Anticipatory coordination through action observation and behavior adaptation

Michele Piunti[1], Cristiano Castelfranchi and Rino Falcone[2]

Abstract. To establish coordination at a cognitive level, we need to come through to some of the strict assumptions of the traditional deliberative agents. We also need to cope with some of the real world boundaries, where, for instance, knowledge and perception are affected by uncertainty and message exchange as most direct form of subjective coordination may not be reliable everywhere. Intentional, subjective aspects of coordination may concern prediction mechanisms (i.e. future state representations, mind reading), but also true cognitive expectations that agents should exploit to reconsider their intentions, and their use in reading other agents goals (for better achieving their own). On the basis of the cognitive theories of social actions and behavioral implicit communication, we here propose an observation based approach providing agents with explicit anticipatory coordination capabilities in order to exploit signs coming from other agents and, contextually, adapt behavior in anticipatory terms. Pro-activeness, adaptiveness, opportunism come out from the means-end reasoning of individual actors: agents embedding such adaptive skills are leaning to make coordination as an emergent property of their interactions.

1 INTRODUCTION

The research behind this work wishes to provide artificial agents engaged in real world applications with anticipatory coordination abilities. We here refer to subjective approaches to coordination[3], meaning those aspects of the activity of an agent specifically devoted to deal with a dynamic environment and its social interferences. In these contexts, agents continuously cope with opportunities to exploit and threats/obstacles to avoid. To coordinate herself with a give event or act, an agent has to perceive or foresee it thanks to some perceptual hints, index or sign. In real world applications most direct form of subjective coordination through message exchange is not universally serviceable. Sometimes agents may not desire to exchange information (i.e. hostile agents), otherwise also cooperative agents may be reluctant to send explicit messages due to heterogeneous models and technologies, environment and resource constraints. Direct messaging further introduce limitations and costs, namely weight for additional equipments and transmitters, bound of communication range, unreliability of services, need for standardized protocols.

On the contrary, we argue that coordination between agents is not necessarily based on explicit communication. An action performed by any one agent potentially updates the perception (and the epistemic states) of other agents, thus observing and interpreting the world where agents are pursuing their goals is an intrinsic opportunity for coordination activities. Beliefs about other's mental states are also a result of the process of interpretation of other's behavior, that can be considered as the observable *sign* for his internal state[7]. We guess one of the main functions of observation in agent living in a common world populated by other agents is coordination, while one of the main form of coordination is observation-based. Indeed, just behavior without any modification or any additional signal or mark can be exploited as a premonitory sign, thus recognition capabilities make possible for an observer to predict future actions of an observed agents. By so doing, recognizer agents should exploit these capabilities to conceive an explicit form of expectation.

In order to enhance coordination for social tasks, several coordination techniques have been developed, including those based on social conventions and norms [26], decision and game theoretical strategies [13, 14], stigmergy infrastructures [1]. Several techniques for goal and plan recognition have been proposed and applied to different application domains. The idea to exploit observation to acquire coordination hints is not new in literature [17, 12]. Less effort has been given to goal directed behavior adaptation on the basis of the anticipated outcomes of the interactions. The inferential knowledge carried out by intended plan recognition mechanisms makes possible to ascribe mental states (goals) to others: in so doing, agents are enabled to anticipate actions performed by others, thus to reconsider their intentions and/or exploit those actions as an enlarged, exogenous repertoire of actions at disposal [10]. To establish coordination at an intentional level, the interferences between agent activities have to be endogenously valued as *positive* (A2's actions realize A1's goals or create opportunities) or *negative* (A2's actions create obstacles to A1 or thwart A1's goals). To do this, a recognizer agent has to subjectively classify the expected external outcomes: valuing these expectations as positive (or negative) is made according to their contribute (or determent) to the recognizer ongoing purposes and mental states (e.g. Goal, Beliefs).

In this work we propose a design model for cognitive agents endowed with the ability to *predict the outcomes* of other agents actions, to *build a model* of future events and to *react in advance* according to these expected events. We do approach the problem by enabling agents with *mind-reading* abilities: in so doing we introduce an enriched perception module used for observing and recognizing signs, actions, and practical behavior. We further provide agents with a mean for intending other agents in form of mental states. As we show in the next sections, the subjective utilities for the expected

[1] Institute of Cognitive Sciences and Technologies - I.S.T.C. - C.n.r. and DEIS, Università di Bologna - Alma Mater Studiorum, Italy, email: michele.piunti@istc.cnr.it

[2] Institute of Cognitive Sciences and Technologies - I.S.T.C. - C.n.r., Italy, email: {cristiano.castelfranchi, r.falcone}@istc.cnr.it

[3] In the context of MAS, it rely on the viewpoint of the individual agent that can perceive and understand the actions of its peers. [24] defined subjective and objective coordination respectively as an endogenous, psychological capability for coordinating agent vs. exogenous, infrastructural system to coordinate agents.

outcomes produce a fully represented expectation that can be used to rationally change behavior, and to cause avoidance or exploitation of alternative courses of action.

