

Strategies for Distributing Goals in a Team of Cooperative Agents

Laurence Cholvy¹ and Christophe Garion²

¹ ONERA Toulouse, 2 avenue Édouard Belin, 31055 Toulouse, France
cholvy@cert.fr

² SUPAERO, 10 avenue Édouard Belin, 31055 Toulouse, France
garion@supaero.fr

Abstract. This paper addresses the problem of distributing goals to individual agents inside a team of cooperative agents.

It shows that several parameters determine the goals of particular agents. The first parameter is the set of goals allocated to the team; the second parameter is the description of the real actual world; the third parameter is the description of the agents' ability and commitments. The last parameter is the strategy the team agrees on: for each precise goal, the team may define several strategies which are orders between agents representing, for instance, their relative competence or their relative cost. This paper also shows how to combine strategies. The method used here assumes an order of priority between strategies.

1 Introduction

Reaching a complex goal often needs to consider a group of agents which must cooperate in order to achieve this goal [1]. For instance, nations often group into coalitions in order to maintain peace in a conflicting area, that means sharing information about the situation, providing emergency medical treatment, providing displaced civilian services, providing engineering infrastructure support etc [2].

The goal allocated to the group is some proposition that one desires the group to make true, or equivalently, goals define some desirable worlds the group must reach. But, as Boutilier noted it in [3], goals are not always achievable. It may happen for instance that none of the agents in the team has the ability to make this proposition true. Furthermore, goals may be defeated for reasons other than inability. It is often natural to specify general goals, but list exceptional circumstances that make the goal less desirable than the alternatives. Rather than a categorical distinction between desirable and undesirable situations, it is more general to rank worlds according to their degree of preference. The most preferred worlds correspond to goal states in the classical sense. However, when such states are unreachable, a ranking on alternatives becomes necessary.

In a previous paper [4] we have considered this general case and we have defined a goal distribution process which allocates goals to individual agents

of a group, according to the preferences representing the goals allocated to the group, the actual world and the agents' ability and commitments.

The agents we consider are cooperative in the sense that they do not contradict each other in their commitments (for instance, we discard the case when one agent commits himself to make a proposition true whilst another one commits himself to make this proposition false) and their commitments do not contradict a goal of the group. Notice that this process of goals distribution is not based on a negotiation between the agents like in [5, 6, 7]. More precisely, it may be viewed as managed by a central authority which knows how is the real world, what are the agents' abilities and the agents' commitments and which allocates to each agent some goals which correspond to the most preferred situations the group can thus reach.

In this present paper, we refine this work and we show that a fourth parameter can be used for determining the goals of particular agents. This last parameter is the strategies the team agrees on. We will see that a strategy depends on a particular goal and is an order between the agents that are able to achieve it. For instance, given a particular goal, agents may be ordered according to their relative competence for achieving it. Agents may also be ordered according to their relative cost for achieving this goal.

We will also show how to combine strategies. For instance, given a particular goal, we could want to order agents by taking into account their relative competence *and* their relative cost. The method used here for combining strategies assumes a priority order between the strategies.

This paper is organized as follows. In section 2, we summarize the process described in [4]. This process is illustrated on an example in section 3. Section 4 focuses on the notion of strategies and combination of strategies. This point is illustrated in section 5. Finally, section 6 is devoted to a discussion.

2 Distribution of Goals Addressed to a Group of Agents

2.1 Preferences Representation

To represent preferences, we use a logic [3], [8] whose language L_B is based on a set of atomic propositional variables $PROP$ with the usual connectives and two modal operators \Box , $\bar{\Box}$.

Models are of the form $\mathcal{M} = \langle W, \leq, val \rangle$. W is a set of possible worlds, \leq is a total *preference* preorder on W (a reflexive and transitive relation on $W \times W$). If w and w' are two worlds of W , then $w \leq w'$ means that w is at least as preferred as w' . Finally, val is a valuation function on W ¹. For any formula φ of W , $val(\varphi)$ is the set of worlds of W which classically satisfy φ .

Let $\mathcal{M} = \langle W, \leq, val \rangle$ be a model. Satisfaction of modal formulas is defined as follows:

¹ I.e. $val : PROP \rightarrow 2^W$ and val is such that $val(\neg\varphi) = W - val(\varphi)$ and $val(\varphi_1 \wedge \varphi_2) = val(\varphi_1) \cap val(\varphi_2)$.

- $\mathcal{M} \models_w \Box\varphi$ iff $\forall w' \in W \quad w' \leq w \implies \mathcal{M} \models_{w'} \varphi$.
- $\mathcal{M} \models_w \Box\varphi$ iff $\forall w' \in W \quad w' \not\leq w \implies \mathcal{M} \models_{w'} \varphi$.

$\Box\varphi$ is true at a world w if and only if φ is true at all worlds at least as preferred as w (including w). $\Box\varphi$ is true at world w if and only if φ is true at all the worlds less preferred than w .

