to reach their goals. For this purpose, they can negotiate. A negotiation is a process in which several stakeholders verbalize contradictory demands and then move towards agreements by concession-making or search of new alternatives [18]. The accommodation of the viewpoints of agents with conflicting interests (in the sense that they cannot be simultaneously fully satisfied) must lead to a mutually acceptable agreement, i.e. a plan of action which deals with a resource allocation, a matching between agents or, more generally a collective decision. In this perspective, the negotiation can be seen as a distributed search in a space of potential agreements [13].

Many negotiation scenarios have been considered in the literature [22, 21, 10, 23, 4, 2] with different domains (the set of issues being negotiated and the valid values for each issue), different contexts (previous similar negotiations can be used to predict the interlocutors' preferences or several interdependent negotiations are going on at the same time) and different multi-agent settings (the number and behaviour of participants). Despite this variety, the preference model which is the most widely used in the negotiation literature is the simplest form of cardinal utility function. The soundness and the relevance of this semantics still remain controversial in decision theory (see [17] for a recent survey on preference modelling) and in social choice theory (e.g. [20]). First, the numerical representation of utility leads to many mathematical aggregation operators which can be misinterpreted in a comparative perspective. Second, the utility refers to the intrinsic evaluation of alternatives rather than the comparative assessment of them. In this paper, we aim to fill the gap between (social) decision theory and negotiation practice.

We propose here an operational and distributed process to take a collective decision. For this purpose, we propose a multi-party simultaneous negotiation game. We capture the agents' rationality with incomplete (weak and partial) preferences over a single issue rather than with utility functions. In this way, some alternatives may be incomparable. Since we aim to provide a fair outcome, i.e. one which does not give priority to one agent and minimizes the maximum efforts of the agents, our protocol allows agents to offer more than two alternatives per round. Moreover, we adopt concession-based strategies where agents

first submit their preferred alternatives and, if that fails, they concede, i.e. they withdraw their expectation. We prove that our negotiation games always successfully end and that the agreements reached by cooperative agents are fair, i.e. minimize the maximum efforts of the agents. In order to illustrate how our negotiation game can tackle distributed problem solving, we introduce a toy problem: two or more agents are paratroopers landed in an unknown maze. Since they aim to meet as soon as possible, they negotiate a rendezvous during their exploration. Our experimentations show that the added value of negotiation is a fast resolution since this adaptive heuristic shares the efforts performed by the agents.

This paper is structured as follows. Section 2 introduces the basic notions of preferences in the background of our work. In Section 3, we define what is a fair collective decision based on individual preferences. Section 4 describes the core of our proposal, i.e. our negotiation game. Section 5 outlines a toy problem in order to illustrate how our negotiation game can tackle distributed problem solving. Section 6 overviews relevant related works and highlights the characteristics of our negotiation framework. Finally, Section 7 concludes with some directions for future work.

2 Preliminaries

In this section, we introduce the formalization of individual preferences and collective preferences in the background of our work¹.

In many cases, an agent has no strict preferences between two alternatives, either since he considers both of them are satisfying or since there is too much uncertainty to compare them. That is the reason why we envisage weak preferences.

Definition 1 (Preference relation) *Let \mathcal{X} be a non-empty finite set of alternatives. We call (weak) **preference relation** over \mathcal{X} , a preorder over \mathcal{X} (denoted \succsim). If $x \succsim y$, we say that “the alternative x is at least as good as y ”. From this weak preference relation, we can define:*

- the **strict preference relation**, denoted \succ , such that :

$$\forall (x, y) \in \mathcal{X}^2, x \succ y \Leftrightarrow (x \succsim y \wedge \neg(y \succsim x)) \quad (1)$$

$x \succ y$ is read as x is strictly preferred to y ;

¹Tab. 1 summarizes the set of symbols used in this paper.

- the **equivalence relation**, denoted \sim , such that:

$$\forall (x, y) \in \mathcal{X}^2, x \sim y \Leftrightarrow (x \succsim y \wedge y \succsim x) \quad (2)$$

$x \sim y$ is read as x is equivalent to y ;

- the **incomparability relation**, denoted \bowtie , such that:

$$\forall (x, y) \in \mathcal{X}^2, x \bowtie y \Leftrightarrow (\neg(x \succsim y) \wedge \neg(y \succsim x)) \quad (3)$$

$x \bowtie y$ is read as x and y are incomparable.

A preference relation is reflexive and transitive (i.e. a preorder). The equivalence relation between two alternatives indicates the agnosticism of the agent. The latter estimates that there is no real difference between them with respect to his tastes. A preference relation is complete if there is always a relation between two alternatives:

$$\forall (x, y) \in \mathcal{X}^2, (x \succsim y) \vee (y \succsim x) \quad (4)$$

By contrast, we consider here partial preference relations. In this way, some alternatives are incomparable. The incomparability can be interpreted as a way for the decision-making agent not to commit to an uncertain judgment. We can note that, contrary to the equivalence relation, the incomparability relation is not transitive.

When the preference relation is partial, the *optimal* alternative ($x \in \mathcal{X} \mid \forall y \in \mathcal{X}, x \succsim y$) may not exist. In order to decide, i.e. select the best alternatives, we define the non-dominance to distinguish, in a set of alternatives, the ones for which the agent prefers no other.

Definition 2 (Non-dominated alternatives) *Let (\mathcal{X}, \succsim) be a set of alternatives associated with a preference relation. The **non-dominated alternatives** are defined such that:*

$$c_{ND}(\mathcal{X}, \succsim) = \{x \in \mathcal{X} \mid \forall y \in \mathcal{X}, \neg(y \succ x)\} \quad (5)$$

We can remark that there is always at least one non-dominated alternative. Moreover, any optimal alternative is also non-dominated. The reverse is not necessarily true.

We consider now a set of agents (two or more) taking a collective decision. Given the individual preferences of the agents, the quality of a collective decision, e.g. the agreement reached by a negotiation, can be evaluated in different ways. A consensus is an alternative which is one of the best for all the agents.

Definition 3 (Consensus) Let \mathcal{X} be a non-empty finite set of alternatives and Ω a set of agents such that each agent $i \in \Omega$ is associated with a preference relation \succsim_i over \mathcal{X} . A **consensus** is an alternative x such that:

$$x \in \bigcap_{i \in \Omega} c_{ND}(\mathcal{X}, \succsim_i) \quad (6)$$

Such a decision, which is in the intersection of the non-dominated alternatives, can be taken if there is no conflict. Therefore, a consensus may not exist.

The Pareto rule is a less restrictive approach which allows to aggregate the individual preferences, even if they are conflicting, in order to compare the alternatives from the viewpoint of the group. 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 . Since we consider incomplete (weak and partial) preferences, we adopt here the possible and necessary Pareto dominance introduced in [6].

Definition 4 (possible/necessary Pareto dominance) Let \mathcal{X} be a non-empty finite set of alternatives and a set of agents Ω where each agent $i \in \Omega$ is associated with a preference relation \succsim_i over \mathcal{X} .

- The **necessary Pareto preference** ($\succ^{NP} \subseteq \mathcal{X} \times \mathcal{X}$) is defined such that $\forall (x, y) \in \mathcal{X}^2$:

$$x \succ^{NP} y \Leftrightarrow [(\forall i \in \Omega, x \succsim_i y) \wedge (\exists i \in \Omega, x \succ_i y)] \quad (7)$$

- The **possible Pareto preference** ($\succ^{PP} \subseteq \mathcal{X} \times \mathcal{X}$) is defined such that $\forall (x, y) \in \mathcal{X}^2$:

$$x \succ^{PP} y \Leftrightarrow [(\forall i \in \Omega, \neg(y \succ_i x)) \wedge (\exists i \in \Omega, \neg(y \succsim_i x))] \quad (8)$$

- An alternative x is a **necessary Pareto-optimum** for Ω iff:

$$\forall y \in \mathcal{X}, \neg(y \succ^{PP} x). \quad (9)$$

- An alternative is a **possible Pareto-optimum** for Ω iff:

$$\forall y \in \mathcal{X}, \neg(y \succ^{NP} x). \quad (10)$$

The necessary Pareto preference is a strict extension of the Pareto rule. Indeed, if two alternatives are incomparable for at least one agent, then the alternatives will be considered Pareto-incomparable beyond a reasonable doubt. By contrast, the possible Pareto preference

is a more liberal extension. Indeed, the possible Pareto-dominance only requires that one agent removes the incomparability. Note that the necessary Pareto preference is a restriction of the possible Pareto preference ($\succ^{NP} \subseteq \succ^{PP}$). Similarly, there are two kinds of Pareto-optimality. A necessary Pareto-optimum is not possibly Pareto-dominated. By symmetry, a possible Pareto-optimum is not necessary Pareto-dominated. Contrary to the necessary Pareto-optimum, there is always a possible Pareto-optimum. That is the reason why we will only consider the set of possible Pareto-optima (denoted $c_{ND}(\mathcal{X}, \succ^{PP})$).

Example 1 *We consider here \succsim_1, \succsim_2 and \succsim_3 the preference relations of three agents over the set of alternatives $\mathcal{X} = \{a, b, c, d, e, f\}$. The associated preference graphs are represented in Fig. 1 as directed graphs where a node represents an alternative and there is an edge from x to y when x is at least as good as y . We can note that the non-dominated alternatives for each agent are $c_{ND}(\mathcal{X}, \succsim_1) = \{a, b, c\}$, $c_{ND}(\mathcal{X}, \succsim_2) = \{c\}$ and $c_{ND}(\mathcal{X}, \succsim_3) = \{a, b\}$. The set of possible Pareto-optima is $c_{ND}(\mathcal{X}, \succ^{PP}) = \{a, b, c, d, e\}$.*

3 Fair agreement

In this section, we aim to define what is a fair collective decision based on individual preferences which are incomplete (weak and partial).

Inspired by [12], we consider that an agent has several levels of discrimination. Therefore, we introduce a numerical value which measures the efforts performed by the agent for taking an alternative for granted.

Definition 5 (Rank) *Let (\mathcal{X}, \succsim) be a set of alternatives associated with a preference relation over \mathcal{X} . The **rank** of an alternative x is defined such that:*

$$r(x, \mathcal{X}, \succsim) = \begin{cases} 1 & \text{if } x \in c_{ND}(\mathcal{X}, \succsim) \\ 1 + r(x, \mathcal{X} \setminus c_{ND}(\mathcal{X}, \succsim), \succsim) & \text{otherwise} \end{cases} \quad (11)$$

Roughly speaking, the rank of an alternative corresponds to its depth in the preference graph. The smaller the rank is, the more satisfactory the alternative is. By convention, the rank of a non-dominated alternative is 1. As [12], our metric does not conform to the axiom of independence of irrelevant alternatives (eliminating some of the unchosen alternatives should affect the selection of the best option). In other words, we assume here a well-posed decision problem. Finally, our metric is compliant with a partial preference relation.

In order to determinate the alternatives which are fair, we adopt here the leximax rule.