The rest of this paper is organized as follows: in section 2 we define the bounds of traditional cognitive architecture in dynamic, social scenarios, showing how anticipatory competencies may overcome a set of restrictive assumptions; in section 3 we describe the architecture for anticipatory agents able to exploit observation and plan recognition as building block for anticipatory behavioral coordination; in section 4 we describe a case study through an experiment; in section 5 we conclude with final discussions.

2 FROM GOAL-DIRECTED, DELIBERATIVE AGENTS TO ANTICIPATORY, INTERACTIVE AGENTS

We refer to true cognitive, *goal-governed* and deliberative systems, able to manage subjective, internal representations for beliefs and goals[4]. Traditionally deliberative agents operate to reach a desired state of affairs from the current state by chaining their environment through a given set of actions and plan operators[5]. The wide adopted BDI-based architectures [23] focalize systems in deliberation among set of goals and means-end analysis[6] between alternative courses of actions, but let intention making and execution of plans in a functional, even purely reactive form. In this sense, deliberative agents process their information reacting in a procedural way: they choose in repertoire the plan to execute according to filtering of conditions (belief formulae, utility functions, priorities etc.), whilst the available plan library is handcrafted at design time.

Early implementations of BDI-like systems operate according to the following restrictive assumptions [22]:

1. **Static world assumption**: the world is not changing during the reasoning process.
2. **Infinite resource assumption**: even if the world is changing the agent has sufficient resources to appraise all the relevant changes and consequently revise the belief base. The agent can also plan faster than the rate at which the world is changing leaving the plans still relevant.
3. **Complete knowledge assumption**: the agent has the capacity to perceive the complete state of itself and its environment and the information describing the environment results consistent and without noise at each point in time.
4. **Determinism assumption**: each planned action will completely realize the expected outcome.
5. **Single agent assumption**: actions performed by other entities do not influence agent activities. There are no other agents to aid or thwart agent plans.

In addition, coordination competencies are generally based on the reaction upon a direct perception of some events or act (*reactive coordination*) and often treated along with the general problem of intention reconsideration [18, 25].

[4] This classification makes sense against the category of merely *goal-oriented*, functional systems, without any internal anticipatory representation for the goal of the action, where the teleonomic character of the behavior is in its adaptive function (e.g. managed by some learning algorithms). This class of systems does define no native support for dealing with the future through representations of future states.

[5] That state also indicates the 'goal state', more precisely the representation of the goal indicating the satisfaction of a subjective desire in a future state.

[6] Deliberation is the process by which agent select the goal to be pursued; means-end is responsible to compose plans (the means) in order to achieve the previously adopted goal (the end).

2.1 Breaking assumptions through anticipation

Multidisciplinary convergencies indicate agents with anticipatory capabilities be more effective to overcome a larger set of real world requirements. We define *anticipation* as the ability to coordinate the behavior with the future: more formally, anticipation enable agent to *react in advance* (at an instant t) to an event (or to a world state) that will be realized at $t + t'$. Practical anticipatory behavior should be exploited on the basis of the knowledge of the current situation *but also* on some form of expectation about future states and events. Given this, the behavior does not only depend on past and present, but also on some knowledge about the future: [3] introduced anticipatory agents entertaining *expectations* as mental representations of the future. Expectations enable agents to be anticipatory just by working on them, for virtually exploring alternatives, opportunities, events, results. Expectations are not simply predictions neither belief on the future: they are given as axiological anticipatory mental representations, also endowed with *valence* against some concern, drive, goal of the agent. As in [20], we point out that, in a cognitive system, expectations play several important roles: i) precede and control the execution of actions. ii) bias sensory processing (attention, active perception) and resource allocation. iii) are used to bias goal selection and intention reconsideration.

Furthermore, in Multi Agent Systems agents play in a shared environment and have to operate in a world eliciting interferences, where the action of an Agent $A2$ could affect the goal of another Agent $A1$. We guess that modeling mental states of individual agents allows interaction with the counterpart in the minds of other agents. Our challenge in this work is to enable the expectations about A2 actions to be used by A1 as a sign, an help in deciding to react in advance, anticipating and exploiting events and outcomes performed by the other. By so doing, we design agents able to interact following an *anticipatory coordination*, based on the anticipation of interferences, opportunities and dangers.

2.2 Interaction for goal directed agents

Interaction between agents may result at a certain grade of cooperativeness, competitiveness or in some grade in between: it may result positive or negative for agents that are helped or damaged, favored or threatened by the (effects of) the actions of the others[7]. In the cooperative case, agents are more inclined to behave pursuing joint goals: on the one side they intend to exploit actions performed by others for their purposes, on the other hand they want to help each other to achieve common goals. On the contrary, in the competitive interactions, agents intend to thwart the others: on the one side they show avoidance of undesired outcomes, on the other side they perform hostile behaviors in order to prevent adversarial threatful purposes.