Boutillier then defines two dual modal operators : $\Diamond\varphi \equiv_{def} \neg\Box\neg\varphi$ means that φ is true at some equally or more preferred world and $\overline{\Box}\varphi \equiv_{def} \neg\Box\neg\varphi$ means that φ is true at some less preferred world.

$\overline{\Box}\varphi \equiv_{def} \Box\varphi \wedge \Box\neg\varphi$ and $\overline{\Diamond}\varphi \equiv_{def} \Diamond\varphi \vee \overline{\Box}\varphi$ correspond respectively to classical necessity and possibility.

A formula φ is valid in \mathcal{M} (noted $\mathcal{M} \models \varphi$) iff $\forall w \in W \quad \mathcal{M} \models_w \varphi$.

Conditional preferences are formulas of the form $I(\beta|\alpha)$ which means that “ideally, if α is true, then β is true”. The connective $I(-|-)$ is defined by $I(\beta|\alpha) \equiv_{def} \overline{\Box}\neg\alpha \vee \overline{\Box}(\alpha \wedge (\Box\alpha \rightarrow \beta))$. $I(\beta|\alpha)$ is valid in \mathcal{M} iff either α is false in every world of W , or there is some world w which satisfies α and such that every world at least as preferred as w satisfies $\alpha \rightarrow \beta$.

2.2 Description of the Actual World

Following a centralized approach, we consider that the team of agents is aware of a shared description of the situation. The actual world is thus described by a finite and consistent set of formulas of $PROP$. It is denoted KB^2 . $Cl(KB)$ denotes its closure by classical logical consequence³.

2.3 Controllability

Let $\mathcal{A} = \{a_1, \dots, a_n\}$ be a finite set of agents. Like Boutillier, for each agent a_i , we partition the literals of $PROP$ into two sets: C_{a_i} (the literals that a_i can control) and \overline{C}_{a_i} (the literals uncontrollable by a_i). We assume that each agent of the group controls at least one literal:

Assumption 1. $\forall a_i \in \mathcal{A} \quad C_{a_i} \neq \emptyset$

The notion of controllability for a group of agents is then defined by:

Definition 1. Let $lit(PROP)$ be the set of literals in the propositional language. The set of controllable literals by the group of agents is $C = \bigcup_{a_i \in \mathcal{A}} C_{a_i}$ and the set of uncontrollable literals by the group of agents is $\overline{C} = lit(PROP) - C$.

Extension to general propositions is given below:

Definition 2. Let w and w' be two worlds of W . Let us note $w' - w = \{l : w' \models l, w \models \neg l \text{ and } l \text{ is a literal}\}$. A proposition φ is:

² Knowledge Base.

³ In the original work, this closure is defined as a default closure.

- *controllable* iff $\forall w \in W (w \models \neg\varphi \exists w' \in W w' \models \varphi \text{ and } (w' - w) \subseteq C)$;
- *influenceable* iff $\exists w \in W (w \models \neg\varphi \exists w' \in W w' \models \varphi \text{ and } (w' - w) \subseteq C)$. In this case, we say that φ is influenceable in w .
- *uninfluenceable* iff it is not influenceable.

2.4 Contexts

Definition 3. A world $w \in W$ is a context for some influenceable proposition φ iff φ is influenceable in w or $w \models \varphi$.

The contexts of an influenceable proposition φ are the worlds in which either φ is false but the agent can change the valuations of some controllable literal to make φ true, or the worlds in which φ is already true.

Definition 4. The set on non-contextual propositions of KB is defined by:

$$NC(KB) = \{\varphi \in Cl(KB) : Cl(KB) \text{ is not a context for } \neg\varphi\}$$

$NC(KB)$ represents the propositions whose truth value will not be changed by some agents' actions (because the group of agents has no ability to do that). We suppose here that $NC(KB)$ is complete.

2.5 CK-Goals of the Group

Definition 5. Let \mathcal{P} be a set of conditional preferences. φ is a CK goal⁴ for \mathcal{A} iff $\mathcal{P} \models I(\varphi|NC(KB))$ and KB is a context for φ .

2.6 Commitments

Given a literal controllable by an agent, this agent can express that it will do an action that will keep or make this the literal true (we say that the agent commits itself to achieve the literal), or the agent can express that it will not do an action that can make the literal true (we will say that the agent commits itself not to achieve the literal), or finally, the agent can express nothing about the literal (we will say that the agent does not commit itself neither to achieve the literal nor not to achieve the literal).

To represent the commitments of each agent a_i , we will use three subsets of C_{a_i} : Com_{+,a_i} , Com_{-,a_i} and P_{a_i} . defined as follows. If l is a literal, if l is controllable by a_i and $l \in Com_{+,a_i}$, it means that “the agent a_i commits itself to achieve l ”; if l is a literal, if l is controllable by a_i and $l \in Com_{-,a_i}$, it means that “the agent a_i commits itself not to achieve l ”. Finally, $P_{a_i} = C_{a_i} - (Com_{+,a_i} \cup Com_{-,a_i})$ is the set of controllable literals by a_i and for which a_i does not commit itself to anything (i.e. a_i does not commit itself neither to achieve them nor not to achieve them).