Definition 6 (leximax preference relation) Let \mathcal{X} be a non-empty finite set of alternatives, Ω a set of n agents and $x, y \in \mathcal{X}$ two alternatives. We note $\vec{r}_\Omega(x) = (r_1(x), \dots, r_n(x))$ the vector of ranks in decreasing order for the agents in Ω . We say that:

- the alternative x is strictly leximax-preferred to y (denoted $x \succ^{lex,r} y$) iff:

$$\exists k \leq n, \forall i < k, r_i(x) = r_i(y) \text{ and } r_k(x) < r_k(y) \quad (12)$$

- the alternative x is leximax-equivalent to y (denote $x \sim^{lex,r} y$) iff:

$$\forall i \leq n, r_i(x) = r_i(y) \quad (13)$$

It is worth noticing that all the agents consider the same set of alternatives. The leximax rule refines the minimax rule by generalizing the principle in case of ties to the second worst-off agent and so on (similarly to the leximin). We can notice that $(\succsim^{lex,r}, \mathcal{X})$ is a total preorder. Indeed, this collective preference relation allows to compare any couple of alternatives in order to rank them.

A collective decision is fair if the efforts performed by the agents are splitted in a leximax-optimal way.

Definition 7 (Fair alternatives) Let \mathcal{X} be a set of alternatives and $\Omega \neq \emptyset$ a set of agents. Let $(\succsim^{lex,r}, \mathcal{X})$ be the leximax preference relation for Ω . The set of **fair alternatives** in \mathcal{X} for Ω is $c_{ND}(\mathcal{X}, \succsim^{lex,r})$.

All the fair alternatives are Pareto-optima (cf Def. 4).

Proposition 1 (possible Pareto-optimality of fair alternatives)

$$c_{ND}(\mathcal{X}, \succsim^{lex,r}) \subseteq c_{ND}(\mathcal{X}, \succ^{PP}) \quad (14)$$

Proof 1 Let x be a fair alternative in \mathcal{X} for Ω . We prove by contradiction that x is a possible Pareto-optimum for Ω . We assume that x is not a possible Pareto-optimum, i.e. there always exists an alternative $y \in \mathcal{X}$ which necessary Pareto-dominates x . By Def. 4, we deduce that (1) $\forall i \in \Omega, y \succsim_i x$ and (2) there is an agent $j \in \Omega$ such that $y \succ_j x$. From (1) and Def. 5, we deduce that $\forall i \in \Omega, r(x, \mathcal{X}, \succsim_i) \geq r(y, \mathcal{X}, \succsim_i)$. From (2) and Def. 5, we deduce there is an agent $j \in \Omega$ such that $r(x, \mathcal{X}, \succsim_j) > r(y, \mathcal{X}, \succsim_j)$. From Def. 6, we deduce that $y \succ^{lex,r} x$ and so x is not a fair alternative which is a contradiction with our assumption.

A group of agents can always take a fair decision.

Proposition 2 (The existence of a fair alternative) *Let \mathcal{X} be a non-empty set of alternatives and Ω a set of agents. Let $(\succsim^{lex,r}, \mathcal{X})$ be the leximax preference relation for Ω . There is always an alternative $x \in \mathcal{X}$ such that $x \in c_{ND}(\mathcal{X}, \succsim^{lex,r})$.*

Proof 2 *The leximax preference relation is acyclic over a non-empty finite set and so, this set contains at least one maximal element.*

In our example, we observe that the fair alternatives are Pareto-optima but the reverse is not true.

Example 2 (Fair alternatives) *Let us consider the preference relations of the three agents represented in Fig. 1. Tab. 4 shows the corresponding ranks and the vector of ranks in decreasing order. The alternatives d and e are fair and, as said previously, possible Pareto-optima. c is also a possible Pareto-optimum but this alternative is not fair.*

In summary, a group of agents can always take a fair decision. Intuitively, such a collective decision aims to split the efforts in a fair way by minimizing the efforts of the worst-off agents and in case of tie by minimizing the second worst-off agent, and so on. Moreover, this collective decision rule is Pareto-inclusive, i.e. all the fair alternatives are Pareto-optima².

4 Negotiation

We present in this section an operational and distributed process to take a collective decision. For this purpose, we propose a multi-party simultaneous negotiation game. Firstly, we introduce the negotiation protocol, the acceptability criteria and two concrete bidding strategies. Then, we present our theoretical results and several empirical ones.

4.1 Negotiation game

Our **protocol** captures the rules of encounter.

Definition 8 (Protocol) *Let \mathcal{X} be a set of alternatives and $i \in \Omega$ one of the participant. At each round k , the utterance of the negotiating agent i is such that:*

²It is worth noticing that, inspired by the Social Choice Theory, we can define the egalitarian rank ($r_e(x, \mathcal{X}, \Omega) = \max_{i \in \Omega} r(x, \mathcal{X}, \succsim_i)$) in order to evaluate the worst-off agent. However, the collective decision rule which consists of minimizing the egalitarian rank is not Pareto-inclusive.

1. the agent can propose ($u_{i,k} = \text{propose}(i, \Omega, A)$) or accept ($u_{i,k} = \text{accept}(i, \Omega, A)$) a set of alternatives ($A \subseteq \mathcal{X}$). We note $p_i(k)$ the set of alternatives proposed by agent i at round k and $\vec{k}_\Omega(a)$ the vector of rounds in decreasing order during which the alternative a has been proposed. The set of alternatives proposed by agent i before the round $k \geq 1$ is denoted $\mathcal{P}_i(k) = \bigcup_{l < k} p_i(l)$;

2. the repetition of an alternative in the proposals of the agent i is forbidden,

$$u_{i,k} = \text{propose}(i, \Omega, A) \text{ if } A \cap \mathcal{P}_i(k) = \emptyset; \quad (15)$$

3. the agent can keep silent if he has already proposed all the alternatives,

$$u_{i,k} = \text{propose}(i, \Omega, \emptyset) \text{ if } \mathcal{P}_i(k) = \mathcal{X}; \quad (16)$$

4. the end occurs when at least one alternative has been proposed by all the agents,

$$u_{i,k} = \text{accept}(i, \Omega, A) \text{ with } A \neq \emptyset \text{ if } A \subseteq \bigcap_{i \in \Omega} \mathcal{P}_i(k) \quad (17)$$

Our protocol is symmetric and so it does not give priority to the initiator of the negotiation. Moreover, the protocol allows an agent to submit multiple offers at each round since he may not always decide between them.

The **acceptability criteria** of the agents is computed by Algorithm 1. Roughly speaking, the agents consider as acceptable only the alternatives which have been proposed by all the participants. Moreover, it gives priority to the alternatives which appear earlier in the history. At each round k , the agents compute the vectors of rounds for each alternative ($\vec{k}_\Omega(x)$, $x \in \mathcal{X}$). When at least one alternative has been proposed by all the agents, the agents compute the corresponding leximax-optima (cf Algorithm 2) and accept them: this is an *agreement*³. Otherwise, each agent makes a proposal in accordance with his strategy.

The content of a proposal, which may contain several alternatives, is determined by the **bidding strategy** which interfaces with the protocol through the condition mechanism of utterance and interpretation of the dialogical moves⁴.

In accordance with our protocol, the bidding strategy must be concession-based. The agent first submits his preferred alternatives. If that fails, the agent concedes, which consists

³An agreement is not necessarily unique. In order to reach a single collective decision, the agents can select together one alternative by using a (pseudo-)randomized function. We will see in Section 4.3, it is not required in most of practical situations.

⁴Obviously, when an agent receives a proposal, he updates his representation of the history.

ALGORITHM 1: Computation of the acceptability criteria

Data: the set of alternatives \mathcal{X} , the vector of rounds $(\vec{k}_\Omega(x))_{x \in \mathcal{X}}$ **Input:** the next round k , the previous offers $(p_i(k-1))_{i \in \Omega}$ $Commons := \emptyset;$ **foreach** $x \in \mathcal{X}$ **do** update the vector of rounds **foreach** $i \in \Omega$ **do** **if** $x \in p_i(k-1)$ **then** $\vec{k}_\Omega(x).push(k)$; **end** **end****end** $Commons := \{x \in \mathcal{X} \mid card(\vec{k}_\Omega(x)) = card(\Omega)\}$ // alternatives proposed by all the agents**if** $Commons \neq \emptyset$ **then** compute the agreement $Agreements := leximax(Commons, \{\vec{k}_\Omega(x) \mid x \in Commons\});$ $accept(i, \Omega, Agreements)$ **end****else** biddingStrategy() ;

ALGORITHM 2: Computation of the alternatives having a leximax-optimal vector of rounds

Data: a set of alternatives A , the vector of rounds $(\vec{k}_\Omega(x))_{x \in A}$ **for** $l := 1$ **to** $card(\Omega)$ **do** $O := \operatorname{argmin}_{x \in A} \{\vec{k}_\Omega(x).[l]\}$; **if** $card(O) = 1$ **then return** O ;**end****return** O

of the withdrawal of its expectation, and so he submits his unreleased preferred alternatives. It is worth noticing that contrary to the most existing concession strategies, this principle does not required perfect information about the opponent’s preferences. In order to apply this kind of egocentric concession [8], several concrete strategies for agent i are compliant with our protocol. We highlight here two strategies: i) the first one, called “Cooperative”, is suitable for agents having a common goal; ii) the second one, called “Temporizing” is suitable for agents in competition.

The *cooperative strategy* for our multi-party negotiation protocol (denoted C) is described by Algorithm 3. If the agent interacts in order to reach a fair agreement, then he will propose during the round k all the alternatives with the rank k . In other words, the agent proposes at each round all the non-dominated alternatives among those he has not released.

ALGORITHM 3: Cooperative strategy of agent i

Data: The set of alternatives \mathcal{X}

$Unreleased := \mathcal{X} \setminus \mathcal{P}_i(t);$

$ND := c_{ND}(Unreleased, \succsim_i);$

$propose(i, \Omega, ND);$

The *temporizing strategy* is more suitable for a selfish agent. The latter does not care about the social virtue of the agreement but only about his preferred alternatives. Algorithm 4 is the first filter for the temporizing strategy. As the cooperative strategy, this strategy should be legal: the alternatives must not be repeated in his proposals. Moreover, it is irrational to propose an alternative which is dominated by an unreleased one. Therefore, the strategy contents itself to the non-dominated alternatives among those remaining. While the cooperative strategy allows to submit multiple offers at each round, a temporizing agent proposes only one. For this purpose, we introduce a function called *uniqueChoice* which returns a singleton (cf Algorithm 5). Due to the restriction over the non-dominated alternatives among the remaining ones, the agent minimizes at each round the depth level in the preference graph among the nodes in *Unreleased*. In other words, the agent only proposes the remaining alternatives with a minimal rank⁵.

In brief, the temporizing strategy (denoted T) combines: i) the selection of rational alternatives; and ii) the choice of a single alternative with a minimal rank.

⁵It is worth noticing that we have also explored more sophisticated strategy, e.g. selecting the most recent alternatives in the proposals of his interlocutors. However, this kind of refinements does not affect neither the theoretical nor the experimental results presented in the next section.