[7] noticed a deeper form of interaction in attempting to influence the behavior of the others by changing their mental states. In observable environments actions acquire a communicative function by preserving their practical end through their long term effects and modification in world states. By considering each action with its necessary world contexts in terms of preconditions and outcomes, A1 may induce A2 to abort her behavior by giving misleading signs or removing the necessary conditions, or may persuade A2 to do something by intentionally signalling opportunities or creating the

[7] Notice that these notions can meaningful be applied only to systems endowed with some form of goal, where the effects of the action of an agent are relevant and impact on the goals of another.

necessary pre-conditions for A2's actions. To this end, A2 can intentionally change A1's mind through implicit communication via stigmergic traces, long term physical outcomes, environment modifications. Hence, A1 may not coordinate only by reading A2's mind (i.e. perceiving her behavior during its performance) but can exploit other post-hoc traces and outcomes of it in observable changes of the environment [27].

From the viewpoint of A1 interfering with A2, there are two strategies:

1. To adapt hers own behavior to A2's behavior, in order to exploit positive interferences (or to avoid negative ones);
2. To attempt to change A2's behavior by inducing A2 to do what she needs or to abort activities damaging A1.

Tab.1 distinguishes four different alternatives for anticipatory coordination. The first rows shows the cases of behavior adaptation: in cooperative (positive) coordination, A1 changes her (practical, purposive) plan in order to profit by a favourable circumstance; in competitive (negative) coordination, A1 is aimed at avoiding a threat. The

	Competitive (Negative Interference)	Cooperative (Positive Interference)
Adapting behavior	avoid adversary activities	exploit teamwork activities
Changing other's mind	misleading signs, stigmergic traces	collaborative signs, stigmergic traces

Table 1. Anticipatory coordination holds to different effects according to the type of interaction between the involved agents.

second row shows the cases of direct influence by changing mental states of the other: A1 may induce A2 to abandon her threatening goal in order to avoid some risky effect or may persuade A2 to pursue some action in order to obtain its profitable outcomes.

2.3 Behavior as 'sign' for Anticipatory Coordination

Behavioral Implicit Communication theory [4, 5] introduces practical behavior as an important form of contextual communication between agents, without explicit messaging, neither direct speech acts. In strong BIC, agents (sources) behave intentionally with the additional motivation to make others (addressees) understand their purposes, i.e. to capture some meaning from implicit messages and, consequently, change their minds.

As for the adaptive strategy, we here refer to a weakest awareness between agents: on the one side they know to be monitored by others but do not ascribe an additional motivation in doing actions also for being recognized; on the other side, they have the goal/plan of interpreting observed behaviors, to coordinate with them and anticipate events. We present a computational model for coordinating with other predicted behavior, thus ignoring, for the moment, the possibility to induce changes in others behavior. The first layer of our design model requires the observer to *perceive (or infer) interferences*. This can be made through general plan recognition techniques:

- As in most plan recognition assumptions, agents refer to an internal knowledge and continuously match perceptual hints with it, in order to recognize other agent actions and plans.
- Through plan recognition mechanisms, agents attain *signification* (namely the semiotic ability to "ascribe sense" to the observed behaviors) and infer expectations on actions and world changes performed *by others*.

The second layer requires to *adapt behavior in anticipatory terms*, by avoiding threats or exploiting opportunities. Agent changes her own plan (sub-goal) and produces a new plan which is based on her beliefs (predictions) about the goal of another. To do this, she uses a further model for evaluating expectations and reconsider intentions:

- Agents *evaluate* positive and negative circumstances. Evaluating enable agents to read the world (i.e. actions performed by others) in terms of *positive and negative expected outcomes*.
- Agents reconsider their intentions and mental states on the basis of the new (valued) expectations.

By so doing, expectations become true representations of the future, upon which agents may concern, deliberate, reason and reconsider their plans, thus coordinating their behavior with the not *yet* existent.

Adapting behavior by working on future states elicits two main kind of appraisal. In the positive case, the agent anticipates an unexpected *help*: she can remove from the planned workflow the action that will be executed by others (agent A1 exploits A2's action, intentionally delegates and relies on it [11]). In the negative case, agent A1 anticipates an unexpected *determent*: to economize resources, she has to reconsider the ongoing intention, aborting the current action and adopting an alternative one (if present). Notice that the use of plan recognition methods introduces uncertainty in the reasoning process (coming from incomplete knowledge and errors in observation evidences, risk evaluation, learning processes etc.).

3 AGENTS AND PLAN RECOGNITION MODELS

In the following sections we present the architecture, including the plan recognition module, and we describe the reasoning process for the anticipatory coordination.

3.1 Design

As for the agents kernel we adopted the Jadex engine [2], a multi-threaded BDI framework leading to loosely coupled Beliefs, Goal and Plans representation, including their mutual relations. Jadex deliberation is driven by the evaluation of logic formulae (put in form of Belief formulae) and arcs of inhibition between goals to dynamically resolve their priority. The sensor component directly gets data from the environment simulator: when an entity is sensed, its symbolic description is provided by a preceptor filter and then is used for belief revision (Fig.2). For simplicity, we assume that visual information retrieved from the environment simulator and symbolic information handled by sensor are given at the same level of representation.