We impose two constraints on those sets.

⁴ For Complete-Knowledge goal as introduced by Boutilier.

Constraint 1. $\forall a_i \in \mathcal{A}$ Com_{+,a_i} is consistent.

Constraint 2. $\forall a_i \in \mathcal{A}$ $Com_{+,a_i} \cap Com_{-,a_i} = \emptyset$

Those two constraints express a kind of consistency for the agent's commitments. The first constraint expresses the fact that an agent does not commit itself to achieve both l and $\neg l$. The second constraint expresses the fact that an agent cannot commit itself both to achieve l and not to achieve l .

Definition 6. $Com_{+,\mathcal{A}}$ is the set of positive commitments of the agents:

$$Com_{+,\mathcal{A}} = \bigcup_{a_i \in \mathcal{A}} Com_{+,a_i}$$

$Com_{-,\mathcal{A}}$ is the set of "negative" commitments of the agents:

$$Com_{-,\mathcal{A}} = \{l \in KB : \forall a_i \in \mathcal{A} \neg l \text{ controllable by } a_i \Rightarrow \neg l \in Com_{-,a_i}\}.$$

The meaning of $Com_{-,\mathcal{A}}$ is the following: if all the agents that control a literal l commit themselves not to achieve l and $\neg l \in KB$, we will consider that $\neg l$ will remain true. We suppose that there is no external intervention.

An assumption that we do on the agents' commitments is: *every CK goal of \mathcal{A} is consistent with the union of $Com_{+,\mathcal{A}}$ and of $Com_{-,\mathcal{A}}$.*

Assumption 2. For every formula φ such that $\mathcal{P} \models I(\varphi|NC(KB))$ and KB is a context for φ , then $Com_{-,\mathcal{A}} \cup Com_{+,\mathcal{A}} \cup \{\varphi\}$ is consistent.

This restriction allows to eliminate some problematic cases like the case where an agent which controls l commits itself to achieve l and another one which controls $\neg l$ commits itself to achieve $\neg l$ (i.e. $Com_{+,\mathcal{A}}$ not consistent). It also eliminates the case where a literal, which is not consistent with the group's CK goals, is true in KB and will remain true because the agents of the group which could make it false do not commit themselves to it, or finally, the case where the positive and negative commitments of the group are not consistent with some CK goal of the group. If this assumption is not verified, the agents must review their commitments.

2.7 Effective Goals

If the assumption 2 is verified, then the agents' commitments are consistent with the group's CK goals. The goals of each agent do not only depend on $NC(KB)$, but also on the commitments of the other agents.

Definition 7. We define:

$$D(KB) = NC(KB) \cup Com_{+,\mathcal{A}} \cup Com_{-,\mathcal{A}}$$

$D(KB)$ contains the propositions of KB for which KB is not a context, i.e. $NC(KB)$, plus the set of positive commitments of the agents, i.e. $Com_{+,\mathcal{A}}$ and the set of "negative" commitments of the agents, i.e. $Com_{-,\mathcal{A}}$. We have proved that $D(KB)$ is consistent. This set will be used in the *conditional* part of $I(-|-)$ to deduce the effective goals of each agent as follows:

Definition 8. Let \mathcal{P} be a set of preferences addressed to the group \mathcal{A} . φ is an effective goal for a_i , denoted by $EGoal_{a_i}(\varphi)$, iff $\mathcal{P} \models I(\varphi|D(KB))$ and KB is a context for φ for a_i .

As we use the $I(-|-)$ operator, we are sure that an agent cannot have contradictory goals.

Effective atomic goals are defined by :

Definition 9. Let \mathcal{P} be a set of conditional preferences. A set of atomic goals is a set of controllable literals $\mathcal{L} = \{l_1, \dots, l_n\}$ such that:

- $\forall i \in \{1, \dots, n\}$ $Cl(KB)$ is a context for l_i .
- for all CK goal φ given \mathcal{P} , $\mathcal{P} \models NC(KB) \wedge \mathcal{L} \rightarrow \varphi$.

In the following, $Ag(\varphi)$ will denote the set of agents who have φ as effective goal:

Definition 10. $Ag(\varphi) = \{a_i \in \mathcal{A} : \Sigma \models I(\varphi|D(KB))$ et KB is a context for φ for $a_i\}$.

3 Example

Let us consider a group of two agents a_1 and a_2 and assume that the preferences imposed to the group $\{a_1, a_2\}$ are the following: if the door is sanded, then it should be lacquered and not covered with paper and if the door is not sanded, then it should be covered with paper and not lacquered.

