A Scheme for Agent Collaboration in Open Multiagent Environments

Ei-Ichi Osawa
Sony Computer Science Laboratory Inc.
3-14-13 Higashi-gotanda
Shinagawa-ku, Tokyo, 141
JAPAN

Abstract

In multiagent planning, an agent sometimes needs to collaborate with others to construct complex plans, or to accomplish large organizational tasks which it cannot do alone. Since each agent in a group may have incorrect beliefs about the world and incomplete knowledge, and because agent's abilities differ, constructing a coordinated collaborative plan among agents is a difficult proposition. In previous work [Osawa and Tokoro 92], we developed a scheme for constructing collaborative plans from the. possibly incomplete, individual plans of agents. This scheme was designed to provide availability-based assignment of goals to agents, and opportunistic collaboration to distributed planning in open multiagent environments based on the contract net. In this paper, we formalize incomplete individual plans and collaborative planning among rational agents using the Multi-World Model, and provide a utility-based model for rational choice of actions. Agents can effectively balance workloads based on the utility theory. A condition for incomplete collaborative plans is also presented.

1 Introduction

In multiagent planning, agents try to achieve goals, which can be independent, shared, or competitive. Researchers have attempted to address the problem of coordinating interacting plans so as to increase efficiency. The subject of coordination has been of continuing interest in multiagent planning [Corkill 79, Georgeff 83, Zlotkin and Rosenschein 90, Martial 90. Osawa and Tokoro 92, Ephrati and Rosenschein 92, Kinny e*t al.* 92]. Martial has investigated how planning agents can positively cooperate in distributed environments [Martial 90]. Many previous papers on distributed coordinated planning mainly focused on how to resolve conflicts [Corkill 79, Georgeff 83]. Martial, however, studies situations where a positive effect can be reached, as modeled by his *favor relation.* We also focus on the positive effect of cooperation in terms of collaborative plan construction, and have developed a scheme for constructing collaborative plans among agents based upon their, possibly incorrect, beliefs and partial (incomplete) knowledge of the world [Osawa and Tokoro 92]. The partiality of agents[1] skills and inconsistencies among agents[1] beliefs in open multiagent environments are not well treated in most of the previous work involving multiagent planning.

In Martial's method, agents broadcast their plans at any time and at different levels of abstraction, so that they may refine their plans in a coordinated way. His method is based on the assumption that there is a collection of autonomous intelligent agents which communicate about, planned actions ahead of time. In our scheme, the investigation of possible positive cooperation (collaboration) is taken into account when the need for help actually arises. The scheme is designed to provide flexible decomposition of goals, availability-based assignment of goals to agents, and opportunistic collaboration to distributed planning based on the contract net proposed by Davis and Smith [Davis and Smith 83]. These features of the scheme are designed to cope with the uncertainty and dynamic nature of open multiagent environments.

In the collaborative planning, agents are presumed to rationally take three factors (obstacle elimination, workload balancing, and cost effectiveness) into account to decide which actions will be performed by each agent in a collaborative plan. The rationality makes it possible for every agent to make a good choice among alternative plans in individual plan construction. Additionally, it enables agents to expect certain decisions and behavior from other agents in collaborative activities. The role of rationality in collaborative planning is illustrated in previous work [Osawa and Tokoro 92], however, the formal treatment, of collaboration and the rational decisions among agents still needs to be addressed.

In this paper, we formalize collaborative planning among rational agents using the Multi World Model [Nagao 93], and provide a utility-based model for rational choice of actions with which agents can effectively balance the above mentioned factors. A condition for incomplete collaborative plans is also presented.

The organization of this paper is as follows. In Section 2, we will present the outline of the collaborative plan scheme proposed in [Osawa and Tokoro 92]. Section 3 gives a formal model of individual planning of rational agents based on the Multi-World Model [Nagao 93]. In Section 4, we formalize the process of investigating the possibility of collaboration among individual plans. Section 5 gives a formal model for a rational choice of actions from the initial individual plan. Section G contains our conclusions. Relation to other work has already been discussed in this section.

2 Collaborative Plan Scheme Outline

Large, multiagent systems can be viewed as open distributed environments. Thus, agents have inconsistent and partial world views. In multiagent cooperative plan construction, several agents mutually generate collaborative plans by inference based on their own beliefs and partial knowledge about the world. Therefore, mutual planning is confounded by disparities in agents' world knowledge.

In a multiagent system, an agent may have a goal or task which it cannot do alone. *Contract-net* protocol [Davis and Smith 83] provides a way for an agent who needs help (the requestor) to dynamically decompose the task into subtasks, and to allocate the subtasks to other agents (requestees) through negotiation. The contract-net protocol also provides dynamic and opportunistic control.

In open distributed environments, services, processing capacity, and the connection topology of computing elements are continuously changing. At the same time, the granularity of agents and plans are changing dynamically. Also agents are heterogeneous. Although contract-net type organization schemes are usually preferable in open distributed environments because of their dynamic nature, a multiagent system embodies additional complexity which makes application of the contract-net difficult.