A Mental States component is used to manage working memory, to allocate configuration of epistemic resources and to express attitudes and bias towards the actual state of affairs. We define Mental States through a set of related behavioral and mental changes increasing agents opportunism and proactiveness towards the environment changes. By using a functional approach, we have further defined some important roles that these affective states play for anticipation (for more details see [20]).

3.2 Plan Recognition

For the plan recognition mechanism, we assume a shared symbolic representation of purposive actions through hierarchical plans and we introduce a background process in perception filtering module (Fig. 1). Plan representation, used to match perceptions, assumes the

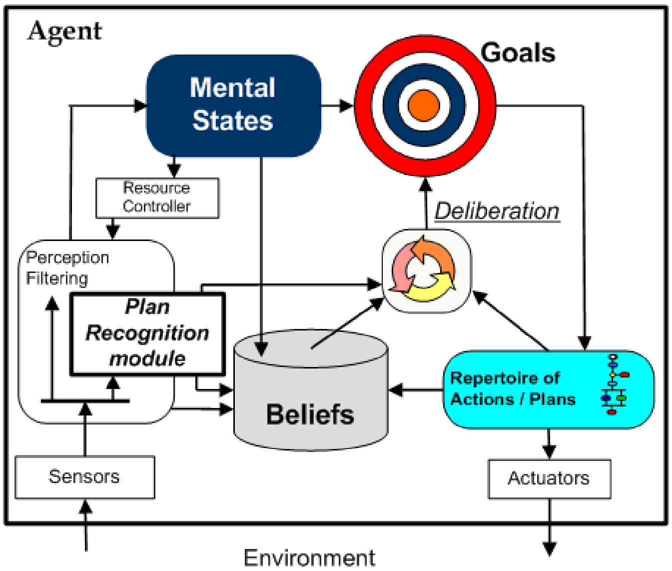

Figure 1. Agent architecture includes a BDI core.

role of the subjective belief about the preferences and the practical behavior usually performed by others.

It has been argued that plan recognition problem can be treated as the general problem of abduction. In this sense, an observer makes hypothesis following a diagnostic approach (observe action "A" to deduce a goal "G"). Thus, for a given set of actions, observer matches perception with the internal representation of plans to discover its best explanation. Perception filtering (Fig. 1) uses a Probabilistic Horn Abduction meta-interpreter [21] implemented within a tuProlog engine [9]. Once the prior probabilities are given, PHA calculates the list of the 'best explanations' ordered by their crescent likelihood. The output of the recognizer is in domain of probability: the execution model manages the confirmation (failure) of the observed actions and provide the confirmation (failure) of the hypothesis. This directly results in reinforcing (inhibiting) the probability of that hypothesis (for more details on these techniques, see [8, 15, 12]).

3.3 Anticipatory process

Representation of plans is given in a single root, directed, acyclic connected graph, where roots indicate top level goals and leafs denote self-contained actions (plan steps). At any given time, the observed agent is assumed to be executing a plan decomposition path, from root to leafs, through plan tree branches (Fig. 2). We introduce:

- A repertoire of actions $\mathcal{A} = \{a_1, a_2, ..., a_A\}$ that constitute agents practical behavior as shared knowledge between observer agents. We assume plan representation used for recognition fully consistent to the plan library of operators used for practical behavior.
- A set of world states $\mathcal{S} = \{s_1, s_2, ..., s_S\}$ can be used to evaluate local background and world contexts. Notice that first order logic formulae upon $\mathcal{S}$ also constitute the preconditions for the execution of actions and for the activation of goal and plans.
- A set of outcomes $\mathcal{O}_a = \{o_{a_1}, o_{a_2}, ..., o_{a_A}\}$ indicating the world state as it is assumed to be *after* the execution an action. Notice that $\mathcal{O}_a \subset \mathcal{S}$, where each outcome indicates the expected effects of the related action in repertoire.
- A prediction function $\pi : \mathcal{A}^n \times \mathcal{S} \times \mathcal{T} \to \mathcal{A} \times \mathcal{O}_a \times \mathcal{P}$ that expresses observer's hypotheses that, given at an instant t, n observed evidences for actions in $\mathcal{A}^n$, with the world context in $\mathcal{S}$, a certain action will be performed by an observed agent at $t + t'$, with the respective outcome in $\mathcal{O}_a$ and a probability in the distribution $\mathcal{P}$.