The representation of this scenario is the following: $\mathcal{P} = \{I(l \wedge \neg p|s), I(p \wedge \neg l|\neg s)\}$. For each model of \mathcal{P} , $I(l \wedge \neg p|s)$ means that there is a world which satisfies s and such that all preferred worlds satisfy $s \rightarrow l \wedge \neg p$. $I(p \wedge \neg l|\neg s)$ means that there is a world which satisfies $\neg s$ and such that all preferred worlds satisfy $\neg s \rightarrow \neg l \wedge p$.

1. Suppose that $KB = \{s, \neg l, \neg p\}$ i.e. the door is sanded but not lacquered nor covered with paper. We have $Cl(KB) = KB$.

Suppose that $\neg s$ is uncontrollable by the agents (i.e. the agents have no “means” to unsand the door). Furthermore, suppose that $C_{a_1} = \{l\}$ and that $C_{a_2} = \{p, \neg p\}$ (i.e. a_1 can lacquer the door, a_2 can cover it with paper or remove the paper if necessary). In this case, $NC(KB) = \{s\}$, because KB is a context for l and for p . $l \wedge \neg p$ is a CK goal of the group⁵.

If the agents do not commit themselves to anything, $D(KB) = \{s\}$, and then $EGoal_{a_1}(l)$ and $EGoal_{a_2}(\neg p)$ hold. a_1 has for atomic goal set $\{l\}$ (i.e. its only goal is to lacquer the door) and a_2 has $\{\neg p\}$ for atomic goals set (i.e. its only goal is not to cover the door with paper). This implies $Ag(l) = a_1$, $Ag(\neg p) = a_2$ and $Ag(\neg l) = Ag(s) = Ag(\neg s) = Ag(p) = \emptyset$.

⁵ In fact, it is theonly one that is interesting. We can also deduce for instance that $(l \wedge \neg p) \vee p$ is a CK goal of the group.

2. Suppose now that $KB = \{\neg s, \neg l, \neg p\}$, $C_{a_1} = \{l, \neg l\}$ and $C_{a_2} = \{s, p, \neg p\}$. In this case, $NC(KB) = \phi$ and $(l \wedge \neg p) \vee (\neg l \wedge p)$ is a CK goal of the group.

If $D(KB) = \phi$ (i.e. the agents do not commit themselves to anything), no effective goal can be derived, because a_2 controls s and could make s true.

But if a_2 commits itself not to achieve s (i.e. it commits itself not to sand the door), then $Com_{-}(\{a_1, a_2\} = \{\neg s\}$ and $EGoal_{a_2}(p)$ and $EGoal_{a_1}(\neg l)$ can be deduced: a_2 has for effective goal to cover the door with paper and a_1 has for effective goal to keep the door unlacquered. I.e $Ag(p) = a_2, Ag(\neg l) = a_1$ and $Ag(l) = Ag(\neg p) = Ag(s) = Ag(\neg s) = \emptyset$.

4 Strategies

The process described in the previous sections allocates goals to agents by taking into account their ability and their commitments. Here, we show how to extend this process in order to take into account more characteristics of the agents (like for instance, their competence, their cost or the required duration for achieving a goal). But, in order to be as general as possible, these characteristics are represented by preference order among the agents and are associated with each goal. These preference orders are called strategies.

4.1 Mathematical Preliminaries and Notations

Definition 11. Let E be a set. \leq_E is an order on E iff \leq_E is a reflexive, anti-symmetrical and transitive relation on E .

Definition 12. Let E be a set and \leq_E an order on E . Then $\min_{\leq_E}(E) = \{e_i \in E : \forall e_j \in E e_j \leq_E e_i \Rightarrow e_j = e_i\}$.

We define also the minimum of a set for a family of orders.

Definition 13. Let E be a set and $\leq_E = \{\leq_E^i : i \in \{1, \dots, n\}\}$ a set of orders on E . Then $\min_{\leq_E}(E) = \bigcap_{i \in \{1, \dots, n\}} \min_{\leq_E^i} E$.

4.2 Notion of Strategy

Definition 14. A strategy is a function $\mathcal{S} : lit(PROP) \rightarrow \mathcal{A} \times \mathcal{A}$ such that for any literal l , $\mathcal{S}(l)$ is an order $\leq_{\mathcal{S}(l)}$ on $Ag(l)$.

Being a function, a strategy associates a literal with at most one order which will be used to select one or several agents. For instance, let us consider a group of three agents $\{a_1, a_2, a_3\}$ achieving a task l . We know that a_1 and a_2 are more competent than a_3 to do l . We can define a strategy \mathcal{S} reflecting this relative level of competence by imposing that $a_1 \leq_{\mathcal{S}(l)} a_3$ and $a_2 \leq_{\mathcal{S}(l)} a_3$ hold.