ALGORITHM 4: Temporizing strategy of agent i

Data: The set of alternatives \mathcal{X}

$Unreleased := \mathcal{X} \setminus \mathcal{P}_i(t)$;

$ND := c_{ND}(Unreleased, \succsim_i)$;

$E := uniqueChoice(ND)$;

$propose(i, \Omega, E)$;

ALGORITHM 5: The function $uniqueChoice$ of agent i

Data: the set of rational alternatives Rat

Output: a set of alternatives

$Min := \operatorname{argmin}_{x \in Rat} \{r(x, \mathcal{X}, \succsim_i)\}$;

return $\{random(Min)\}$;

4.2 Theoretical results

We demonstrate here some properties of the agreements reached by our negotiation game.

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

Proposition 3 (Success of the negotiation) *Our negotiation game ends (successfully).*

Proof 3 *Since the set of alternatives is finite and the protocol forbids the silence and the repetition of alternatives in the proposals of a given participant, at least one alternative is proposed by all the agents. Therefore, the latter is accepted and the end occurs. The set of alternatives with the leximax relation over the vector of rounds is a total preorder over a non-empty finite set and so, this set contains at least one maximal element, i.e. an agreement.*

Since the agents can adopt different strategies, the negotiation process can be parameterized with several game configurations.

Definition 9 (Strategy profile) *Let \mathcal{X} be a set of alternatives and Ω a set of n agents. The **strategy profile** S_Ω is the tuple $S_\Omega = (S_1, \dots, S_n)$ where S_i is the concrete strategy for agent $i \in \Omega$ ($S_i \in \{C_i, T_i\}$). The corresponding set of agreements is denoted $\sigma(S_\Omega)$.*

For the sake of brevity, we denote:

- C_Ω when all the agents adopt the cooperative strategy (C_1, C_2, \dots, C_n) ;
- T_Ω when all the agents adopt the temporizing strategy (T_1, T_2, \dots, T_n) ;
- $C_{\Omega-i}$ when all the agents adopt the cooperative strategy excepted agent i who adopts the temporizing one $(C_1, \dots, C_{i-1}, T_i, C_{i+1}, \dots, C_n)$;

- $T_{\Omega-i}$ when all the agents adopt the temporizing strategy excepted agent i who adopts the cooperative one $(T_1, \dots, T_{i-1}, C_i, T_{i+1}, \dots, T_n)$;

We can note that our approach reflects two basic axioms of social choice theory: anonymity and neutrality. Actually, negotiators are treated equally and our collective decision procedure does not have a built-in bias for or against any alternatives.

In the following example, the negotiating agents reach a fair agreement when they all adopt the cooperative strategy. However, the agreement may be not a possible Pareto-optimum when the third agents temporizes.

Example 3 (Negotiation between three agents) *Let us consider the preference relations of three agents represented in Fig. 1.*

The left side of Tab. 5 shows the trace of a negotiation where they adopt the cooperative strategy. At round $k = 1$, agent 1 starts with the alternatives a, b and c , agent 2 proposes the alternative c and agent 3 proposes the alternatives a and b . Since none alternative has been proposed by all the agents, a second round takes place. At the third round, the alternatives d, e and f have been proposed and so the end occurs. The agreements are the alternatives for which the vectors of rounds are leximax-optima : d and e .

The right side of Tab. 5 shows the trace of a negotiation where the first two agents are cooperative while the third one temporizes. The cooperative agents have the same behavior as previously, and so make the same proposals. The third agent who is temporizing submits only one alternative per round. Indeed, he can propose either the alternative a or the alternative b at the first round. Then, he submits the one which has not be released. At the third round, whatever the alternative he submits (e or f), this alternative has been also proposed by the two other agents, and so the negotiation game is over. The last proposal of agent 3 is the agreement reached by the negotiation: either e , or f . We remark that the alternative d , which is worst than e and f for agent 3, is no more in the set of agreements. If agent 3 temporizes, he will reach a better agreement. It is worth noticing that the alternative f , which can be an agreement, is not a possible Pareto-optimum.

A negotiation game between cooperative agents reach all the fair alternatives.

Theorem 1 (Fair agreement) *Let \mathcal{X} be a set of alternatives and Ω a set of agents.*

$$\sigma(C_\Omega) = c_{ND}(\mathcal{X}, \succsim^{lex,r}) \quad (18)$$

Since cooperative agents share the efforts in the negotiation game, a fair process leads to a fair agreement.

Lemma 1 *Let $x \in \mathcal{X}$ be an alternative and θ the last round. If $r(x, \mathcal{X}, \succsim_i) \leq \theta$ then $x \in \mathcal{P}_i(r(x, \mathcal{X}, \succsim_i))$.*

Proof 4 *Proof by induction.*

Base case: *At round 1, cooperative agent i submits the non-dominated alternatives (cf Algorithm 3), i.e. all the alternatives $x \in \mathcal{X}$ such that $r(x, \mathcal{X}, \succsim_i) = 1$ (cf Def 5).*

Induction step: *let the round k be given and suppose the cooperative agent has submitted all the alternatives $x \in \mathcal{X}$ such that $r(x, \mathcal{X}, \succsim_i) \leq k$.*

Then, cooperative agent i submits all the non-dominated alternatives which are legal at round $k+1 < \theta$. Formally, he submits $p_i(k+1) = c_{ND}(\mathcal{X} \setminus \mathcal{P}_i(k+1), \succsim_i)$. By induction hypothesis, $p_i(k+1) = c_{ND}(\mathcal{X} \setminus \{x \in \mathcal{X} \mid r(x, \mathcal{X}, \succsim_i) \leq k\}, \succsim_i)$. Due to Algo. 3, we deduce that cooperative agent i has submitted at round $k+1$ or previously all the $x \in \mathcal{X}$ such that $r(x, \mathcal{X}, \succsim_i) \leq k+1$.

Conclusion: *agent has submitted all the alternatives $x \in \mathcal{X}$ such that $r(x, \mathcal{X}, \succsim_i) \leq \theta$ at round θ .*

Proof 5 (Theorem 1) (a) *We prove by contradiction that $\sigma(C_\Omega) \subseteq c_{ND}(\mathcal{X}, \succsim_{lex,r})$. Let $x \in \sigma(C_\Omega)$ be an agreement. We assume that $x \notin c_{ND}(\mathcal{X}, \succsim_{lex,r})$. From Definition 6, there exists an alternative $y \in \mathcal{X}$ such that $y \succ_{lex,r} x$. From lemma 1 and Algorithm 1, we deduce $x \notin \sigma(C_\Omega)$ which is in contradiction with our assumption.*

(b) *We prove by contradiction that $c_{ND}(\mathcal{X}, \succsim_{lex,r}) \subseteq \sigma(C_\Omega)$. Let $x \in c_{ND}(\mathcal{X}, \succsim_{lex,r})$ be a possible Pareto-optimum. We assume that $x \notin \sigma(C_\Omega)$. From Proposition 3, there exists agreement $y \in \sigma(C_\Omega)$. From Theorem 1 (a), we deduce $y \in c_{ND}(\mathcal{X}, \sim_{lex,r})$. According to Definition 6, $\vec{r}_\Omega(x) = \vec{r}_\Omega(y)$. Lemma 1 leads to $\vec{k}_\Omega(x) = \vec{r}_\Omega(y)$. From Algorithm 1, we deduce $x \in \sigma(C_\Omega)$ which is in contradiction with our assumption.*

It is worth noticing that, as illustrated in our previous example, an agreement may be not a possible Pareto-optimum if at least one agent temporizes.

We demonstrate here that the temporizing strategy is more suitable than the cooperative for a selfish agent. For this purpose, we fix the strategies of the interlocutors and the (pseudo-)random function (*random* in Algorithm 5). Indeed, we compare the agreements with respect to the strategy adopted by the agent, all other things being equal.

When all his interlocutors cooperate, the agent can never get a better deal by being cooperative rather than temporizer.

Theorem 2 (Non-dominance of the cooperative strategy) *Let \mathcal{X} be a set of alternatives and $i \in \Omega$ an agent. If the (pseudo-)random function for agent i is arbitrarily fixed then:*

$$\forall x \in \sigma(C_\Omega) \setminus \sigma(C_{\Omega-i}), \nexists y \in \sigma(C_{\Omega-i}), (x \succ_i y) \quad (19)$$

Proof 6 Let $x \in \sigma(C_\Omega) \setminus \sigma(C_{\Omega-i})$ be an agreement reached by cooperative agents which is not reached when agent i temporizes. We prove by contradiction that $\forall y \in \sigma(C_{\Omega-i}), \neg(x \succ_i y)$. Let $y \in \sigma(C_{\Omega-i})$ be an agreement reached when all the agents adopt the cooperative strategy expected agent i who adopts the temporizing one. We assume that $x \succ_i y$. We know that: i) in the game configuration C_Ω , all the agents have proposed x ($\exists k > 0, x \in \mathcal{P}_i(k) \wedge x \in \bigcap_{j \neq i} \mathcal{P}_j(k)$) since x is an agreement; ii) in the game configuration $C_{\Omega-i}$, all the agents have proposed y ($\exists k' > 0, y \in \mathcal{P}_i(k') \wedge y \in \bigcap_{j \neq i} \mathcal{P}_j(k')$) since y is an agreement. Each cooperative agent j makes the same proposals in the two game configurations. Since $x \notin \sigma(C_{\Omega-i})$, i proposes y before x ($\exists k' > 0, y \in \mathcal{P}_i(k') \wedge x \notin \mathcal{P}_i(k')$). This last assertion is in contradiction with our assumption ($x \succ_i y$).

4.3 Empirical results

Our experimentations aim to evaluate: i) in which practical conditions the temporizing strategy is more suitable for a selfish agent and; ii) if a random function is required in order to reach a single decision at the end of the negotiation process. We first introduce the experimental technicalities and then we present the experimental results with their analysis.

In our experimentations, the pseudo-random generation of a preference relation is parameterized by:

- $card(\mathcal{X})$, the number of alternatives;
- $minEqRa$, the minimal ratio of equivalent couples ($0 \leq minEqRa \leq 1$);
- $maxIncRa$, the maximal ratio of incomparable couples ($0 \leq maxIncRa \leq 1$).

From these parameters, we generate a preference graph (cf Algorithm 6). Initially, we consider a graph without any edge and we add: i) some equivalence couples until the rate of equivalent couples is greater or equal to $minEqRa$ and; ii) some preference couples until the rate of incomparable couples is smaller or equal to $maxIncRa$. We note that $minEqRa$ and $maxIncRa$ are both threshold. Obviously, we generate preference relations which are acyclic.

We consider 100 agents whose preferences have been pseudo-randomly generated. An experience is a set of 4 negotiation games : $C_\Omega, C_{\Omega-i}, T_\Omega$ and $T_{\Omega-i}$. In order to compare them, we fix the random function of temporizer agents by using independent lexicographic orders. The parameters for preference generation are: $card(\mathcal{X}) = 25$, $minEqRa \in \{0, 0.1, 0.2\}$ and $maxIncRa \in [0, 0.6]$. For each set of parameters, we generate 1000 runtimes by generating a new preference relation.