For an observer agent, plan recognition process provide the prediction of the next action performed by an observed agent. To this end, it refers to two sources of information: we do assume for each performed action in $\mathcal{A}$ an associate tuple of conditions on its observable features; observer agents further relate these features to some clarifying contextual world states in $\mathcal{S}$ (as noticed in [15], the use of world states significatively helps to disambiguate situations and reduce the overall complexity of the process). Given prior probabilities on plan branching, as they are reported in plan representation as meta-belief, π introduces a grade of (un)certainty in observer's prediction. Hence, when considering what goal the observed agent might be pursuing, PHA meta-interpreter provides the best (most likely) explanation in terms of recognized goals, also evaluating the world state (in $\mathcal{S}$). When allowed by world constraints and observability, agent's perception filtering *observes actions* performed by others and relate it to the *world context*, translating the data stream from sensors in symbols simultaneously referring to the prior knowledge of plans. By matching perception with the internal representation, a PHA-based mechanism provide *concurrent hypothesis*: observation process persists until the set of evidences in $\mathcal{A}^n$ become sufficient to disambiguate the corresponding goal: once the best explanation overcomes a fixed threshold, observer agent shapes an *expectation*, by balancing the observed predicted goal with own purposes. By so doing, observer *appraises and gives a subjective value* to the expectation: in positive terms, if the expectation is due to positive interferences (i.e. helps the pursuing of her goal); in negative terms, if the expectation is due to negative interferences (i.e. agent anticipates threats, obstacles).

In the second phase, observer agent adapts the behavior by reconsidering her intentions. We assume:

- A repertoire of *counteractions* $\mathcal{CA} = \{ca_1, ca_2, ..., ca_C\}$ that can be related to the observer goals and carried out to react to the prediction given by π.
- A set of outcomes $\mathcal{O}_{ca} = \{o_{ca_1}, o_{ca_2}, ..., o_{ca_C}\}$ indicating the expected effects for each counteraction.
- An *outcome function* $\varphi : \mathcal{CA} \times \mathcal{S} \times \mathcal{T} \to \mathcal{O}_{ca} \times \mathcal{P}$ that returns the probability for realizing the outcome of the counteraction ca_i (performed instead of a_j) when the actual world's state is in $\mathcal{S}$. It models the uncertainty and the confidence of the observer in deciding which counteraction to take respect to the determinism of its outcome.
- An *utility function* $\upsilon : \mathcal{O}_{ca} \to \mathcal{U}$ giving the *utility value* of a certain outcome as an heuristic composition of subjective importance and desirability of the outcome, thus it is strictly related to the ongoing goal of the observer. For the observer agent, utility measures the desirability of any given outcome. Its value can be related to different domains (i.e. game-theoretic, normative) and coupled with different measures as perception of risk, urgency etc.

More formally, let h_j be a tuple $\langle a_j, s_j, t_j \rangle$: given the above definitions, $\pi(h_j)$ is the probability (provided by the plan recognition module) of a certain hypothesis j, $\varphi(ca_i, s_j, t_j)$ the confidence on the expected outcome for the counteraction ca_i (the probability that the counteraction will have its intended outcome), and $\upsilon(ca_i)$ the expected utility (given in decision theoretic account, in case of success of the respective counteraction). Agents select the reconsidered action to take by comparing, for each counteraction ca_i, the following expression:

$$\pi(h_j) \times \varphi(ca_i, s_j, t_j) \times \upsilon(ca_i) \tag{1}$$

The above expression anticipate the effects and the subjective utility of a counteraction to take, given the anticipated effect of an action performed by the other. By so doing we do introduce subjective expectations in terms of agent's native epistemic states (Beliefs) and motivational states (Goals) (Fig. 1).

In other terms, agents adapt their plans by managing an *expected degree of adequacy* for counteractions in repertoire. Its value is a composition of an *epistemic state* (an uncertain, graded belief) and a motivational state (a graded utility and a subjective importance for the counteraction ca_i realizing a certain goal). By selecting the proper counteraction, agents may take advantage of the anticipated events, enhancing opportunism and pro-activeness, or decide to avoid them, by abandoning their activities and saving resources for alternative pursuable goals.

4 EXPERIMENT

In order to test different architectural solutions so that different strategies can be significantly compared, we engaged agents in a foraging task $T\langle LOI_n, V_n, R_L, \mathcal{A}, \mathcal{S}, r, S_r, s, \mathcal{D}\rangle$ in a 2D environment. The scenario presents a set of n Location Of Interest (LOIs) and requires a group of agents $\mathcal{A}$, each with adaptive sensor range r, sensor rate S_r, speed s, to find n types of valuables V_n, pick up them (one at a time) and bring back to the repository location R_L. Sentry agents $\mathcal{S}$ have the goal to guard LOIs and hinder agent foraging. Each valuable type V_X is coupled with a respective location LOI_X.

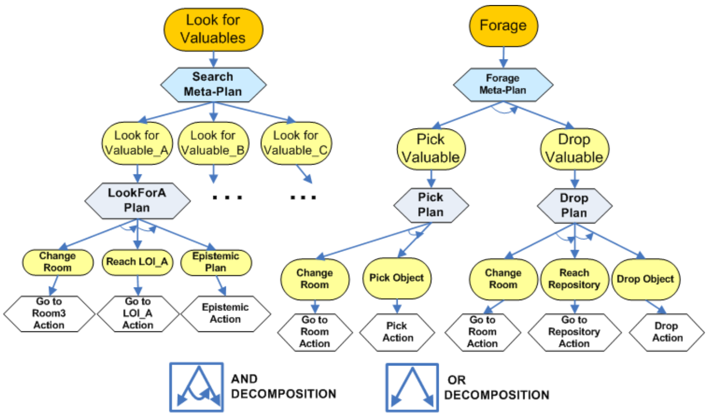

Figure 2. Plan representation used to match observations: it shows practical behavior and purposive actions for agents engaged in the foraging task.