4.3 Effective Goals

The notion of effective goals can then be refined by taking into account the notion of strategy as follows:

Definition 15. $\mathcal{A}' \subseteq \mathcal{A}$ is optimal for l according to the strategy \mathcal{S} iff $\mathcal{A}' = \min_{\leq_{\mathcal{S}(\varphi)}} Ag(l)$

This is denoted by $OGoal_{\mathcal{A}'}^{\mathcal{S}}(l)$. This means intuitively that \mathcal{A}' is the subgroup of agents preferred according to \mathcal{S} in order to achieve l .

Let us notice some basic properties:

- As $\mathcal{A}' \subseteq Ag(l)$, every agent in \mathcal{A}' is such that l is an effective goal for it;
- Consider a literal l such that l is an effective goal for only one agent. In this case, according to the previous definition, this agent will be optimal for l whatever the strategy we consider (if we assimilate the agent and the subgroup constituted by this single agent).
- Let l be a literal which is not an effective goal. In this case, $Ag(l) = \emptyset$. Thus, for any strategy \mathcal{S} , $\min_{\leq_{\mathcal{S}(l)}} Ag(l) = \emptyset$. So $OGoal_{\emptyset}^{\mathcal{S}}(l)$ holds and no agent is optimal for l .

4.4 Families of Strategies

We present in the following two main classifications of strategies.

Selective and Non-selective Strategies. Selective strategies are strategies which select a single agent among the agents for which l is an effective goal.

Definition 16. A strategy \mathcal{S} is a selective strategy for l iff $|\min_{\leq_{\mathcal{S}(\varphi)}} Ag(\varphi)| = 1$.

Example 1. Let us resume the example provided in section 3. Suppose that $KB = \{\neg p, \neg l, \neg r\}$, $C_{a_1} = \{p, l\}$ and $C_{a_2} = \{l, r, \neg r\}$. If a_1 commits itself to do p and a_2 commits itself to do l , then $D(KB) = \{p, l\}$. Thus $OGoal_{a_1}(p \wedge l)$ and $OGoal_{a_2}(l \wedge \neg r)$. Both a_1 and a_2 have l for effective goal.

First, notice that as $Ag(p) = \{a_1\}$ and $Ag(\neg r) = \{a_2\}$, for every strategy (S) , $OGoal_{\{a_1\}}^S(p)$ and $OGoal_{\{a_2\}}^S(\neg r)$ hold.

Consider here a selective strategy \mathcal{S} . Suppose that $\min_{\leq_{\mathcal{S}(l)}} Ag(l) = \{a_1\}$, then $OGoal_{\{a_1\}}^{\mathcal{S}}(p)$, $OGoal_{\{a_2\}}^{\mathcal{S}}(\neg r)$ and $OGoal_{\{a_1\}}^{\mathcal{S}}(l)$ hold.

Non-selective strategies are strategies which allocate a goal to several agents.

Definition 17. A strategy \mathcal{S} is a non-selective strategy for l iff $|\min_{\leq_{\mathcal{S}(\varphi)}} Ag(\varphi)| > 1$.

Example 2. In the previous example, suppose now that \mathcal{S} is a non-selective strategy for l , then $|\min_{\leq_{\mathcal{S}(\varphi)}} Ag(\varphi)| > 1$. But $Ag(\varphi) = \{a_1, a_2\}$, thus $\min_{\leq_{\mathcal{S}(\varphi)}} Ag(\varphi) = \{a_1, a_2\}$. In this case, we cannot deduce that $OGoal_{\{a_1\}}^{\mathcal{S}}(l)$ nor $OGoal_{\{a_2\}}^{\mathcal{S}}(l)$ holds. But $OGoal_{\{a_1, a_2\}}^{\mathcal{S}}(l)$ holds.

Voluntary and Non-voluntary Strategies. Voluntary strategies assign a task to the agents which committed themselves to achieve it. The formal definition is the following:

Definition 18. *Let l be a literal and \mathcal{S} a strategy. \mathcal{S} is a voluntary strategy φ iff $\forall a_i \in Ag(l) \forall a_j \in Ag(l) a_i \leq_{\mathcal{S}(l)} a_j$ iff $Eng_+(a_i) \models l$ and $Eng_+(a_j) \not\models l$.*

By using such an order, all the agents in $\min_{\leq_{\mathcal{S}(\varphi)}} Ag(l)$ commit themselves to achieve l . Non-voluntary strategies do not assign a goal to the agents which commit themselves not to achieve it. These strategies are less restrictive than the previous ones: an agent which did not commit itself to do l nor to not do l can be selected.

Definition 19. *Let l be a literal and \mathcal{S} a strategy. \mathcal{S} is a non-voluntary strategy for l iff $\forall a_i \in Ag(l) \forall a_j \in Ag(l) a_i \leq_{\mathcal{S}(l)} a_j$ iff $Eng_-(a_i) \not\models l$ and $Eng_-(a_j) \models l$.*

4.5 Combining Strategies

We can wonder on what we will define strategies. The first possibility is to use “primitive” strategies, i.e. strategies which are defined on only one criteria. This criteria can be for instance the relative competence of the agents, the cost of each agent in term of resources or the time an agent will take in order to achieve the task.