ALGORITHM 6: Pseudo-random generation of a preference relation

Data: the preference graph (\mathcal{X}, E) with the nodes \mathcal{X} and the edges E , the minimal ratio of equivalent couples $minEqRa$, the maximal ratio of incomparable couples $maxIncRa$
 $nbCouples := card(\mathcal{X}) \times (card(\mathcal{X}) - 1)$;
 $eqRa := card(\{(x, y) \in \mathcal{X}^2 | (x \sim y)\}) / nbCouples$;
while $eqRa < minEqRa$ **do** add equivalent couples
 $Inc := \{(x, y) \in \mathcal{X}^2 | (x \bowtie y)\}$;
 $(a, b) := random(Inc)$;
 $E := E \cup \{(a, b)\} \cup \{(b, a)\}$;
 $eqRa := card(\{(x, y) \in \mathcal{X}^2 | (x \sim y)\}) / nbCouples$;
end
 $incRa := card(\{(x, y) \in \mathcal{X}^2 | (x \bowtie y)\}) / nbCouples$;
while $eqRa > maxEqRa$ **do** add preference couples
 $Inc := \{(x, y) \in \mathcal{X}^2 | (x \bowtie y)\}$;
 $(a, b) := random(Inc)$;
 $E := E \cup \{(a, b)\}$;
 $incRa := card(\{(x, y) \in \mathcal{X}^2 | (x \bowtie y)\}) / nbCouples$;
end

Inspired by [5], we generalize the preference relation over set of alternatives in order to compare the set of agreements in the different game configurations⁶.

Definition 10 (Preference relation over set) *Let (\mathcal{X}, \succ) be a set of alternatives associated with a preference relation. Let A and B be two set of alternatives ($A, B \subseteq \mathcal{X}$). A is **strictly preferred** to B (denoted $A \succ B$) iff:*

$$\begin{cases} \exists a \in A \setminus B, \forall b \in B \setminus A \ a \succ b \text{ and} \\ \nexists b \in B \setminus A, \exists a \in A \setminus B \ b \succ a \end{cases} \quad (20)$$

Fig. 2(a) shows the outputs when all the interlocutors of agent i cooperate. In this case, the temporizing strategy is better than the cooperative one ($\sigma(C_{\Omega-i}) \succ_i \sigma(C_{\Omega})$) in more than 65 % of the runtimes. The dominance of T over C is maximal for the minimal equivalence ratio of 0.2 and the incomparability ratio close to 0.1. For this set of parameters, the temporizing agent postpones the agreement by submitting the alternatives with a low rank, while the interlocutors cooperate and submit several alternatives at each round. When the uncertainty of the agents is higher, the dominance of strategy T is less salient. This is due to the fact that the comparison of the agreements requires that the alternatives are themselves comparable (cf Def. 10).

Fig. 2(b) shows the outputs when all the interlocutors of agent i temporize. In this

⁶We have tested several extensions of the preference relation over sets. Its definition has no influence on the experimental results.

case, the temporizing strategy is better than the cooperative one ($\sigma(T_\Omega) \succ_i \sigma(T_{\Omega-i})$) in less than 5 % of the runtimes. The impact of the temporizing strategy is reduced because the interlocutors submit only one alternative per round.

Fig. 3(a) shows the percentage of runtimes where the agreements are fair if only one agent temporizes ($\sigma(C_{\Omega-i})$). In the worst case, the agreements are fair only in about 20 % of the runtimes. As soon as one agent temporizes, the fair agreements are infrequent. Fig. 3(b) shows the percentage of runtimes where the agreements are fair if all the agents temporize ($\sigma(T_\Omega)$). We observe that the agreements are fair only in about 15 % of the runtimes.

In Fig. 4, we compare, for each experiment, the average cardinality of non-dominated alternatives, i.e. the outcome of single-agent decision-making and the average cardinality of the set of fair agreements, i.e. the outcome of the negotiation between cooperative agents (with the same ratio of equivalence and incomparability). At the individual level, the more alternatives with *ex æquo* rank - incomparable or equivalent - we have the higher the cardinalities are, and so the more arbitrariness is required. At the collective level, with 100 agents and 25 alternatives, the set of agreements is, in most cases a singleton. Therefore, the agents do not need to randomly select together a “final” decision.

In summary, our experimentations show that the dominance of the temporizing strategy requires very restricted conditions. When all the other agents cooperate with suitable equivalence rate and incomparability rate, the agent can expect to reach a better deal. However, when a significant part of the agents temporize, the impact of the temporizing strategy is limited. In the meanwhile, the adoption of the temporizing strategy by few agents is sufficient to lose the fairness of the agreements. Finally, the outcome of a negotiation is less arbitrary than the outcome of single-agent decision-making, since the uncertainty of one agent can be balanced by the preferences of his interlocutors, even if the outcomes are different due to the concessions.

5 Usecase

In this section, we introduce a toy problem in order to illustrate how our negotiation game can tackle distributed problem solving. Landed in an unknown maze, paratroopers aim to meet as soon as possible and so, they communicate to fix a rendezvous point. Even if they have a common goal, these agents have conflicting interests since they want to move as few as possible. For this purpose, the agents negotiate by exchanging offers, i.e. potential cells where to meet.

First, we introduce the details of our scenario and how we apply our negotiation game. Then, we present our experimentations which compare negotiation with other heuristics and

we study the impact of information-sharing.

5.1 Scenario

We consider that the n agents in Ω are paratroopers landed in a maze. This environment is represented by a connected⁷ graph $G = (V, E)$ where each vertex in V is a cell and there is an edge in E between two vertices if the corresponding cells are adjacent and not separated by a wall. In our scenario, each agent has a local perception of the environment: he discovers the walls during the exploration. Formally, the agent i knows at each simulation step t its location $x_i(t) \in V$ and it has an internal representation of the environment $G_i(t) = (V, E_i(t))$. Initially, an agent knows the topology of the environment but it has no knowledge about the walls, i.e. about the lack of an edge between two adjacent vertices. An agent i assumes that initially there is no wall. Then in $E_i(0)$ all the edges between two adjacent vertices exist. An agent i perceives the walls in his perception halo. If this halo is restricted to the current cell, then

$$E_i(t+1) = E_i(t) - \{(x_i(t), y) \in E_i(t) \mid (x_i(t), y) \notin E\} \quad (21)$$

Based on his local perception of the environment, the agent i is able to estimate the distance towards any cell y at step t . This distance (denoted $d_i(y, t, E_i)$) is the length of the shortest path to reach y in $E_i(t)$. $d_i(y, t, E)$ denotes the real length of the shortest path for the agent i to reach y (when he knows all the walls). At each simulation step t , the agent i can stand still or move to an adjacent cell if there is no wall which prevents this move:

$$x_i(t+1) \in \{y \in V \mid d_i(y, t, E) \leq 1\}. \quad (22)$$

Since they want to move as few as possible, the preference relation of an agent (cf Def. 1) depends on his estimations of the distances. At each simulation step t , the preference relation of the agent i (denoted $\succsim_{d,t,i}$) is defined such that:

$$\forall (x, y) \in V^2, x \succsim_{d,t,i} y \Leftrightarrow d_i(y, t, E_i) \geq d_i(x, t, E_i) \quad (23)$$

We note that this preference relation is weak and complete. The corresponding set of non-dominated alternatives is denoted $c_{ND}(V, \succsim_{d,t,i})$.

The agents want to meet as soon as possible even if they do not know *a priori* the preferences of the other agents. That is the reason why we can consider that our domain is unpredictable. Since [3] has pointed the low importance of the opponent model, we focus on the bidding strategy and acceptance criteria.

⁷We assume that all the cells are accessible.

In order to capture the optimality of a rendezvous, we aim to minimize the number of simulation steps, i.e. the maximum number of moves performed by one agent and, in case of equality, we want to decrease the number of moves performed by the second agent making the more moves, and so on. As argued in Section 2, the leximax rule is suitable since it captures the idea of fairness by balancing the efforts performed by the agents. We denote $\vec{d}_\Omega(x, t) = (d_1(x, t, E_1), \dots, d_n(x, t, E_n))$ the vector of the estimated distances towards x in decreasing order at step t and $c_{ND}(V, \succsim^{lex, d, t})$ the corresponding set of leximax-optimal cells (cf Definition 6). Indeed, the computation of the optimal rendezvous requires the knowledge of the whole maze, i.e. its walls, by taking into account $d_i(y, t, E)$. Under this assumption, agents can compute the shortest path between them and so, the optimal rendezvous.

Even if all the cells are accessible, the maze may be imperfect. Since the path between two cells is not necessarily unique, the agents have to fix a rendezvous point. Otherwise they can turn around and never meet. Contrary to the classical rendezvous problem [1], agents are allowed to communicate. Each agent i can inform the other participants about the existence of walls ($tell(i, \Omega, wall(c_i, c_j))$ with $c_i, c_j \in V$) or they can propose/accept potential rendezvous ($propose(i, \Omega, C)$ or $accept(i, \Omega, C)$ with $C \subseteq V$). Obviously, the estimated distances and so the negotiation outcome depends on the partial knowledge of the agents, i.e. their perception of the environment.

Example 4 *In order to illustrate this problem, we consider a 3×3 maze (cf Fig. 5). At the first simulation step, Alice is on c_1 while Bob is on c_9 . We consider that the perception halo consists of one cell and so, Alice does not perceive any wall while Bob sees the wall between c_8 and c_9 . Since Alice does not perceive any wall, she estimates that the distance towards the cell c_7 is 2. On the other side, Bob estimates the distance towards this cell is 4.*

5.2 Rendezvous negotiation

In order to negotiate a rendezvous, we adopt the operational and distributed process presented in Section 4.1. Since we consider distributed problem solving, we assume that agents are collaborative and so they adopt the cooperative strategy in order to reach a fair agreement. Even if all the agents share the same strategy, they have their own preferences since their location and perception are different.

According to our concrete bidding strategy, the agent i makes at round k after the simulation step t the following proposals:

$$propose(i, \Omega, \{y \in V; y \in c_{ND}(V \setminus \mathcal{P}_i(k), \succsim_{d, t, i})\}) \quad (24)$$

For simplicity, the set of agreements for such a negotiation is denoted $\sigma(\Omega)$. If the former rendezvous belongs to $\sigma(\Omega)$, then the rendezvous stays the same. Otherwise, the agents randomly select a rendezvous in $\sigma(\Omega)$.

Example 5 *Let us consider the negotiation which occurs during the first simulation step of Example 4. We represent the negotiation traces in Tab. 6. Alice starts with the proposal $\{c_1\}$ while Bob starts with $\{c_9\}$. The first rendezvous is either c_3 or c_5 .*

All the agreements reached by a negotiation game are local-optimal.

Theorem 3 (Local-optimal agreements) *Let $G = (V, E)$ be a maze and Ω a set of cooperative agents. At each simulation step t ,*

$$\sigma(\Omega) = c_{ND}(V, \succsim^{lex,d,t}) \quad (25)$$

Proof 7 *Similar to the proof of Theorem 1.*

Since the negotiation phases and the moving phases are interleaved, the agents can adopt different coordination tactics. We consider here that the agents negotiate at the beginning of the simulation and whenever the path length towards the rendezvous increases for one agent due to the discovery of a wall.