Two such problems in the contract-net occur in decomposition and task allocation. When the requestor first decomposes the task, its fixed decomposition of the task may not suit the open distributed environment,. Not only may it not know what agents are currently available, but it also may not know the changing skills of potential requestees. The requestor then selects one agent per subtask through negotiation, and allocates the subtask to that agent. No single agent may have a plan to achieve the subtask alone. Even though subcontracting is possible, this fixed task allocation strategy, which assigns a subtask to only one agent, may result in an ineffective hierarchy of subcontracts.

If we apply the contract-net protocol to hierarchical multiagent planning, the problems become more serious. The requestor wants some agent to accomplish a goal, but if it does not have sufficient knowledge to decompose a complex goal properly in an open distributed environment, it cannot ask any single agent to achieve the goal. Its task allocation strategy fails. Therefore, we need a more flexible strategy for selecting requestees.

Suppose that the requestor can somehow select several agents as collaborative requestees. This raises some questions. What information should the requestor provide to those requestees? In other words, what information is necessary for the requestees to mutually construct collaborative plans? Also, how should the mutual plan construction be coordinated and organized?

The scheme proposed in [Osawa and Tokoro 92] is designed to provide flexible decomposition of goals, availability-based assignment of goals to agents, and opportunistic collaboration to distributed planning partially based on the contract net proposed by Davis and Smith [Davis and Smith 83]. These features of the scheme are designed to cope with uncertainty and the dynamic nature of open multiagent environments.

In the collaborative plan scheme, an agent who needs help dynamically organizes a group. The agent first announces a request for proposals (RFP) by sending a message to a bulletin board agent. Agents who read the RFP and can construct an individual plan for the request, even if incomplete, send their individual plans to the originating agent, hereafter referred to as the requestor. The requestor then investigates possible collaboration among potential requestees. If collaboration seems possible, the requestor gives collaborative awards, along with *suggestions for collaboration,* to the requestees. A suggestion for collaboration given to a requestee agent contains: (1) Explicit obstacles of the other collaborating agents which the agent may possibly resolve; (2) Actions which collaborating agents may perform. The suggestions set up a partial model for predicting the other agent's actions. Using these suggestions, along with its initial individual plan and beliefs, each collaborating agent constructs a collaborative plan through inference. In collaborative plan construction, each agent decides on the actions it should perform, the actions the other agents would perform, and the actions both agents will achieve jointly. In the process, each agent takes three factors into account: the elimination of obstacles of other agents, balancing of the workload among agents, and cost effectiveness. This whole process can be summarized as follows.

1. Requestor sends a request for proposal (RFP) to the bulletin board agent
2. Free agents[1] request the bulletin board to provide a stored RFP
3. Bulletin board sends RFP to requesting free agents
4. Free agents generate individual plans
5. Free agents send individual plans, if any. to the requestor

G. Requestor investigates the possibility of collaboration (computes suggestions for collaboration)

7. Requestor sends collaborative awards to requestees (out. of free agents)
8. Requestees construct collaborative plans

3 Individual Planning of Rational Agent

In this section, we present a. model for individual plan generation by rational agents. In the model, the beliefs of an agent at time *t* are modeled by a first-order axiomatic system, which is called a world. Operators are represented by a transition from one world to another. Therefore, a plan can be viewed as a chain of operators which connect several worlds. This model is based on the *Multi-World Model* [Nagao 93].

3.1 Belief Model of Agent

Definition 1 (Belief) The set of beliefs of agent a at time t, w_a^t, is a first-order axiomatic system.

Definition 2 (Entailment) $w_a^t \vdash p$ means that world w_a^t entails proposition p.

---

[1] A free agent is one without a current task.

$\vdash$ is used for entailment from a single world. Later, we introduce entailment based on consistent inheritance of propositions from previous worlds, which is called entailment with consistent inheritance.

3.2 Operator

Each agent maintains a library of operators that it may execute. Operators in the library are generic functions that are represented in the following form:

$$\begin{aligned}&\mathbf{op}(Agent, Parameters, T, \tau)\\ &\quad precond\text{: } pre_1, \cdots, pre_n\\ &\quad effect\text{: } eff_1, \cdots, eff_m,\end{aligned}$$

where $pre_1, \cdots, pre_n$ $(eff_1, \cdots, eff_m)$ are propositions that hold before (after) the operator is executed. Preconditions and effects are sometimes written as *preconds* and *effects.* Also, each operator is associated with a temporal variable T, and an execution time cost r which is the expected cost of the operator. The cost of the operator is predicted from an agent's working environment.

The arguments of an operator in the library, *Agent, Parameters, T, preconds,* and *effects,* are instantiated when the operator is invoked.

The functions *agt, pars, time, cost, pre,* and *eff,* which are used in the following definitions, are functions that take an instantiated operator and return its respective argument, *agent, parameters*, *t, T, preconds,* and *effects.*

Operators are defined as follows.