Valuables are dynamically generated by the environment simulator close to the respective LOI, according to a probability distribution $\mathcal{D}$. Environment also present a layout of walls and doors creating room, corridors and pathways. Agents do not have an a priori knowledge of the distributions and use a library of paths and plans to move between locations.

Fig. 2 shows representation of hierarchical plans for foraging agents purposive behavior. Top level nodes (*Look for Valuables* and *Forage*) are expanded into sequences of lower level nodes, each of which is further expanded into yet lower level nodes. Thus, single plans are not just a sequence of basic actions, but may also dispatch sub-goals. The leaves of the plan structure form a non-hierarchical plan of practical actions that agents execute and observe themselves. Notice that, according to world constraints (i.e. wall, obstacles, sentries), goal/plan hierarchies may result with interleaved sequences of leaves and generate interleaved sequences of actions.

In our experiments we use three kind of valuables and three associated LOI ($n = 3$). Sensor component directly gets data from the environment simulator: when an entity is sensed, its symbolic description is provided by the simulator and the preceptor module filters it for belief revision and further reasoning processes. For simplicity, we assume that both perception data and symbolic information handled by the filter are given at the same level of representation. Foraging agent's plan knowledge (used by recognizer) is built upon internal Prolog representation. Intention reconsideration and re-planning processes are triggered by the activation of an hypothesis, namely when $\pi(h_j)$ overcomes the corresponding threshold: the process of valuing is managed at meta-level reasoning, with a meta-plan, by which the observer evaluates *on-line* the various available options.

Tab. 2 shows, from the point of view of the single agent A1 and for each sequence of observed actions and world contexts, the set of options in repertoire. Each option is a counteraction and encapsulates the respective confidence of success (due to indeterminism) and the subjective expected utility (in case of success).

In the second row of the table we show the case when agent A1 receives the evidence that agent A2 is transporting a valuable Obj_x, while the context is that A1 is looking for the same Obj_x. In this case, the time further devoted by both agents in looking for the same valuable would be wasted but A1 provide an explicit counteraction to save resources and optimize global behavior. An internal signal (from A1 perception filtering) indicates $\pi(h_j)$ is overcoming the fixed threshold: it triggers the meta-level reasoning process where confidences φ and utilities υ of counteractions in repertoire are evaluated. From a cooperative perspective, A2 not only has the goal to transport the valuable, but also the goal to make A1 aware of something: although she is not sending an explicit message, he has the goal of changing A1's mental states, updating her beliefs in order to modify behavior. In this case, the first counteraction to drop the ongoing search is taken because of its optimal expectation, hence A1 will spend her resources to look for a different kind of valuable, namely Obj_y near LOI_y (see Fig. 3).

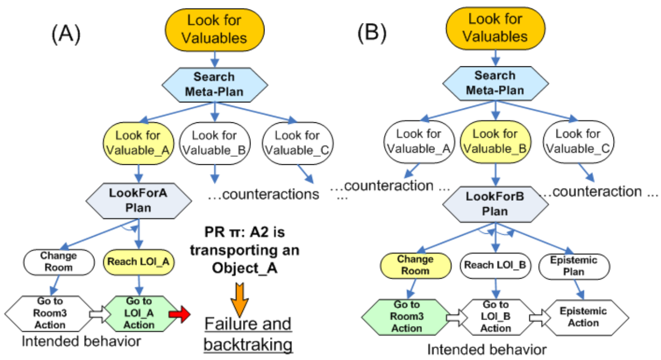

Figure 3. Adaptive Behavior: when A1 recognize A2, she receives an internal event (from Plan Recognition module), and breaks the current plan (A). The selection of the alternative course of actions (B) is driven by the evaluation of expectations for each counteractions in repertoire.

Along experiments, default values for $\varphi(ca_i)$ and $\upsilon(ca_i)$ are given in fuzzy terms, from ZERO [0.0] indicating absence of confidence

Observed action	Context World state	Options (ca_i) Repertoire of counteractions	Confidence of counteractions (φ)	Utility of the outcome (υ)
A2 approaching Obj	A2 closer to Obj	Abort Speed up Persist	MAXIMUM [1.0] LOW (0.3) MAXIMUM [1.0]	LOW [0.3] HIGH (0.8) ZERO [0.0]
A2 transporting Obj_x	A1 look for Obj_x	(Drop search Obj_x and search Obj_y) Persist	MAXIMUM [1.0] MAXIMUM [1.0]	MAXIMUM [1.0] ZERO [0.0]
A1 and A2 approaching the same Obj	Sentry close to A1	Distract the Sentry Abort Persist	LOW [0.3] MAXIMUM [1.0] ZERO [0.0]	MAXIMUM [1.0] LOW [0.3] ZERO [0.0]

Table 2. Intention Reconsideration through on-line evaluation of hypothesis and counteraction selection (agent A1 observes and anticipates agent A2). Confidences and utilities are given in fuzzy terms between ZERO [0.0] and MAXIMUM [1.0].

and utilities to MAXIMUM [1.0] indicating full utility and confidence value. Notice that agents evaluate also the hypothesis to remain committed and persist without adopting new intentions.