There are of course lots of other primitive criteria on which a strategy can be based. Most important is the fact that “in real life”, such decisions are not taken considering only one primitive factor, but several criteria which are combined in order to determine the “best” agents to select. To take this into account, we have to combine strategies.

For doing so, we suggest to use a *priority relation* between strategies. This comes to associate levels of importance to criteria. For instance, we could want to choose the agents which are, for a given task, the most competent to achieve it and the less costly, assuming that the competence is a criteria which is more important than the cost.

In the following, we present a mathematical framework for combining strategies.

Our objective is the following: we consider two orders \leq_1 and \leq_2 on the same set E and we want to obtain one or several orders $\leq_{1 \circ 2}$, called *orders combined considering \leq_1 having priority on \leq_2* , which verify first the order \leq_1 and then are “completed” by a part of \leq_2 . We suggest to use the technique developed in belief bases priority merging [9] by representing the order relation by a binary predicate of a first order logic.

Definition 20. *Let $E = \{e_1, \dots, e_n\}$ be a finite set. Let $\leq_E = \{\leq_i : i \in \mathbb{N}\}$ the set of possible orders on E . E and \leq_E are represented by the first-order language \mathcal{L}_E and the theory \mathcal{T}_E defined as in the following:*

1. the language \mathcal{L}_E is constituted by classical logical symbols (an enumerable set of variables, connectives, quantifiers), a set of constants symbols defined

by $\{e_1, \dots, e_n\}$, a set of predicate symbols $\{\preceq_i : \preceq_i \in \leq_E\} \cup \{=\}$ where each \preceq_i and $=$ are binary predicate symbols.

$$2. \mathcal{T}_E = \{\neg(e_i = e_j) : (i, j) \in \{1 \dots n\}^2 \ i \neq j\} \cup \bigcup_{\preceq_i \in \leq_E} \text{RAT}(\preceq_i)$$

where $\text{RAT}(\preceq_i) = \{\forall x \preceq_i(x, x), \forall x \forall y \preceq_i(x, y) \wedge \preceq_i(y, x) \rightarrow x = y, \forall x \forall y \forall z \preceq_i(x, y) \wedge \preceq_i(y, z) \rightarrow \preceq_i(x, z)\}$.

The theory \mathcal{T}_E lists the Unique Name Axioms and the mathematical properties of orders. For the sake of simplicity, we will denote $\preceq_i(x, y)$ by $x \preceq_i y$ in the following.

When someone wants to represent an order on a set, he/she does not describe the order by extension. On the contrary, he/she gives the relations which are verified by the elements of the set, the remaining relations are deduced by using the mathematical properties of orders. Thus, we will consider a set of *explicit* literals which will allow to generate the whole order (by using transitivity, antisymmetry and reflexivity).

For instance, if we consider a set $E_1 = \{a_1, a_2, a_3\}$, then the explicit set $\{a_1 \leq_{E_1} a_2, a_2 \leq_{E_1} a_3\}$ allows to build the order \leq_{E_1} on $\{a_1, a_2, a_3\}$ such that $a_1 \leq_{E_1} a_2, a_1 \leq_{E_1} a_3, a_2 \leq_{E_1} a_3, a_1 \leq_{E_1} a_1, a_2 \leq_{E_1} a_2, a_3 \leq_{E_1} a_3, a_2 \not\leq_{E_1} a_1, a_3 \not\leq_{E_1} a_2$ and $a_3 \not\leq_{E_1} a_1$.

We will characterize orders by generating them from explicit sets associated to a theory.

Definition 21. Let E be a set, \mathcal{L}_E and \mathcal{T}_E as previously defined. \mathcal{E}_i , set of formulas of the kind $e_i \preceq_i e_j$ with $e_i \in E$ and $e_j \in E$ is an explicit set iff $\text{Cl}(\mathcal{T}_E \cup \mathcal{E}_i)$ is consistent.

The order \preceq_i on E called order generated by \mathcal{E}_i is defined by: $\forall e_i \in E \forall e_j \in E \ e_i \leq_E e_j$ iff $\text{Cl}(\mathcal{T}_E \cup \mathcal{E}_i) \vdash e_i \preceq_i e_j$

It is easy to prove that we obtain an order by using the axioms in \mathcal{T}_E .

For instance, in the previous example, we can see that $\{a_1 \preceq_{E_1} a_2, a_2 \preceq_{E_1} a_3\}$ is an explicit set generating \leq_{E_1} .

We now define how to combine two orders, one having priority on the other. We use the explicit sets defining the orders for building maximal consistent sets of first-order formulas. Notice that we can obtain several explicit sets.