Definition 3 (Operator (definition attempt)) Operator op of agent a is a transition from world w_a^t to world $w_a^{t+\tau}$, if and only if worlds w_a^t and $w_a^{t+\tau}$ satisfy the following.

$$op : w_a^t \longrightarrow w_a^{t+\tau},$$

$$s.t.(\forall p \in pre(op)\ w_a^t \vdash p) \ \wedge \ (\forall p \in eff(op)\ w_a^{t+\tau} \vdash p)$$

3.3 Abduction

Abduction is a special kind of transition between worlds. An operation that translates a world into another world by introducing an hypothesis *p (p* is atomic) is called abduction.

Definition 4 (Abduction) Abduction $ab(a, p, t, \sigma)$ of agent a is a transition from world w_a^t to world $w_a^{t+\sigma}$, which entails proposition p.

$$ab : w_a^t \longrightarrow w_a^{t+\sigma}, s.t.\exists p\ w_a^{t+\sigma} \vdash p \wedge w_a^t \not\vdash p$$

The difference between an operator and an abduction is that the former is obtained by instantiating some generic function in the library, while the latter is not limited in that way. Abduction is used to introduce unsatisfied operator proposition *preconds* into a world*. These resulting propositions are called hypotheses. The hypothesis introduced by abduction is associated with its cost, since the hypothesis will be achieved by executing some operator. The semantics of the cost will be discussed in Subsection 3.6 of this section.

3.4 Operator Sequence

A transition from one world to another by way of a chain of operators and abductions is called an *operatar sequence*.

Definition 5 (Operator Sequence) An Operator sequence $[x_i]_{(i=1,\cdots,n)}$ that translates world w_a^t into world $w_a^{t'}$ is called an operator sequence and is represented by $w_a^t \Longrightarrow w_a^{t'}$.

Now, we define entailment with consistent inheritance and extend the definition of the operator. In the previous definition of the operator, the *preconds* of the operator are restricted to be solely entailed from the world in which the operator will be applied. However, if there is a chain of several worlds which are connected through a sequence of operators, not only the world in which the operator will be applied, but also some previous world in the chain, will entail a proposition in *preconds.* Therefore, we need to extend the definition of operators. For that purpose*, we first define entailment with consistent inheritance. This inference rule is analogous to the default rules in nonmonotonic reasoning [R.eiter 8()].

Definition 6 (Entailment with Consistent Inheritance) Proposition p is entailed from world w_a^t with consistent inheritance, $w_a^t \Vdash p$, if and only if

$$(\exists t'(\leq t)\ w_a^{t'} \vdash p) \ \wedge \ (\nexists t''(t > t'' > t')\ w_a^{t''} \vdash \neg p),$$

where there is some operator sequence $w_a^{t'} \Longrightarrow w_a^t$, except in the case where world $w_a^{t'}$ is identical to world w_a^t. Also, $w_a^{t''}$ is a world which exists in the chain of worlds linking $w_a^{t'}$ and w_a^t.

With this definition, we redefine the operators.

Definition 7 (Operator) Operator op of agent a is a transition from world w_a^t to another world $w_a^{t+\tau}$, and satisfies the following condition.

$$op : w_a^t \longrightarrow w_a^{t+\tau},$$

$$s.t.(\forall p \in pre(op)\ w_a^t \Vdash p) \ \wedge \ (\forall p \in eff(op)\ w_a^{t+\tau} \vdash p)$$

3.5 Plan

Definition 8 (Goal) Goal g of agent a is a set of first-order atomic formulae that agent a wants to satisfy.

Definition 9 (Plan) Let $w_a^{t_0}$ be the initial beliefs of agent a, and g be the goal of agent a. The plan of agent a that satisfies goal g, $plan(a, g)$, is the sequence of operators that satisfies the following condition.

$$plan(a, g) = [x_i]_{(i=1,\cdots,n)} : w_a^{t_0} \Longrightarrow w_a^t, s.t.\forall p \in g\ w_a^t \Vdash p$$

Definition 10 (Incomplete Plan) Let $plan(a, g)(= [x_i]_{(i=1,\cdots,n)})$ be a plan that satisfies goal g of agent a. If at least one x_i is abduction, $plan(a, g)$ is called an *incomplete plan.*

Let $\mathcal{I}$ denote a unary predicate over plans that is true if and only if its argument is incomplete.

In general, there can be several plans, including incomplete ones, that satisfy goal g of agent a. Let $PLAN(a, g)$ denote the set of these plans.

If we allow abduction, an agent may introduce arbitrary hypotheses, sonic of which might, be irrelevant to the agent's goal. In order to avoid abducting these irrelevant hypotheses, agents need to have a control strategy for abduction. This is done based on the cost of subplans, including abducted hypotheses, computed dynamically in the course of planning. With the cost, agents are able to calculate the utility of the goal. The details of this are discussed in Subsection 3.6

3.6 Cost, Worth, and Utility

We will now define the cost of plans, worth of goals, and utility of goals.

Definition 11 (Cost of Plan) The cost of plan $plan(a,g)(=[x_i]_{(i=1,\cdots,n)})$, $cost(plan(g))$, is calculated as follows:

$$cost(plan(a,g)) = \sum_{i=1,\cdots,n} cost(x_i),$$

where $cost(x_i)$ is the cost of operator x_i If x_i is an abduction, $cost(x_i)$ is the abduction cost.