Theorem 4 (Simulation success) *A simulation of cooperative agents using our negotiation game to find a rendezvous in a maze always ends.*

Proof 8 (Theorem 4) *Let us consider a simulation of a set of agents Ω who negotiate a meeting point in a maze $G = (V, E)$. Let $c \in V$ and $c' \in V$ be the negotiation outcomes at simulation steps t and $t' > t$. Each agent moves towards c from t to t' . Then only two situations occur:*

- *either none of the agents has discovered a wall on the shortest path towards c . Since c was a leximax-optimum at step t (cf Theo. 3) and each agent is closer to c , then c is still a leximax-optimum at t' . Since all the agents moved towards c , $\vec{d}_\Omega(c, t') < \vec{d}_\Omega(c, t)$.*
- *or at least one agent has discovered a new wall on the shortest path towards c then $\cup_{i \in \Omega} E_i(t') = \cup_{i \in \Omega} (E_i(t) \setminus \{(x_i(t'), y) \mid \text{wall discovered between } t \text{ and } t'\})$. It follows that $\cup_{i \in \Omega} E_i(t') \subsetneq \cup_{i \in \Omega} E_i(t)$, and then $|\cup_{i \in \Omega} E_i(t')| < |\cup_{i \in \Omega} E_i(t)|$*

Let us note that since $\forall i, E \subset E_i(t)$, $|\cup_{i \in \Omega} E_i(t)| \geq |E| \times |\Omega|$. Therefore $\vec{d}_\Omega(c, t)$ and $|\cup_{i \in \Omega} E_i(t)|$ are both lower-bounded and strictly decreasing, then the simulation ends.

5.3 Empirical results

Our experimentations aim to: (i) compare the distributed problem solving by negotiation with other heuristics; and ii) study the impact of information-sharing. We first introduce the experimental technicalities and then we present the experimental results with their analysis.

Our simulator **NegoMaze**, implemented in Java, contains:

- a model of environment to represent and generate different mazes;
- a model of agents who perceive the environment, compute the shortest path towards any cell and move;
- the implementation of our negotiation game used in Section 4.3.

The demonstrator is a GUI based on the Swing API. It allows to display the mazes (cf Fig. 6), the negotiation process and to setup the simulation (the number of agents and their coordination tactics, their perception halo and the algorithm to generate the maze). Moreover, we have developed a “test tube” to run experimentations.

The pseudo-random generation of $n \times n$ mazes is performed by:

- a pseudo-randomized version of Prim’s algorithm. This algorithm results in simple mazes (i.e. with many short dead ends and so the solution is usually pretty direct);
- a recursive backtracker algorithm. This algorithm results in hard mazes (i.e. with fewer but longer dead ends and so the solution is usually very long and twisty).

In our first experiments, we consider $n \in [2; 100]$ by step of 5. For each n , we generate 100 experiments on simple mazes and 25 experiments for hard mazes. For each experiment, we pseudo-randomly generate a new $n \times n$ maze and 10 agents’ position. The corresponding problem is solved by: i) fixing the optimal rendezvous (denoted **Opt**); ii) negotiating only if required (denoted **Neg**); iii) fixing the rendezvous at the center of the agents’ starting position (denoted **Cen0**); iv) fixing at each step the rendezvous at the centroid of all the agents’ current position in the maze (denoted **Cent**). While the optimal resolution (i) requires the knowledge of the whole maze, the other ones are performed according to the agents’ perception. Contrary to **Cen0**, **Neg** and **Cent** are adaptive heuristics: the rendezvous is recomputed during the simulation.

In Fig. 7, we consider the number of total steps performed by the 10 agents and the maximum number of steps performed by one agent (i.e. the number of simulation steps). Negotiation is not a better heuristic than fixing the rendezvous at the centroid to reduce the whole effort (top). In other words, the global satisfaction of the agents is not improved by

negotiation. However, in simple mazes, the adaptive heuristics, in particular the negotiation, require less simulation steps (at bottom left). The benefit increases with the size of the maze. The added value of negotiation, i.e. a fast resolution, is lost in hard mazes (at bottom right).

In Fig. 8, we consider the ratio between the maximum number of steps performed by one agent (i.e. the number of simulation steps) and the number of total steps divided by the number of agents (i.e. the average number of steps for one agent). This *fairness ratio* measures if the meeting point is reached at the expense of the worst-off agent. For example, the fairness ratio is around 1.15 for **Neg** in 100×100 hard mazes. Similar results are observed for smaller and/or simpler mazes. The heuristic **Cen0** minimizes the number of total steps at the expense of one agent since the fairness ratio is around 2. It is worth noticing that even if the fairness ratio is smaller for **Cen** (around 1.4), it is still 20 % higher than with negotiation. Therefore, negotiation is fairer since it shares the efforts performed by the agents better than the other heuristics.

In Fig. 9, we consider the number of simulation steps when the agents share their perception during the exploration. It allows us to shift from an unpredictable domain to a predictable one where the preferences of the other agents are taken into account. It can be noticed that an agent informs his interlocutors about the existence of a wall if the latter was not already shared. If we compare these results with Fig. 7, it is clear that the information-sharing phase improves the speed of the resolution. In hard mazes, when agents do not share information, the adaptive heuristic (e.g. negotiating) have worse performance than **Cen0**. Actually, errors (i.e. dead ends) cost a lot. Due to the recomputation of the rendezvous, the errors are cumulative. Moreover, the agents, who do not share information, can repeat the errors made by the others. When the problem is hard, an adaptive heuristic is not efficient without additional information. However, the adaptive heuristics requires less simulation steps than **Cen0** when the agents share their perception (cf top of Fig. 9). The harder the problem is, the more information-sharing is important. The number of simulation steps becomes stable when more than four agents sharing information try to meet in simple mazes (bottom left of Fig. 9). Moreover, the number of simulation steps decreases in hard mazes (bottom right of Fig. 9). The more agents we have, the more the environment is shared. Communications allow to avoid to repeat errors.

In Fig. 10, we evaluate the *information profit ratio*, i.e the number of simulation steps with pure negotiation divided by the number of simulation steps with information-based negotiation. The added value of information-sharing is higher in difficult mazes: the resolution is approximately 1.5 times faster for simple mazes and 3.5 times faster for hard mazes. It is not surprising that information-sharing help to solve hard problems. The harder the problem is, the more helpful information is.

In Fig. 11, we consider the communication cost of negotiation. For this purpose, we generate 20×20 simple and hard mazes. We distinguish the cases in which agents share (or not) their perceptions. We partially destroy walls until the density (current number of walls / initial number of walls) reaches d %. We generate 100 experiments for each d ($d \in [0; 100]$). During each of experiments, four agents at each corner negotiate in order to meet. We measure the number of proposals during the various negotiations and the number of communicated walls when the perceptions are shared. We observe that, even if the communication of walls requires more messages, the number of communicated walls is insignificant with respect to the number of proposals. More surprising, the communication cost of the pure negotiation heuristic is greater than the information-based negotiation one since the latter is more efficient and then requires less simulation steps.

Finally, we compare our protocol with the monotonic concession protocol [21] which is single-offer. Using this protocol is similar than if all the agents adopt the temporizing strategy (cf Section 4.1). Even if the cooperative strategy is required to prove the fairness of the agreements, the outcomes of these two strategies are very similar. Indeed, this is due that, in our scenario, the preference relation of an agent has a low equivalence ratio and a zero incomparability ratio (the preference relation is complete).

6 Related Work

In this section, we highlight the differences and the similarities of our work with respect the existing ones in the literature based on the ingredients of the negotiation environment (protocol, strategy and preferences) and the quality of the agreements.

Many negotiation environments have been considered in the literature:

- the negotiation object can be single-issue or multi-issue, with discrete or continuous domains;
- the preferences of the two or more agents can be captured with utilities or qualitative preference models;
- the protocol can be simultaneous (players act at the same time) or extensive (players make alternating moves), in discrete or continuous time, with a deadline;
- the information of the participants about their opponents can be perfect or imperfect.

[11] distinguish two main approaches for designing a multi-agent negotiation. In the game-theoretical approach (e.g. [24]), the negotiation environment (cf Tab. 2) is limited in order to

prove the properties of the outcome (e.g. the efficiency) and the process (e.g. the stability). In the heuristic approach, the negotiation environment (cf Tab. 3) is more realistic and the properties are empirically evaluated. Inspired by [14], we embrace in this paper both of these approaches since our results are theoretical and empirical. Since [25], the argumentative approach, in which offers are supported/defeated by arguments, has been very fruitful (see [19] for a survey). Even if we have consider this approach in our previous work [15], we restrict ourself here to a very simple alternating-offer protocol. Moreover, since we consider that the compliance of the agents to the agreements can be taken for granted, the monitoring and enforcement of the deal, performed by a normative system in [26], is also out of the scope of this paper.

Protocol. Most of the existing work are based on the following two protocols:

1. the alternating offers protocol [22] which is asymmetric and extensive. At each round, one agent can make a proposal to his opponent, who can accept, make a counter-offer or exit the negotiation. If a counter-offer is made, the negotiation goes on until one of the agents accepts an offer or exits the negotiation. Several work are built on a continuous version of this protocol in which each agent can make an offer at any time;
2. the monotonic concession protocol [21] which is symmetric and simultaneous. At each round, both agents make an offer. An agreement is reached when one agent makes a proposal that his opponent rates at least as high as his own current proposal. Otherwise, a disagreement occurs, and each agent can stand still or concede. A concession consists of a counter-offer which is better for his opponent. If no agent concedes, the negotiation ends unsuccessfully.

Contrary to the alternating offers protocol, the latter requires that each agent knows the preferences of his opponent to determine the legal offers. We propose here a symmetric and simultaneous protocol which does not require perfect information and which allows an agent to submit multiple offers at each round.

Preferences. Most of the existing work [22, 21, 10, 23, 4] are built on a preference model which is cardinal and quantitative, i.e. a utility function (denoted u) or a linear ordering (e.g. [7]), i.e. a complete, transitive and antisymmetric relation. To the best of our knowledge, [2] is the only work considering a preference model which is ordinal, qualitative and partial (denoted \succsim). Even if a utility function can be modeled by a complete preference relation ($a \succsim b$ iff $u(a) \geq u(b)$), the utility function supposes that the intensity of preferences can be measured, or even compared. This preference model suggests that the satisfaction of one agent can be greater than another. It allows monetary compensation. We do not make such assumptions. While [2] focus on CP-net, we consider incomplete (partial and weak)

preference relations in general. Moreover, our negotiation framework is setup to prove the optimality of the agreement (i.e. the fairness) of multi-party negotiations.

Strategy. Based on the monotonic concession protocol and utility functions, the Zeuthen strategy [21] is compliant with the following principle: an agent who has already made many concessions has less to lose if conflicts happen. For this purpose, each agent computes at each round the minimal sufficient concession that gives his opponent the least utility. It is worth noticing that this strategy requires perfect information about the opponent's utility. This is the case for most of the concessions defined in [8]. In the game-theoretical approach, mixed strategies (i.e. assigning a probability to each outcome) allow us to handle imperfect information. Pure strategies which are deterministic can only handle perfect information. In most of the cases, the negotiation deadline allows backward induction to determine the optimal strategy. Our negotiation environment has no deadline and it deals with imperfect information.