Belief thresholds strongly affect agent performances with space, time and activities trade offs. As in [17], the use of a belief-net may introduce learning mechanisms to adjust thresholds, confidences and utilities during the task. Thus, agents that become aware to act in particular environments (i.e. more or less risky) can adopt different strategies simply by tuning their values: by changing utility function may result in different agent personalities (e.g. individualist-autonomous, cooperative-collaborative), by changing confidence function agents become more or less self-confident etc.

5 DISCUSSION AND FUTURE WORKS

In this work we introduced anticipatory agents able to reconsider intentions on the basis of the expectations shaped on other agent recognized goals. The model enables to recognize other agent behavior as a BIC message and further provide abilities for signification and evaluation of related expectations. It further introduces noticeable properties for cognitive interaction:

- Coming through some lacks of the traditional deliberative architectures.
- Exploiting plan recognition mechanisms to really enhance pro-activeness and opportunism.
- Implementing a simplified approach to BIC which allows a wide spectrum of coordination issues to be modeled without relying on speech acts.

By introducing the ability for *intention reconsideration on the basis of expectations*, the model directly elicits anticipation and adaptivity to indeterminism, also allowing a strong subjective social interaction. Agents embeds adaptive capabilities to make anticipatory coordination an *emergent* property of the interactions: sociality is let emerge from the action and intelligence of individual agents.

Our experiments show that forms of silent, anticipatory coordination result in low cost, low complexity, highly effective mechanisms for coordination of agents with finite resources. The symbolic plan recognition engine, based on PHA, is very efficient and can serve concurrent hypotheses hence is able to predict agents pursuing multiple goals, namely interleaved plans. From a behavioral perspective, enhancements are in terms of pro-activeness, situated, real-time adaptivity to complex tasks. From the reasoning perspective, the model helps to disambiguate uncertainty, also providing strong adaptive means-end processing.

We guess this kind of architecture may contribute to the design of self-organizing/emergent societies, where virtual agents interact according to cognitive paradigms like *trust, reliance, delegation* [11, 6].

5.1 From simulation to real applications

We have made a series of assumption to simplify the domain. Moving from simulations to real applications, a series of key issues remains open. Firstly, the design of abstract actions to be recognized implicitly places the problem on the definition of the heuristics for the (reverse) process of recognizing their features: we define our representation as a series of abstract plan in first order logic terms, but to define the granularity of a real action may not be so obvious. We further have supposed complete plan representation handcrafted by the designers: in real-world scenarios this may result an intractable problem, due to complexity of tasks, heterogeneity of agents and multiplicity of their interactions. In addition, incorporating unknown goals and plans in plan representation is tractable only where the domain complexity is low [16].

Secondly, observing ongoing actions in real world application results a more complex task than we supposed: many actions have complex multi featured observable features, rather than few atomic features. Agents should embed components to resolve information processing form sensors to the internal symbolic representation. Furthermore, some of the features to observe may be intermittently lost due to noise or sensory failures. We assumed each action logically revealed without taking into account the information about its duration.

Finally, the computational costs of overwatching (matching observation against all possible actions performed by multiple agents and world contexts), may introduce overhead and serious problems in agent with finite resources.

In simplified domains (e.g. web applications), similar mechanisms for signification and plan recognition can be embedded with a different perspective, by utilizing hybrid approaches and smart infrastructures for a more *objective* coordination [24]. According to this paradigm, objective coordination is induced in MAS by means of ad-hoc abstractions of *coordination artifacts* [19] that may mediate interactions and provide coordination services. Coordination artifacts could be engineered to support previous knowledge of the plans used in the application domain and to dislocate (and automate) observation activities. As showed by many studies (i.e. [27]), this kind of infrastructure alleviates complex burdens for the involved agents: they can refer to the provided services in an uncoupled way and then decide autonomously by evaluating utilities to ascribe to the intentional options.

ACKNOWLEDGEMENTS

The research of this paper has been funded by the EU Project MindRACES within the Cognitive Systems area: Contract Number: FP6-511931. (www.mindRaces.org). We would like to thank aliCE[8] group at University of Bologna and Jadex[9] group at University of Hamburg for their components and useful support.