Now, we more precisely characterize the abduction cost. Individual plans in the; collaborative planning scheme can be incomplete [Osawa and Tokoro 92]. As we stated above, an incomplete plan is a plan which includes hypotheses introduced by abduction. The cost of hypothesis can be viewed as the maximum expected cost that the agent will pay to satisfy the proposition. In other words, the cost of abducted hypothesis for agent *a* can be viewed as the *worth* of *p* for agent a. Worth can be given to any goal as well as any hypothesis.

Definition 12 (Worth of Goal) The worth of a goal for an agent is the maximum expected cost that the agent will pay to satisfy the goal[3]. Function *worth* is a binary function over agents and goals (or hypotheses) that designates the worth of the goal (or hypotheses) for the agent.

If we know the worth of a goal, we can define the utility of the goal.

Definition 13 (Utility of Goal) The utility of goal *g* for agent *a* is calculated by the following formula.

$$utility(a,g) = worth(a,g) - cost(plan(a,g))$$

3.7 Best Plan

Definition 14 (Best plan) The plan in *PLAN*(a,*g)* that has the minimal cost is called the *best plan* for goal *g* of agent *a.*

Definition 15 (Rational agent) Rational agent *a* chooses the best plan out of *PLAN(a.g)* for goal *g* as long as the utility of the. goal is positive. If there is no best plan, rational agents abandon trying to achieve the goal.

If a rational agent is asked to propose a plan for the goal, it will choose the best plan, and propose it as its own individual plan.

(Example) We will use the following example throughout this paper (see figure 1). The goal in this example is for agent a_3 to have block *b* in room r_3. We assume that agent a_3 knows that by performing trans(Agent, a_3,b), it can hold block b. However, since some parts of the precondition of the action, i.e. *(holding (Agent, b)* $\wedge$ *in(Agent,* r_3)), don't hold at this moment, it needs to ask other agents to achieve this goal, conditioned by the fact that block *b* is not in room r_3 at this moment. Therefore, the agent sends a RFP, which includes asking g_{ex} $= (holding(Agent,b) \wedge in(Agent,r_3))$ to be satisfied, to the bulletin board agent.

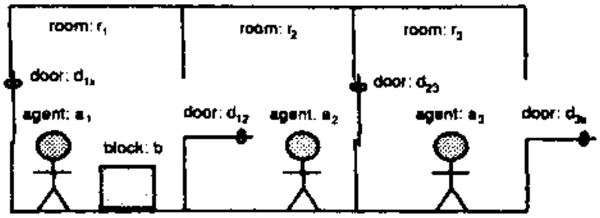

Figure 1: Moving a block between adjacent rooms

We assume that agent a_1 believes that agent a_1 and block b are in room r_1, and room r_1 and r_2 are adjacent, and the door between these two rooms, namely d_{12}, is open, and door d_{23} is closed. Also, we assume that agent a_1 can execute operations **pickup**$(a_1, Object)$, **move**$(a_1, Object, From, To)$, and **trans**$(a_1, Object, Recipient)$.

Suppose agent a_1 sends a request to the bulletin board agent, and receives the RFP. One possible plan agent a_1 may generate is [**pickup**(a_1, b), **move**(a_1, b, r_1, r_2), $open(door_{23})$, **move**(a_1, b, r_2, r_3)], where terms in italics indicate hypotheses introduced by abduction.

This plan is incomplete, since it includes hypothesis $open(door_{23})$ which is introduced by abduction. The hypothesis is a part of the preconditions of operator **move**(a_1, b, r_2, r_3)

4 Investigating and Awarding Collaboration

In this section, we formalize the process of investigating the possibility of collaboration among possible contractors described in Section 2. Throughout this section, we assume that for goal g of the requesting agent several bid plans $plan(a_i, g)$ are proposed. The set of proposed individual plans for goal g is denoted as $PPLAN(g)(=\{plan(a_i, g)\})$.

4.1 Filter

In collaborative planning, the requesting agent wants to choose the best plan, or best collaborative plan, out of $PPLAN(g)$. However if $PPLAN(g)$ includes many individual plans, the computational cost of choosing a plan may cause combinatorial explosion. Therefore, it is preferable to select promising plans from $PPLAN(g)$. *Plan filter*, defined below, provides the mean to do this.

Definition 16 (Plan Filter) Plan filter $filter$ is a binary function over a set of plans P and a unary predicate $\mathcal{R}$ that designates a subset of P, each of which satisfies predicate $\mathcal{R}$.

$$filter : P \times \mathcal{R} \longrightarrow P', s.t. \forall p \in P' \ \mathcal{R}(p)$$

With this filter, requesting agent a first chooses the set of plans out of $PPLAN(g)$ whose costs are lower than agent a's worth for goal g. The subset, $FPLAN(g)$, is obtained as follows:

$$filter(PPLAN(g), leq(X, worth(a,g))) \longrightarrow FPLAN(g)$$

---

[3] This view of worth is discussed in [Zlotkin and Rosenschein 89]. We generally follow their idea. Also, we assume that such an upper bound exists.

where $leq(X,Y) \longleftarrow X \leq Y$.