[10] adopt the heuristic approach for designing strategies in multi-issue negotiations. In this work, agents are involved in multiple bilateral negotiations based on a continuous version of the alternating offers protocol with a deadline. The preferences of agents are captured by multi-attribute utility functions. While the offers are evaluated by such functions, the offers are submitted in accordance with a linear combination of simple functions called tactics. A tactic is a function which computes the value of an issue (price, quantity, quality, etc.) from a criterion (time, available resources, opponent behaviour). For instance, [10] introduce the bouldware tactic which consists in standing still as much as possible and the conceder tactic which consists in quickly conceding up to the reservation value. Inspired by this work, the annual automated negotiating agents competition lays emphasis on the development of successful automated negotiators in realistic bilateral multi-issue environments [4]. As shown in Tab. 3, our negotiation environment is different: we address single-issue and multi-party negotiation.

Efficiency. Most of the existing negotiation processes (strategies and protocols) aim at reaching agreements which are Pareto-optimal. Additionally, a wide range of concepts taken from Social Choice Theory can be used to assess the quality of agreements. Examples include utilitarian social welfare [16], where the aim is to maximize the sum of individual utilities, and egalitarian social welfare [9], where the aim is to maximize the individual welfare of the agent that is currently worst off. However, these works are also based on a preference model which is cardinal and quantitative. Based on a preference model which is ordinal, qualitative and partial we refine the egalitarian rule by generalizing the fairness principle in case of ties to the second worst-off agent and so on.

7 Conclusions

In this paper, we have proposed a multi-party negotiation game where agents exchange offers over a single-issue. This is an operational and distributed process, since the computation of the acceptability criteria does not require the knowledge of the preferences. Contrary to most existing works [22, 21, 10, 23, 4], we capture the agents' rationality with incomplete preferences. In this way, agents can exhibit indecisiveness since they are not always able to compare alternatives due to lack of information, uncertainty, or conflicting preferences. For this purpose, we have introduced the notion of rank for an alternative. It is worth noticing that our metric does not conform to the axiom of independence of irrelevant alternatives (eliminating some of the unchosen alternatives should affect the selection of the best option).

Since we aim to provide a fair outcome, i.e. one which does not give priority to one agent and minimizes the maximum efforts of the agents, our protocol is simultaneous and allows agents to offer more than two alternatives per round. We have adopted concession-based strategies where agents first submit their preferred alternatives and, if that fails, they concede, i.e. they withdraw their expectation.

We have proposed and evaluated two concrete strategies: a cooperative one to reach a common goal; and a temporizing one, more suitable for selfish agents in competition. Even if our negotiation game always ends successfully, all the agents need to be cooperative to reach fair agreements. In practice, the individual advantage of the temporizing strategy requires very restricted conditions. Reciprocally, the fairness of agreements reached by temporizing negotiators highly depends on the structure of the preference relations as in our usecase. Moreover, the outcome of a negotiation is less arbitrary than the outcome of single-agent decision-making, since the uncertainty of one agent can be balanced by the preferences of his interlocutors.

Through a toy problem, we have shown that our negotiation game can tackle distributed problem solving. In particular, we have considered the case of two or more agents in a maze who aim to meet as soon as possible by negotiating in order to minimize and share the efforts. Our experiments show that negotiation does not reduce the whole effort but it allows the agents to fairly share them. Therefore, multi-agent negotiation is an adaptive heuristic which allows us to solve faster than the naive heuristic (i.e. fixing the rendezvous at the initial centroid) simple problems. Negotiation with information-sharing can solve hard problems faster. The harder the problem, the more helpful the information. It may seem counter-intuitive but we have shown that the communication cost of the pure repeated negotiations may be greater than information-based ones since the latter can be more efficient. Even if we consider here only mazes, our framework allows us to explore any graph, whatever

the typology (scale-free, small world). Moreover, our scenario consists of a “test tube” to evaluate and compare negotiation protocols/strategies such as argumentative ones. Finally, our experimentations assess the scalability of our negotiation framework which can tackle large negotiation domains with hundreds of participants.

In future works, we plan to apply our negotiation game for group decision support systems in various domains which require the elicitation of incomplete user preferences. For instance, the design of a meeting scheduling system must be built on a language for expressing preferences about time slots and take into account the incomplete preferences about them, as can be done with our negotiation game.

Acknowledgements

The authors would like to thank the reviewers for their comments on this paper.

References

- [1] Steve Alpern and Shmuel Gal. *The theory of search games and rendezvous*. Springer, 2003.
- [2] Reyhan Aydogan, Tim Baarslag, Koen V. Hindriks, Catholijn M. Jonker, and Pinar Yolum. Heuristic-based approaches for CP-nets in negotiation. In *Complex Automated Negotiations*, pages 113–123. Springer, 2013.
- [3] Tim Baarslag, Alexander Dirkzwager, Koen V. Hindriks, and Catholijn M. Jonker. The significance of bidding, accepting and opponent modeling in automated negotiation. In *Proc. of 21st European Conference on Artificial Intelligence*, pages 27–32, Prague, Czech Republic, 2014. IOS Press.
- [4] Tim Baarslag, Katsuhide Fujita, Enrico H. Gerding, Koen Hindriks, Takayuki Ito, Nicholas R. Jennings, Catholijn Jonker, Sarit Kraus, Raz Lin, Valentin Robu, and Colin R. Williams. Evaluating practical negotiating agents: Results and analysis of the 2011 international competition. *Artificial Intelligence*, 198:73–103, 2013.
- [5] S. Barberà, W. Bossert, and P. K. Pattanaik. *Handbook of utility theory*, chapter Ranking sets of objects. Springer, 2004.
- [6] Sylvain Bouveret, Ulle Endriss, and Jérôme Lang. Fair division under ordinal preferences: Computing envy-free allocations of indivisible goods. In *Proc. of the 17th Eu-*

- ropean Conference on Artificial Intelligence (ECAI), pages 387–392, Lisbon, Portugal, 2010.
- [7] Clemens Puppe Dorothea Herreiner. A simple procedure for finding equitable allocations of indivisible goods. *Social Choice and Welfare*, 19(2):415–430, April 2002.
- [8] Ulle Endriss. Monotonic concession protocols for multilateral negotiation. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 392–399, Hakodate, Japan, 2006. ACM.
- [9] Sylvia Estivie, Yann Chevaleyre, Ulle Endriss, and Nicolas Maudet. How equitable is rational negotiation? In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 866–873, Hakodate, Japan, 2006. ACM.
- [10] Peynam Faratin, Carles Sierra, and Nick .R. Jennings. Negotiation decision functions for autonomous agents. *Robotics and Autonomous Systems*, 24(3-4):159–182, 1998.
- [11] Shaheen Fatima and Iyad Rahwan. *Multiagent Systems*, chapter Negotiation and Bargaining, pages 143–176. MIT Press, 2013. second edition.
- [12] Leo A Goodman and Harry Markowitz. Social welfare functions based on individual rankings. *American Journal of Sociology*, 58(3):257–262, Nov. 1952.
- [13] N.R. Jennings, P. Faratin, A.R. Lomuscio, S. Parsons, M.J. Wooldridge, and C. Sierra. Automated negotiation: Prospects methods and challenges. *Group Decision and Negotiation (GDN)*, 10(2):199–215, 2001.
- [14] Sarit Kraus. *Strategic Negotiation in Multiagent Environments*. MIT Press, 2001.
- [15] Maxime Morge and Paolo Mancarella. Arguing over goals for negotiation: Adopting an assumption-based argumentation decision support system. *Group Decision and Negotiation*, 23(5):979–1012, 2014.
- [16] Antoine Nongailard, Philippe Mathieu, and Brigitte Jaumard. A multi-agent resource negotiation for the utilitarian welfare. In *Engineering Societies in the Agents World IX, 9th International Workshop (ESAW), Revised Selected Papers*, pages 208–226, Saint-Etienne, France, 2008.
- [17] Meltem Öztürk, Alexis Tsoukiàs, and Philippe Vincke. *Multiple Criteria Decision Analysis : State of the Art*, chapter Preference Modelling. Springer, 2012.

- [18] Dean G. Pruitt. *Negotiation Behavior*. Academic Press, 1981.
- [19] Iyad Rahwan, Sarvapali D. Ramchurn, Nicholas R. Jennings, Peter Mcburney, Simon Parsons, and Liz Sonenberg. Argumentation-based negotiation. *The Knowledge Engineering Review (KER)*, 18(4):343–375, 2004.
- [20] John Rawls. *A Theory of Justice*. Harvard University Press, 1971.
- [21] J.S. Rosenschein and G. Zlotkin. *Rules of Encounter - Designing Conventions for Automated Negotiation among Computers*. MIT Press, 1994.
- [22] Ariel Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–102, January 1982.
- [23] Tuomas Sandholm and Nir Vulkan. Bargaining with deadlines. In *Proc. of Sixteenth National Conference on Artificial Intelligence (AAAI)*, pages 44–51, Orlando, Florida, 1999.
- [24] Yoav Shoham and Kevin Leyton-Brown. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, 2008.
- [25] Katia Sycra. Persuasive argumentation in negotiation. *Theory and Decision*, 28(3):203–242, 1990.
- [26] A Game Theoretic Approach to Contracts in Multiagent Systems. Guido boella and leendert van der torre. *IEEE Transactions on Systems, Man, and Cybernetics - Part C: Applications and Reviews*, 6(1):68–79, January 2006.

A Authors short biographical notes

Fabien Delecroix has a PhD in Computer Science. He is specialized in Artificial Intelligence and Multiagent systems, in particular in Social Choice Theory. He was Lecturer in Computer Science at the University of Lille 1 between 2011 and 2014.

Maxime Morge is Assistant Professor of Computer Science at the University of Lille 1 since 2009. He is specialized in Artificial Intelligence and Multiagent systems, in particular in automated negotiation. M. Morge has been research associate of the 035200 ARGUGRID project (part of the Sixth Framework IST programme of the EC) at Università di Pisa from 2006 to 2009. He has been a post-doc at Université Lille 1 in 2005 during one year. M. Morge has defended his PhD at Ecole Nationale Supérieure des Mines de Saint Etienne in 2005.

Thomas Nachtergaelle received his MS in Computer Science from the University of Lille 1 in 2015. He is Research Engineer at Worldline Global.

Jean-Christophe Routier is full Professor of Computer Science at the University of Lille 1 since 2007. He is specialized in Artificial Intelligence and in particular in Multiagent systems. He has worked on the modelling of behaviors in situated multi-agent simulations, and on multi-agent distributed problem solving. Jean-Christophe Routier was Assistant Professor of Computer Science at the University of Lille 1 from 1994 to 2007. Jean-Christophe Routier has defended his PhD at the University of Lille 1 in 1994.