Definition 22. Let E be a set and \mathcal{L}_E the first-order language associated with E . Let $\mathcal{E}_1 \ \mathcal{E}_2$ two explicit sets generating respectively the orders \leq_1 and \leq_2 on E . We note $\mathcal{E}_{1 \rightarrow 1 \circ 2} = \{e_i \preceq_{1 \circ 2} e_j : e_i \preceq_1 e_j \in \mathcal{E}_1\}$ and $\mathcal{E}_{2 \rightarrow 1 \circ 2} = \{e_i \preceq_{1 \circ 2} e_j : e_i \preceq_2 e_j \in \mathcal{E}_2\}$.

The explicit set $\mathcal{E}_{1 \circ 2}^i$ is defined by $\mathcal{E}_{1 \circ 2}^i = \{e_j \preceq_{1 \circ 2}^i e_k : (e_j \preceq_{1 \circ 2} e_k) \in (\mathcal{E}_{1 \rightarrow 1 \circ 2} \cup \mathcal{E}_{2 \rightarrow 1 \circ 2}^i)\}$ where $\mathcal{E}_{2 \rightarrow 1 \circ 2}^i$ is a maximal subset of $\mathcal{E}_{2 \rightarrow 1 \circ 2}$ such that $\mathcal{E}_{1 \rightarrow 1 \circ 2} \cup \mathcal{E}_{2 \rightarrow 1 \circ 2}^i \cup \mathcal{T}_E$ is consistent.

We note $\preceq_{1 \circ 2}^i$ the order on E generated by $\mathcal{E}_{1 \circ 2}^i$ and we denote by $n_{1 \circ 2}$ the number of different orders we can obtain from \leq_1 and \leq_2 by giving priority to \leq_1 .

If the different orders can be reduced to a single order, we will note $\leq_{1\circ 2}$ this order.

Definition 23. *If $\exists i \in \{1, \dots, n_{1\circ 2}\}$ such that $\forall j \in \{1, \dots, n_{1\circ 2}\} \mathcal{E}_{1\circ 2}^j \subseteq \mathcal{E}_{1\circ 2}^i$, then we note $\mathcal{E}_{1\circ 2} = \mathcal{E}_{1\circ 2}^i$.*

Example 3. Let us consider $E = \{e_1, e_2, e_3\}$ and examine some examples:

Suppose that \leq_1 is generated by $\{e_3 \preceq_1 e_2\}$ and \leq_2 is generated by $\{e_2 \preceq_2 e_3, e_1 \preceq_2 e_2\}$. Then $\mathcal{E}_{1 \rightarrow 1\circ 2} = \{e_3 \preceq_{1\circ 2} e_2\}$ and $\mathcal{E}_{2 \rightarrow 1\circ 2} = \{e_2 \preceq_{1\circ 2} e_3, e_1 \preceq_{1\circ 2} e_2\}$. The only subset of $\mathcal{E}_{2 \rightarrow 1\circ 2}$ consistent with $\mathcal{E}_{1 \rightarrow 1\circ 2} \cup \mathcal{T}_E$ is $\{e_1 \preceq_{1\circ 2} e_2\}$ (because $\mathcal{E}_{1 \rightarrow 1\circ 2} \cup \mathcal{T}_E \vdash \neg e_2 \preceq_{1\circ 2} e_3$), thus we obtain an order $\leq_{1\circ 2}$ generated by $\{e_3 \preceq_{1\circ 2} e_2, e_1 \preceq_{1\circ 2} e_2\}$. In this case $\min_{\leq_{1\circ 2}} E = \{e_1, e_3\}$.

Suppose now that \leq_1 is generated by $\{e_3 \preceq_1 e_2\}$ and \leq_2 is generated by $\{e_1 \preceq_2 e_3, e_2 \preceq_2 e_1\}$. Then $\mathcal{E}_{1 \rightarrow 1\circ 2} = \{e_3 \preceq_{1\circ 2} e_2\}$ et $\mathcal{E}_{2 \rightarrow 1\circ 2} = \{e_1 \preceq_{1\circ 2} e_3, e_2 \preceq_{1\circ 2} e_1\}$. There are two maximal consistent subset of $\mathcal{E}_{2 \rightarrow 1\circ 2}$ consistent with $\mathcal{E}_{1 \rightarrow 1\circ 2} \cup \mathcal{T}_E$. Thus we obtain two orders: $\leq_{1\circ 2}^1$, generated by $\{e_3 \preceq_{1\circ 2}^1 e_2, e_1 \preceq_{1\circ 2}^1 e_3\}$ and $\leq_{1\circ 2}^2$, generated by $\{e_3 \preceq_{1\circ 2}^2 e_2, e_2 \preceq_{1\circ 2}^2 e_1\}$. In this case, $\min_{\leq_{1\circ 2}} E = \emptyset$.

5 Example

Let us resume the example in section 3 and consider a group of three agents $\{a_1, a_2, a_3\}$. Let us suppose that $KB = \{\neg s, \neg l, \neg p\}$, that $C_{a_1} = \{s, l\}$, that $C_{a_2} = \{l, p, \neg p\}$ and that $C_{a_3} = \{l\}$. if a_1 commits itself to do s , that a_2 and a_3 commit themselves to do l , then $D(KB) = \{s, l\}$. Thus $OGoal_{a_1}(s \wedge l)$, $OGoal_{a_2}(l \wedge \neg p)$ and $OGoal_{a_3}(l)$ hold. The three agents have to lacquer the door.