Let $CFPLAN(g)$ be the set of complete plans in $FPLAN(g)$.

$$CFPLAN(g) = \{p \mid p \in FPLAN(g) \ \wedge \ \neg\mathcal{I}(p)\}$$

4.2 Investigating he Possibility of Collaboration

Definition 17 (Collaborative Plan) Let $w_a^{t_0}$ be the initial belief of requesting agent a, and g be the goal that the agent requires to be satisfied. Collaborative plan $(plan(a_i, g); plan(a_j, g)) \mid_c$ of two agents $a_i, a_j (i \neq j)$ for goal g is defined as follows.

$$(plan(a_i, g); plan(a_j, g)) \mid_c : w_a^{t_0} \Longrightarrow w_a^t, s.t. \forall p \in g \ w_a^t \vdash p$$

where $plan(a_i, g), plan(a_j, g) \in PPLAN(g)$ and c represents a partial order of operators in $plan(a_i, g)$ and $plan(a_j, g)$.

Let $CPLAN(g)$ denote the set of collaborative plans, both of which belong to $PPLAN(g)$. Also, let $CCPLAN(g)$ denote the set of complete collaborative plans in $CPLAN(g)$.

$$CCPLAN(g) = \{p \mid p \in CPLAN(g) \ \wedge \ \neg\mathcal{I}(p)\}$$

4.3 Selecting the Best Plan

Requesting agent a selects the best individual plan or the best collaborative plan by means of the following procedure.

```
If $FPLAN(g) = \phi$
   then return null plan
   else  if $CCPLAN \cup CFPLAN \neq \phi$
         then return $p$,
               $s.t.\ cost(p) = \min_{q \in CCPLAN \cup CFPLAN} cost(q)$
         else return $p$,
               $s.t.\ cost(p) = \min_{q \in CPLAN \cup FPLAN} cost(q)$
```

4.4 Requesting Collaboration (Collaborative Award)

If the best plan is the null plan, agent a abandon trying to achieve goal g. If it is a single individual plan $plan(a_i, g)$, agent a requests a_i to execute $plan(a_i, g)$. If the best plan is a collaborative plan $(plan(a_i, g); plan(a_j, g)) \mid_c$, agent a requests two agents, a_i and a_j, to collaborate with each other. This is called a collaborative award. In this case, agent a informs both a_i and a_j of the collaborative plan $(plan(a_i, g); plan(a_j, g)) \mid_c$ as collaborative awards.

(**Example**) Looking back at the example given in the previous section, suppose three agents a_1, a_2, and a_4[4] propose the following plans $plan(a_1, g_{ex})$, $plan(a_2, g_{ex})$, $plan(a_4, g_{ex})$, respectively, for $g_{ex} = (holding(Agent, b), in(Agent, r_3))$ (agents a_2 and a_4 don't know the location of block b). We also assume that the worth of g_{ex} for agent a_3 is equal to 7, and the cost of all operators and the worth of all hypotheses in the following plans are equal to 1, for simplicity.

$$\begin{aligned}
plan(a_1, g_{ex}) = [&\mathsf{pickup}(a_1, b), \mathsf{move}(a_1, b, r_1, r_2), \\
&open(door_{23}), \mathsf{move}(a_1, b, r_2, r_3)] \\
plan(a_2, g_{ex}) = [&\mathsf{open}(a_2, door_{23}), \mathsf{move}(a_2, nil, r_2, r_1), \\
&onfloor(b, r_1), \mathsf{pickup}(a_2, b), \\
&\mathsf{move}(a_2, b, r_1, r_2), \mathsf{move}(a_2, b, r_2, r_3)] \\
plan(a_4, g_{ex}) = [&\mathsf{open}(a_4, door_{23}), holding(Agent, b), \\
&in(Agent, r_3), \mathsf{trans}(b, Agent, a_4), \\
&\mathsf{move}(a_4, b, r_1, r_2), \mathsf{move}(a_4, b, r_2, r_3)]
\end{aligned}$$

After reception of all of these plans, agent a_3 needs to select the best plan(s). In this case, $FPLAN$, $CPLAN$, $CCPLAN$, and $CFPLAN$ are given as follows.

$$\begin{aligned}
FPLAN &= \{plan(a_1, g_{ex}), plan(a_2, g_{ex}), plan(a_4, g_{ex})\} \\
CPLAN &= \{(plan(a_1, g_{ex}); plan(a_2, g_{ex})) \mid_{c_{12}}\} \\
CCPLAN &= \{(plan(a_1, g_{ex}); plan(a_2, g_{ex})) \mid_{c_{12}}\} \\
CFPLAN &= \phi
\end{aligned}$$

Therefore, agent a_3 selects collaborative plan $(plan(a_1, g_{ex}); plan(a_2, g_{ex})) \mid_{c_{12}}$, according to the procedure described in Subsection 4.3.