Table 1: Table of symbols

Symbol	Description
\succsim	Preference relation (cf Def. 1)
$c_{ND}(\mathcal{X}, \succsim)$	Non-dominated alternatives (cf Def. 2)
\succsim_{NP}	Necessary Pareto preference (cf Def. 4)
\succsim_{PP}	Possible Pareto preference (cf Def. 4)
$r(x, \mathcal{X}, \succsim)$	The rank of the alternative (cf Def. 5)
$c_{ND}(\mathcal{X}, \succsim^{lex,r})$	Fair alternatives (cf Def. 7)
$c_{ND}(\mathcal{X}, \succsim_{PP})$	Possible Pareto optimum (cf Prop. 1)
$k_{\Omega}(a)$	The vector of rounds (cf Def. 8)
$\mathcal{P}_i(k)$	The set of alternatives proposed by agent i before the round k (cf Def. 8)
C_{Ω}	The strategy profile of cooperative agents (cf Def. 9)
T_{Ω}	The strategy profile of temporizing agents (cf Def. 9)
$C_{\Omega-i}$	The strategy profile of cooperative agents excepted agent i (cf Def. 9)
$T_{\Omega-i}$	The strategy profile of temporizing agents excepted agent i (cf Def. 9)
$\sigma(S_{\Omega})$	The set of agreements (cf Def. 9)
$minEqRa$	The minimal ratio of equivalent couples (cf Sec. 4.3)
$maxIncRa$	The maximal ratio of incomparable couples (cf Sec. 4.3)
$d_i(y, t, E)$	The real length of the shortest path for the agent i to reach y (cf Sec. 5.1)
$\hat{d}_i(y, t, E_i)$	The estimated length of the shortest path for the agent i to reach y (cf Sec. 5.1)

Table 2: Analysis grid of the game-theoretical approach

	[21]	[22]	[23]
Nb of Agents	2	2	2
Object	single-issue	single-issue	single-issue
Preference	u	u	u
Protocol	symetric	asymmetric	asymmetric continuous
deadline	no	no	yes
Information	perfect	imperfect	imperfect
Efficiency	Pareto - Nash	not Pareto	not Pareto
Equilibrium	no	sequential	sequential

Table 3: Analysis grid of the heuristic approach

	[10]	[4]	[2]	Here
Nb of Agents	$n \geq 2$	2	2	$n \geq 2$
Object	multi-issue	multi-issue	multi-issue	single-issue
Preference	u	u	\succsim	\succsim
Protocol deadline	asymetric continuous yes	asymetric continuous yes	asymetric continuous yes	symetric no
Information	imperfect	imperfect	imperfect	imperfect
Efficiency	?	?	?	Pareto - Leximax
Equilibrium	?	?	?	no




Figure 1: The preference relations \succsim_1 (at left), \succsim_2 (at center) and \succsim_3 (at right).

Table 4: The vector of ranks for the preferences in Fig. 1

x	$r(x, \mathcal{X}, \succsim_1)$	$r(x, \mathcal{X}, \succsim_2)$	$r(x, \mathcal{X}, \succsim_3)$	$\vec{r}_\Omega(x)$
a	1	4	1	(4,1,1)
b	1	4	1	(4,1,1)
c	1	1	4	(4,1,1)
d	2	2	3	(3,2,2)
e	2	3	2	(3,2,2)
f	3	3	2	(3,3,2)

Table 5: Trace of negotiation in the game configuration C_Ω (at left) and in the game configuration $C_{\Omega-3}$ (at right)

k	C_1	C_2	C_3
1	$\{a, b, c\}$	$\{c\}$	$\{a, b\}$
2	$\{d, e\}$	$\{d\}$	$\{e, f\}$
3	$\{f\}$	$\{e, f\}$	$\{d\}$
$\sigma(C_\Omega) = \{d, e\}$			

k	C_1	C_2	T_3
1	$\{a, b, c\}$	$\{c\}$	$\{a\}$ or $\{b\}$
2	$\{d, e\}$	$\{d\}$	$\{b\}$ or $\{a\}$
3	$\{f\}$	$\{e, f\}$	$\{e\}$ or $\{f\}$
$\sigma(C_{\Omega-3}) = \{e\}$ or $\sigma(C_{\Omega-3}) = \{f\}$			

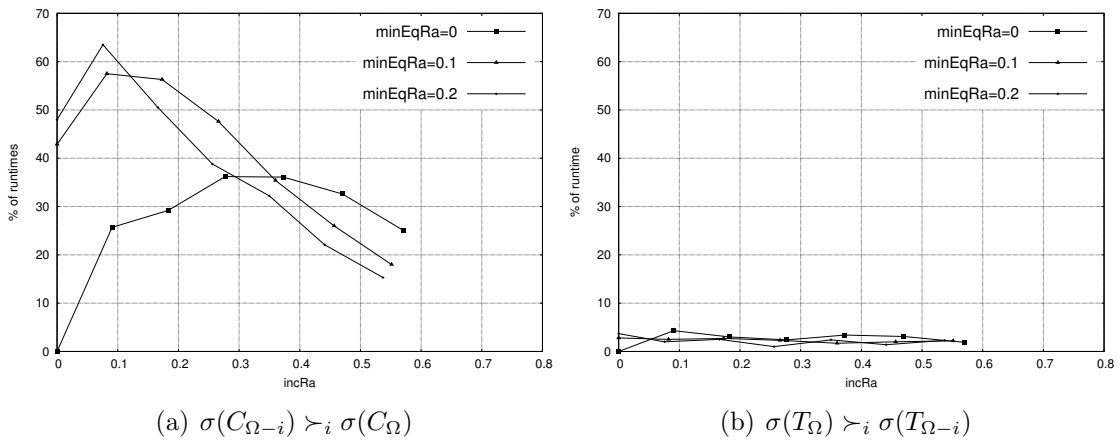


Figure 2: Percentage of runtimes where the temporizing strategy dominates the cooperative one.

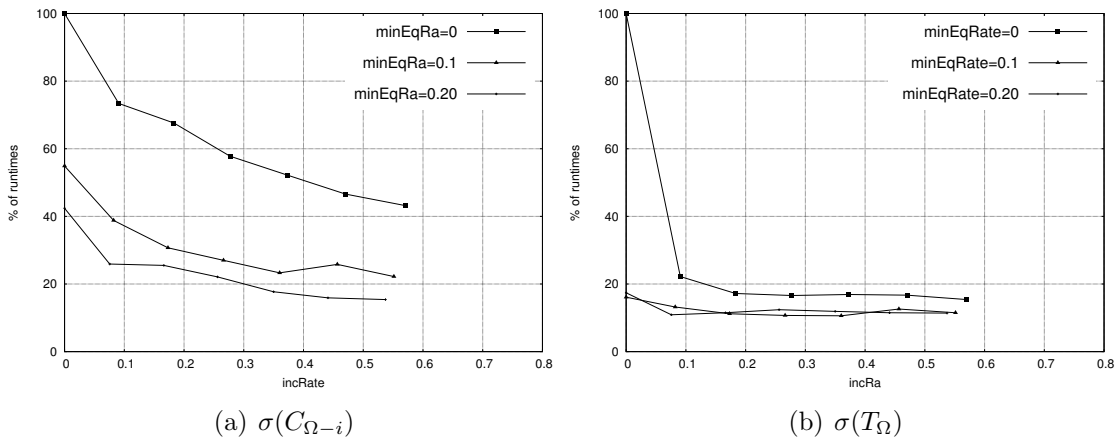


Figure 3: Percentage of runtimes where the agreements are fair.




Figure 4: Cardinality of the non-dominated alternatives for the agent 1 and cardinality of the set of agreements when the all the agents cooperate ($\sigma(C_\Omega)$).

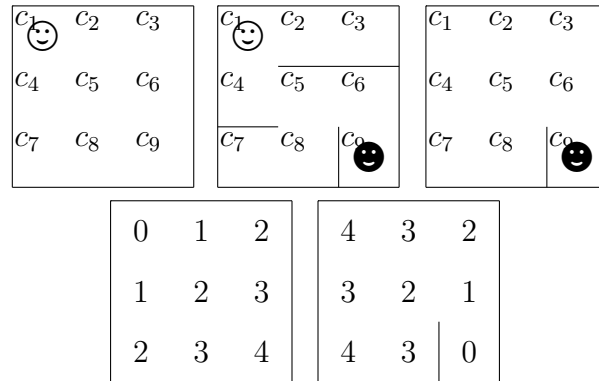


Figure 5: First step of simulation, \ominus represents Alice and $\omin�$ represents Bob: G (at top center) and its internal representation for the agent $\mathcal{G}_{\ominus}(1)$ (at top left) and $\mathcal{G}_{\omin�}(1)$ (at top right) and the estimation by the agents of the distances $d_{\ominus}(y, 1, E_{\ominus})$ (at bottom left) and $d_{\omin�}(y, 1, E_{\omin�})$ (at bottom right).

Table 6: Trace of negotiation between Alice (\odot) and Bob (\ominus) during the first simulation step

k	$u_{\odot,k}$	$u_{\ominus,k}$
1	$\{c_1\}$	$\{c_9\}$
2	$\{c_2, c_4\}$	$\{c_6\}$
3	$\{c_3, c_5, c_7\}$	$\{c_3, c_5\}$
	$\sigma(\Omega) = \{c_3, c_5\}$	

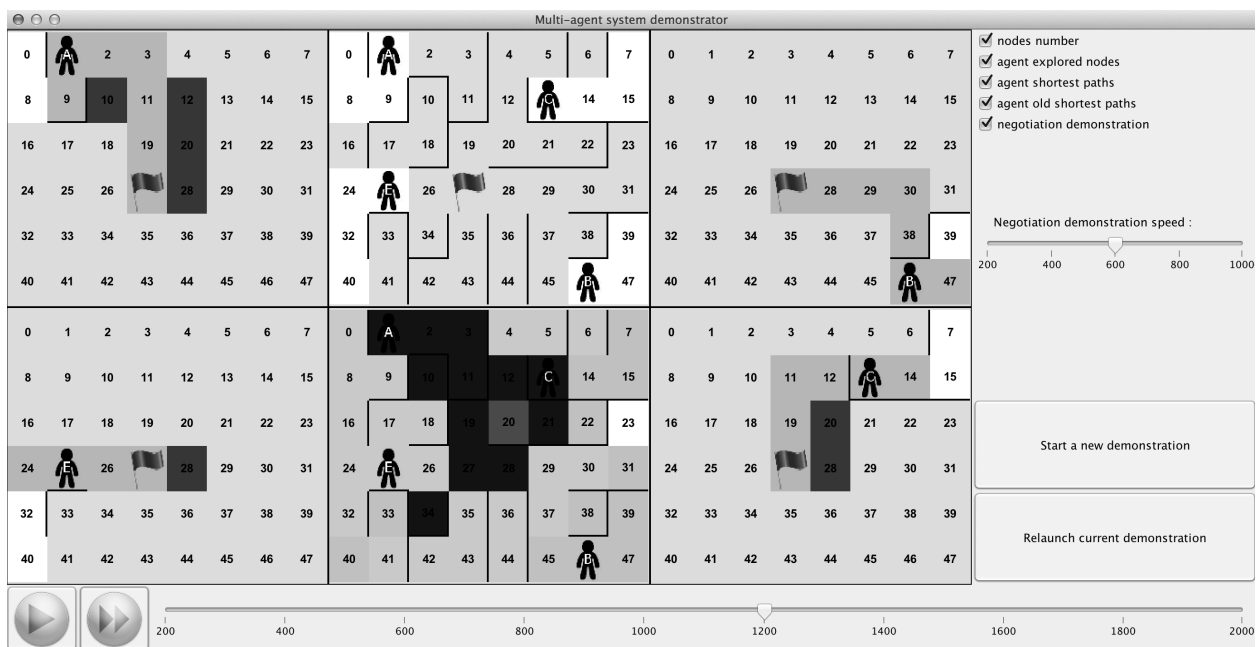


Figure 6: Screenshot of the demonstrator. The panel at center allows to observe the agents' moves, the one below represents the offers and the others represent the local perceptions of the environment by the agents.