Let us suppose that we do not want that several agents have the same task for efficiency reason. We have to find a selective strategy to select only one agent.

A voluntary strategy \mathcal{S}_V for l gives the following order: $a_2 \leq_{\mathcal{S}_V(l)} a_1$ and $a_3 \leq_{\mathcal{S}_V(l)} a_1$. This strategy is not sufficient to select a single agent because it cannot choose between a_2 and a_3 .

Let us suppose that there is a strategy \mathcal{S}_E for l which reflects the relative efficiency of the agents to achieve l : $a_1 \leq_{\mathcal{S}_E(l)} a_2$ and $a_2 \leq_{\mathcal{S}_E(l)} a_3$. In this case, there are two solutions: either $\leq_{\mathcal{S}_E(l) \circ \mathcal{S}_V(l)}$ is chosen (the efficiency of the agent is privileged) and thus a_1 is optimal for l , either $\leq_{\mathcal{S}_V(l) \circ \mathcal{S}_E(l)}$ is chosen (the voluntary agents are privileged) and in this case a_2 is optimal for l .

6 Discussion

This work focused on determining the individual goals of agents from goals addressed to a team of agents, a representation of agents and strategies. In order to do that, we have relied on the support of some previous work and Boutilier's work on qualitative decision theory. We have defined the notion of strategy for allocating tasks to a sub-team of agents and we have shown how to

combine strategies in order to refine the allocation process. We are aware that, as for the strategies combination method, we could have used another one like, for instance an arbitration method [10]. It would have come to select the “less worst” agent given all the primitive criteria (this can be viewed as a maximin selection). This present work does not contribute in combination techniques. Its originality concerns the use of the Qualitative Decision Logic to the case of a team of several agents and the extension of the model of agents since we consider their ability, their commitments and, through the notion of strategy, any other characteristics we want.

However, this work is rather preliminary and it could be extended in several ways.

First, instead of having an unique set KB which represents a common point of view about the real world, we could consider that the agents do not share the same beliefs about the real world. In the worst cases, these beliefs may happen to be contradictory and belief bases merging techniques (cf. [10, 11]) could be used in order to solve the conflicts. We could also consider that there is no central entity and that the agents communicate in order to inform the others about their commitments.

We also intend to work on the notion of strategy in order to obtain general properties on strategies and define global strategies. Moreover, the present strategies are defined for literals only and we could envisage to define them to propositions. However, in this case, relations between for instance $\mathcal{S}(l)$, $\mathcal{S}(l')$ and $\mathcal{S}(l \wedge l')$ should be defined.

References

1. Kraus, S., Shehory, O.: Methods for task allocation via agent coalition formation. *Artificial Intelligence* **101** (1998) 165–200
2. Barès, M.: Formal approach of the interoperability of C4IRS operating within a coalition. In: *Proceedings of the 5th International Command and Control Research and Technology Symposium*. (2000)
3. Boutilier, C.: Toward a logic for qualitative decision theory. In Doyle, J., Sandewall, E., Torasso, P., eds.: *Principles of Knowledge Representation and Reasoning (KR'94)*, Morgan Kaufmann (1994) 75–86
4. Cholvy, L., Garion, C.: Distribution of goals addressed to a group of agents. In Rosenschein, J.S., Sandholm, T., Wooldridge, M., Yokoo, M., eds.: *Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems*, ACM Press (2003) 765–772
5. Kraus, S.: Negotiation and cooperation in multi-agent environments. *Artificial Intelligence Journal - Special Issue on Economic Principles of Multi-Agent Systems* **94** (1997) 79–98
6. Grosz, B., Hunsberger, L., Kraus, S.: Planning and acting together. *AI Magazine* **20** (1999) 23–34
7. Soh, L., Tsatsoulis, C.: Allocation algorithms in dynamic negotiation-based coalition formation. In: *Workshop on Teamwork and Coalition Formation (affiliate workshop of 1st International Conference On Autonomous Agents and MultiAgent Systems)*. (2002)

8. Boutilier, C.: Conditional logics of normality : a modal approach. *Artificial Intelligence* **68** (1994) 87–154
9. Cholvy, L.: Reasoning about merged information. In Gabay, D., Smets, P., eds.: *Handbook of Defeasible Reasoning and Uncertainty Management Systems*. Volume 3. Kluwer Academic Publishers (1998) 233–263
10. Konieczny, S., Pino-Pérez, R.: Merging information under constraints: a qualitative framework. *Journal of Logic and Computation* **12** (2002) 773–808
11. Cholvy, L., Garion, C.: Answering queries addressed to several databases according to a majority approach. *Journal of Intelligent Information Systems* **22** (2004) 175–201