5 Rational Choice of Subplans from Individual Plan in Collaboration

An agent, who is given a collaborative award tries to construct its contribution according to the collaborative plan by refining the individual plan which it proposed. The refinement mainly consists of choosing subplans from the individual plan. A formal model of rational decision with which agents can effectively choose their actions is presented in this section. We first define two meta-operations, *hypothesize* and *commit,* on operators in plans. Second, we show criteria with which agents decide[1] what actions they will execute. With these meta-operations and criteria, we finally present how agents rationally choose their actions in collaborative planning.

5.1 Hypothesizing and Committing

We define two operations, *hypothesize* and *commit** which are utilized in collaborative planning.

Definition 18 (Hypothesizing)
Operation *hypothesize* takes operator $op \ (w_a^t \longrightarrow w_a^{t+\tau})$ and makes it, an abduction *ah* with zero cost.

Suppose a certain operator is included in both individuals' plans. If one of the collaborating agents executes the operator, the other agent can view the cost of the operator as zero.

Definition 19 (Committing) The *commit* operation *(commit(a.op))* commits agent *a* to execute operation *op.*

Committing will be applied to an operator which supports a hypothesis of the other collaborating agents. Operators which are committed cannot be hypothesized.

5.2 Criteria for Choosing Actions from Individual Plan

The following two criteria are taken into account when agents choose actions from their initial individual plans.

- Obstacle detection and elimination: Hypotheses included in individual plans can be regarded as explicit obstacles to the agent's plan, since they can

---

[4] Although agent a_4 doesn't appear in the previous example, we assume the existence of the agent for convenience.

be unsatisfied preconditions of a certain operator. The subplan of the other collaborating agent, which achieves the hypotheses, needs to be chosen by the agent. If the subplan is not chosen, the hypotheses remain unsatisfied, and the overall collaborative plan remains incomplete. If the collaborative plan is incomplete, a new collaborative plan is formed among agents using the the same protocol. The goal of this plan is the unsatisfied hypotheses of the first collaborative plan, in other words other agents help to complete the first plan. Since this entire process can be expensive, it would be better for the collaborating agent, whose subplan supports the hypotheses of the other agent, to choose the subplan to be executed. The choice can be regarded as *obstacle, elimination.*

- Workload balancing: Each collaborating agent estimates the worth of the goal which collaborating agents are trying to achieve. Collaborating agents are willing to expend effort to achieve the goal, however the effort should not exceed the worth of the goal for the particular agent. Therefore, the workload of each collaborating agent should be balanced according to the utility which each agent will gain from the achievement of the goal. This should be done based on the utility equalization principle, which is defined below.

[Utility Equalization Principle] In collaboration among agents, the utility which each agent will gain from the collaborative goal should be as equal as possible.

To make this principle operational, collaborating agents need to know the worth of the goal to the partner. In the following discussion, we assume that the worth of the goal for each agent, *worth*((a_2,g) and *worth*(a_1. g). are known to both collaborating agents, a_1 and a_2.

5.3 Choosing Actions from Individual Plan

We describe how the collaborating agents choose operators from their initial individual plans. In the following description, agent a_1's choice is described. The choice process of agent a_2 is identical. We assume that the collaborating plan is denoted as $(plan(a_1, g) : plan(a_2, g)) |_c$, and each plan *plan*$(a_1, g) = [x_i^1]_{(i=1,\cdots,n)}$, $plan(a_2, g) = [x_j^2]_{(j=1,\cdots,m)}$.

Agent a_1 's choice of operators from its initial individual plan is done through the following two steps:

1. Committing operators that support the hypotheses of the other agents
2. Identifying interchangeable subplans (defined below) and making choices (committing)

Definition 20 (Interchangeable Plans) Two plans $[x_i]_{(i=1,\cdots,n)}(w^t \Longrightarrow w^{t'})$ and $[y_i]_{(i=1,\cdots,m)}(w^{t'} \Longrightarrow w^{t''})$ are said to be interchangeable, if and only if $w^{t'} \equiv w^{t''}$ and $pre(x_1) \equiv pre(y_1)$.

Let INT denote a binary predicate over operators that is true if and only if its arguments are interchangeable plans.

(**Example**) For the following collaborative plans given in the previous example,

$$plan(a_1, g_{ex}) = [\mathsf{pickup}(a_1, b), \mathsf{move}(a_1, b, r_1, r_2),$$
$$open(door_{23}), \mathsf{move}(a_1, b, r_2, r_3)]$$
$$plan(a_2, g_{ex}) = [\mathsf{open}(a_2, door_{23}), \mathsf{move}(a_2, nil, r_2, r_1),$$
$$onfloor(b, r_1), \mathsf{pickup}(a_2, b),$$
$$\mathsf{move}(a_2, b, r_1, r_2), \mathsf{move}(a_2, b, r_2, r_3)],$$

interchangeable subplans are $[\mathsf{pickup}(a_1, b), \mathsf{move}(a_1, b, r_1, r_2), \mathsf{move}(a_1, b, r_2, r_3)]$ and $[\mathsf{pickup}(a_2, b), \mathsf{move}(a_2, b, r_1, r_2), \mathsf{move}(a_2, b, r_2, r_3)]$.