REFERENCES

[1] R. Beckers, O.E. Holland, and Deneubourg J.L., 'From local actions to global tasks: Stigmergy in collective robotics', *Artificial Life IV*, 181–189, (1994).
[2] L. Braubach, A. Pokahr, and W. Lamersdorf, *Software Agent-Based Applications, Platforms and Development Kits*, chapter Jadex: A BDI Agent System Combining Middleware and Reasoning, Birkhuser Book, 2005.
[3] C. Castelfranchi, 'Mind as an anticipatory device: For a theory of expectations', *In BVAI 2005*, 258–276, (2005).
[4] C. Castelfranchi, *From Conversation to Interaction via Behavioral Communication*, chapter Theories and Practice in Interaction Design, 157–179, Erlbaum, 2006.
[5] C. Castelfranchi, 'Silent agents: From observation to tacit communication.', in *IBERAMIA-SBIA*, pp. 98–107, (2006).
[6] C. Castelfranchi and R. Falcone, 'Principles of trust for mas: cognitive anatomy, social importance, and quantification', *Proc. of the Int. Conf. on Multi-Agent Systems (ICMAS'98), Paris*, 72–79, (1998).
[7] C. Catelfranchi, 'Modelling social action for ai agents', *Artificial Intelligence*, **103**, 157–182, (1998).
[8] E. Charniak and R. Goldman, 'A bayesian model of plan recognition', *Artificial Intelligence*, **64(1)**, 53–79, (1993).
[9] E. Denti, A. Omicini, and A. Ricci, 'tuprolog: A light-weight prolog for internet applications and infrastructures', in *Practical Aspects of Declarative Languages, 3rd International Symposium (PADL 2001)*, (2001).
[10] R. Falcone and C. Castelfranchi, 'Chaplin: A chart based plan recognizer', in *Proc. of the 13th International Conf. of Avignon.*, (1993).
[11] R. Falcone and C. Castelfranchi, 'Towards a theory of delegation for agent-based systems', *Robotics and Autonomous Systems*, **24**, 141–157, (1998).
[12] C. Geib and R. Goldman, 'Plan recognition in intrusion detection systems', in *DARPA Information Survivability Conference and Exposition (DISCEX)*, (2001).
[13] M.R. Genesereth, M:L. Ginsberg, and Rosenschein J.S., 'Cooperation without communication', in *Proc. of the National Conference on Artificial Intelligence (AAAI-86)*, (1986).
[14] Piotr J. Gmytrasiewicz, Edmund H. Durfee, and David K. Wehe, 'A decision-theoretic approach to coordinating multi-agent interactions.', in *Proc of International Joint Cconference on Artificial Intelligence (IJCAI-91)*, pp. 62–68, (1991).
[15] R. Goldmand, C. Geib, and C. Miller, 'A new model of plan recognition.', in *15th Conference on Uncertainty in Artificial Intelligence*, (1999).
[16] M. J. Huber, E. H. Durfee, and M. P. Wellman, 'The automated mapping of plans for plan recognition', in *Proc of the 10th Conf. on Uncertainty in Artificial Intelligence*, (1994).
[17] M.J. Huber and E.H. Durfee, 'Deciding when to commit to action during observation-based coordination', in *Proc. of the 1st Int. Conf. on Multi-Agent Systems ICMAS95*, (1995).
[18] D. Kinny and M. P. Georgeff, 'Commitment and effectiveness of situated agents', in *Proc. of the Twelfth International Joint Conf. on Artificial Intelligence (IJCAI-91)*, Sydney, Australia, (1991).
[19] A. Omicini, A. Ricci, M. Viroli, C. Castelfranchi, and L. Tummolini, 'Coordination artifacts: Environment-based coordination for intelligent agents', in *3rd international Joint Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2004)*, New York, USA, (2004).
[20] M. Piunti, C. Castelfranchi, and R. Falcone, 'Surprise as shortcut for anticipation: clustering mental states in reasoning', in *In Proc. of International Joint Conference on Artificial Intelligence (IJCAI07)*, Hyberadad, India., (2007).
[21] David Poole, 'Probabilistic horn abduction and bayesian networks', *Artificial Intelligence*, **64**(1), 81–129, (1993).
[22] A.S. Rao, *A unified view of plans as recipes*, chapter Contemporary Action Theory, Kluwer Academic Publishers, 1997.
[23] A.S. Rao and M.P. Georgeff, 'Bdi agents: From theory to practice', *Proc. of the 1st conf. on MAS (ICMAS95)*, (1995).
[24] M. Schumacher, *Objective Coordination in Multi-Agent System Engineering*, Springer Berlin / Heidelberg, 2001.
[25] M. Schut and M. Wooldridge, 'Intention reconsideration in complex environments', in *Proc. of the 4th international conference on Autonomous agents-AGENTS '00*, (2000).
[26] Y. Shoham and M. Tennenholtz, 'On the synthesis of useful social laws for artificial agents societies', in *Proc. of of 10th National Conference on Artificial Intelligence (AAAI-92)*, (1992).
[27] L. Tummolini, C. Castelfranchi, A. Ricci, M. Viroli, and A.Omicini, '"Exhibitionists" and "voyeurs" do it better: A shared environment approach for flexible coordination with tacit messages', in *Environments for MultiAgent Systems*, Springer-Verlag, (2005).

[8] www.alice.unibo.it
[9] vsis-www.informatik.uni-hamburg.de/projects/jadex/