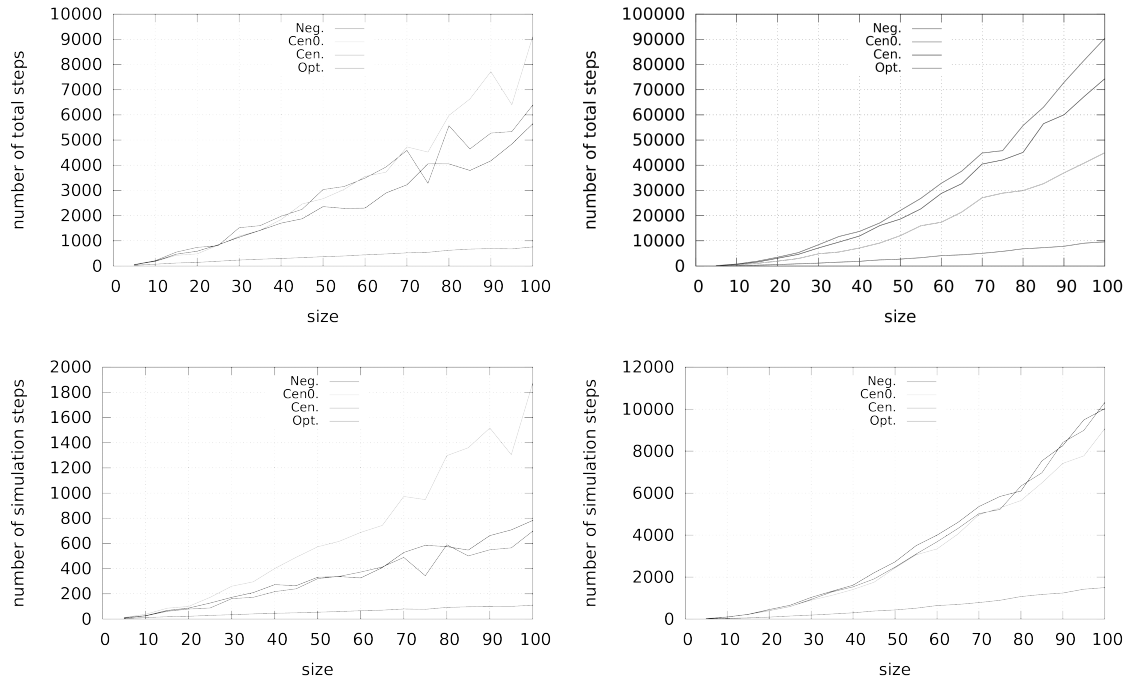


Figure 7: The number of total steps performed by the 10 agents (at top) and the number of simulation steps (at bottom) in some mazes. The simple mazes are generated with the pseudo-randomized version of Prim's algorithm (at left). The hard mazes are generated with the recursive backtracker algorithm (at right). The agents do not share their perception.




Figure 8: The fairness ratio (maximal number of steps / average number of steps) in some mazes. The simple mazes are generated with the pseudo-randomized version of Prim's algorithm (at left). The hard mazes are generated with the recursive backtracker algorithm (at right). The agents do not share their perception.

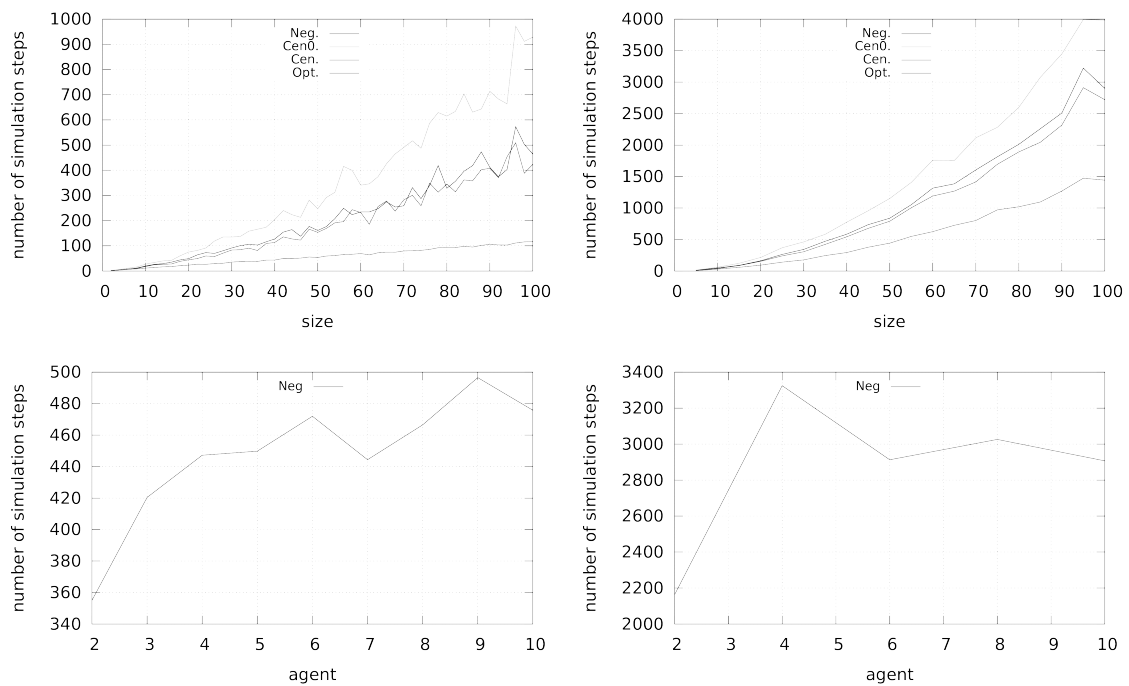


Figure 9: The number of simulation steps wrt the size of the maze (at top) and wrt the number of agents in 100×100 mazes (at bottom) in some mazes. The mazes are simple (at left) and hard (at right). The agents share their perception.

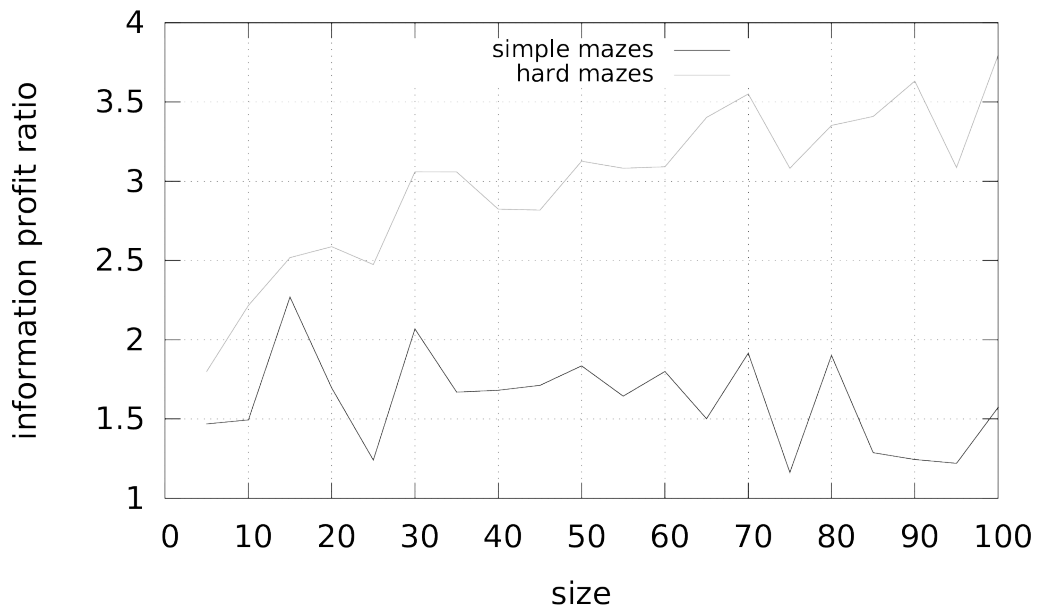


Figure 10: Information profit ratio for simple/hard mazes.




Figure 11: Number of proposals and communicated walls in simple mazes (at left) and in hard mazes (at right).

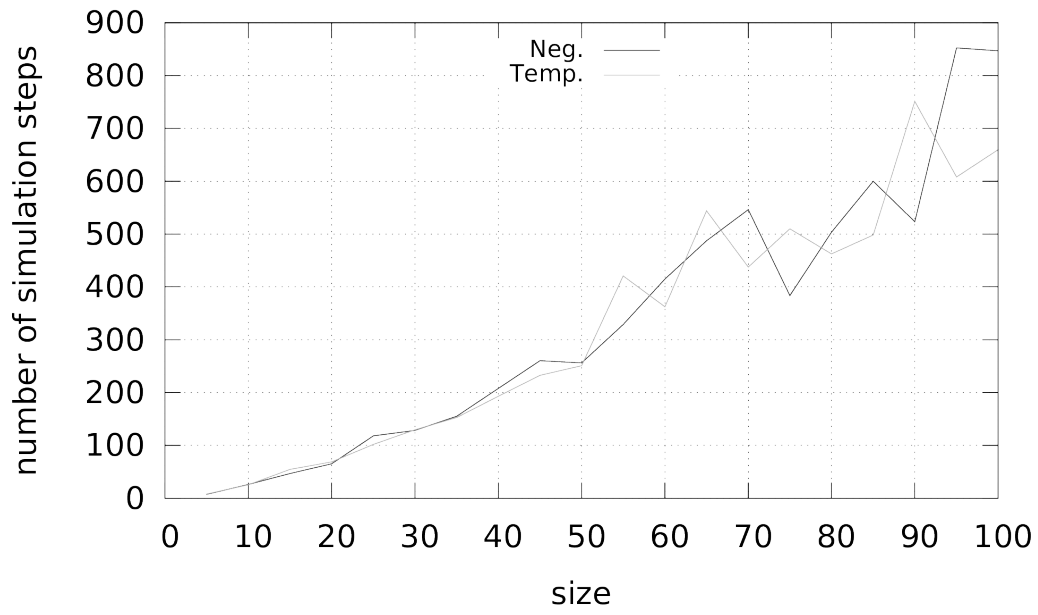


Figure 12: Number of simulation steps with the cooperative strategy (Neg) and the temporizing strategy (Temp).