5.3.1 Committing Operators

Let $\mathcal{ABD}$ denote a unary predicate over operators that is true if and only if its argument is an abduction.

Operators in plan $[x_i^1]_{(i=1,\cdots,n)}$ which support a_2's hypotheses are committed according to the following procedure.

```
For all $x_j^2 \in [x_j^2]_{(j=1,\cdots,m)}$
        if $\mathcal{ABD}(x_j^2)$
        then do commit$(a_1, x_i^1)$,
                s.t. $\neg\mathcal{ABD}(x_i^1) \wedge x_j^2 \in eff(x_i^1)$
```

(**Example**) For $plan(a_1, g_{ex})$ and $plan(a_2, g_{ex})$, since the effect of operator $\mathsf{open}(a_2, door_{23})$ in $plan(a_2, g_{ex})$ includes hypothesis $open(door_{23})$ in $plan(a_1, g_{ex})$, agent a_2 commits itself to execute operator $open(door_{23})$.

5.3.2 Choice of Subplans

Let $\mathcal{CMT}$ denote a unary predicate over operators that is true if and only if its argument contains committed operators.

Agent a_1's choice of subplans from its initial individual plan is done according to the following procedure:

1. From the last operators in plans $[x_i^1]_{(i=1,\cdots,n)}$ and $[x_j^2]_{(j=1,\cdots,m)}$, identify all the subplans $y_k^1 (\subset [x_i]_{(i=1,\cdots,n)})$ and $y_k^2 (\subset [x_j^2]_{(j=1,\cdots,m)})$ which satisfy $\mathcal{INT}(y_k^1, y_k^2)$. The identified subplans of one agent are not allowed to have shared operators. (We assume that the number of these interchangeable operators is N)
2. For all $y_k^1 (k = 1, \cdots, N)$, evaluate the utility equalization condition defined below. If the condition is satisfied, commit all the operators in subplan y_k^1. Otherwise, apply the *hypothesize* operation to all operators in y_k^1 that are not committed ($\mathcal{CMT}(y_k^1)$ is not true). The hypothesized operators are supported by the subplan of the other collaborating agent, y_k^2. If $\mathcal{CMT}(y_k^1)$ is true, then all operators preceding the last committed operator are *committed*. The remaining operators are *hypothesized*.
3. Agent a_1 chooses all the committed operators in $[x_i]_{(i=1,\cdots,n)}$

[Utility Equalization Condition]

$$\begin{gathered}
|\,(worth(a_1, g) - cost(y_k^1) - cost([x_i^1 \mid \mathcal{CMT}(x_i^1)]))\\
-(worth(a_2, g) - cost([x_j^2 \mid \mathcal{CMT}(x_j^2)]))\,|\\
\leq\\
|\,(worth(a_1, g) - cost([x_i^1 \mid \mathcal{CMT}(x_i^1)]))\\
-(worth(a_2, g) - cost(y_k^2) - cost([x_j^2 \mid \mathcal{CMT}(x_j^2)]))\,|
\end{gathered}$$

where if an committed operator is included in y_k^1(or y_k^2), the cost of the operator is omitted in the calculation of the cost of y_k^1(or y_k^2).

(**Example**) Assume that $worth(a_1, g_{ex}) = 6$ and $worth(a_2, g_{ex}) = 6$. For $plan(a_1, g_{ex})$ and $plan(a_2, g_{ex})$, the first interchangeable plans are [$\mathsf{pickup}(a_1, b)$, $\mathsf{move}(a_1, b, r_1, r_2)$, $\mathsf{move}(a_1, b, r_2, r_3)$] and [$\mathsf{pickup}(a_2, b)$, $\mathsf{move}(a_2, b, r_1, r_2)$, $\mathsf{move}(a_2, b, r_2, r_3)$]. Since agent a_2 has already committed itself to operator $\mathsf{open}(a_2, door_{23})$ in $plan(a_2, g_{ex})$, the current utility of agent a_2 is equal to 5 (= 6-1). Meanwhile, the current utility of agent a_1 is equal to 6, since agent a_1 has not committed itself to any operator. Therefore, agent a_1 commits itself to execute subplan [$\mathsf{pickup}(a_1, b)$, $\mathsf{move}(a_1, b, r_1, r_2)$, $\mathsf{move}(a_1, b, r_2, r_3)$] by means of evaluating the utility equalization condition. On the other hand, agent a_2 hypothesizes operators $\mathsf{pickup}(a_2, b)$,$\mathsf{move}(a_2, b, r_1, r_2)$, and $\mathsf{move}(a_2, b, r_2, r_3)$. No other interchangeable subplans remain, and agent a_2 hypothesizes operator $\mathsf{move}(a_2, nil, r_2, r_1)$. As a result, the following two plans, $plan(a_1, g_{ex})'$ and $plan(a_2, g_{ex})'$, which are sequences of committed operators, are obtained.

$$plan(a_1, g_{ex})' = [\mathsf{pickup}(a_1, b), \mathsf{move}(a_1, b, r_1, r_2), \mathsf{move}(a_1, b, r_2, r_3)]$$
$$plan(a_2, g_{ex})' = [\mathsf{open}(a_2, door_{23})]$$

With an appropriate temporal ordering (schedule) c_{12}, these two plans form a complete collaborative plan.

5.3.3 Condition for Incomplete Collaborative plan

All the committed operators with a partial order in the collaborative award form the collaborative plan for goal g. The collaborative plan is represented in the following form.

$$([x_i^1 \mid \mathcal{CMT}(x_i^1)](i = 1, \cdots, n) : [x_j^2 \mid \mathcal{CMT}(x_j^2)](j = 1, \cdots, m)) \mid$$

If either $[x_i^1 \mid \mathcal{CMT}(x_i^1)]_{(i=1,\cdots,n)}$ or $[x_j^2 \mid \mathcal{CMT}(x_j^2)]_{(j=1,\cdots,m)}$ contains any hypotheses that are not supported by operators in the collaborative plan, the resulting collaborative plan is incomplete. The plan can be made complete by forming another collaborative group with agents who have the skills necessary to achieve the hypotheses. This is accomplished by reapplying the collaborative plan scheme, using the hypotheses as goals.

6 Concluding Remarks and Future Work

We have presented a formal model for generating collaborative plans from, possibley incomplete, individual plans in multiagent domains. Also, we have developed a utility-based model of rational choice with which agents can rationally decide which actions will be performed by each agent in a collaborative plan. Given a goal, a rational agent generates the best plan with respect to its utility, which is calculated by subtracting the cost of the plan from the worth of the goal for that agent. The choice of activities in collaboration is guided by two criteria; (1) the maximum completeness (obstacle detection and elimination), and (2) the utility equalization principle. If the resulting collaborative plan still contains an unsupported hypothesis, the plan is incomplete. The plan can be made complete by forming another collaborative group with agents who have the skills necessary to achieve the hypotheses, using the hypotheses as goals.

We are currently working on the following extensions: (1) Implementing the proposed scheme; (2) Theoretical analysis on computational complexity of the collaboration scheme; (3) Incorporating a learning capability into agents, so that successful collaboration can be reutilized again without the overhead of organizing a group.

References

[Corkill 79] Daniel D. Gorki 11. Hierarchical Planning in a Distributed Environment. In Proceedings of the Sixth International Joint Conference on Artificial Intelligence (IJCAI-79), pp. 168 175, 1979.

[Davis and Smith 83] Randall Davis and Reid G. Smith. Negotiation as a Metaphor for Distributed Problem Solving. Artificial Intelligence, Vol. 20, pp.63 109, 1983.

[Ephrati and Rosenschein 92] Eithan Ephrati and Jeffrey S. Rosenschein. Constrained Intelligent Action: Planning Under the Influence of Master Agent. In Proceedings of the Tenth National Conference on Artificial Intelligence (AAAI-92), pp. 263 268, 1992.

[Georgeff 83] Michael P. Georgeff. Communication and Interaction in Multi-Agent Planning. In Proceedings of the. Third National Conference on Artificial Intelligence (AAAI-83). pp. 125 129, 1983.

[Kinny et al. 92] David Kinny, Magnus Ljungberg, Anad Rao, Elizabeth Sonenberg, Gil Tidhar, and Eric Werner. Planned Team Activity. In Proceedings of the Fourth European Workshop on Modeling Autonomous Agents in a Multi-Agent World, pp. 1 20, 1992.

[Martial 90] Frank von Martial. Coordination of Plans in Multiagent Worlds by Taking Advantage of the Favor Relation. In Proceedings of the Tenth International Workshop on Distributed Artificial Intelligence, 1990.

[Nagao 93] Katashi Nagao. Abduction and Dynamic Preference in Plan-Based Dialogue Understanding. In Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence (IJCAI-93), 1993.

[OsaWa and Tokoro 92] Ei-Ichi Osawa and Mario Tokoro. Collaborative Plan Construction for Multiagent Mutual Planning. In Eric Werner and Yves Demazeau. editors, DECENTRALIZED A.I. 3. Elsevier/North Holland, 1992. (Proceedings of the 3rd European Workshop on Modeling Autonomous Agents in a Multi-Agent. World).

[Reiter 80] Raymond Reiter. A Logic for Default Reasoning. Artificial Intelligence, Vol. 13, pp.81 132, 1980.

[Zlotkin and Rosenschein 89] Gilad Zlotkin and Jeffrey S. Rosenschein. Negotiation and Task Sharing in a Non-Cooperative Domain. In Proceedings of the Ninth Workshop on Distributed Artificial Intelliqence, pp. 307 327, 1989.

[Zlotkin and Rosenschein 90] Gilad Zlotkin and Jeffrey S. Rosenschein. Conflict Resolution in Non-Cooperative Domain. In Proceedings of the Eighth National Conference on Artificial Intelligence (AAAI-90), pp. 100 105, 1990.