Planning in the Logics of Communication and Change (Extended Abstract)

Pere Pardo
Institut d'Investigació en Intel·ligència Artificial
Campus de la UAB s/n, E-08193
Bellaterra, Spain
pardo@iiia.csic.es

Mehrnoosh Sadrzadeh
Oxford University
Dept. of Computer Science, Wolfson Building
Parks Rd, Oxford OX13QD, UK
mehrnoosh.sadrzadeh@cs.ox.ac.uk

ABSTRACT

We adapt backward planning to Logics of Communication and Change (LCC), that model how do actions, announcements and sensing change facts and agents' beliefs. An LCC planner takes into account the epistemic effects of planned actions upon other agents, if their beliefs are relevant to her goals. Our results include: a characterization of frame axioms as theorems in $*$-free LCC, and soundness and completeness results for deterministic planning and strong planning in the non-deterministic case.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Intelligent Agents

General Terms

Algorithms, Theory

Keywords

Dynamic Epistemic Logic, Planning, Communication

1. INTRODUCTION

In the present contribution, we adapt backward planning techniques to the Logics of Communication and Change (LCC). An LCC reasoning agent (who can foresee the possible epistemic effects of her actions and communications) is endowed with planning abilities to achieve some goals by means of LCC action models. This greatly expands on the social complexity of multi-agent planning scenarios.

Example 1.1. *Agent* a *is having a party, and would like her friend* b *to assist without their friend* c*. If* b *is secretive, a private announcement to* b *will suffice. However, suppose that* b *tells everything to* c*. Yet, if* a *knows that* c *only assists to parties with beer, while* b*'s interests also include jazz music (the party will include both), a solution may consist in informing* b *only about jazz.*

Example 1.2. *Agent* a *just bet agent* b *10 coins that the next coin toss will land heads (h);* a *can sense and even flip the coin without* b *ever suspecting it. A successful plan seems to be:* toss *the coin; if sense that h, then announce it to* b*; otherwise flip the coin and announce h.*

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), 4-8 June 2012, Valencia, Spain.

Related Work: [1] studies forward planning in LCC [2], under a semantic approach. Because of the large number of LCC actions available (one announcement per formula) forward planning faces the state-explosion problem. Thus, (deduction-based) backward planning seems appropriate.

2. LCC AND FRAME AXIOMS.

Definition 2.1. *The language* $\mathcal{L}_{\mathrm{PDL}}$ *of* $*$*-free* PDL*, for a given sets of atoms* $p \in \mathsf{Var}$ *and agents* $a \in \mathsf{Ag}$ *is:*

$$\varphi \quad ::= \quad \mathsf{Var} \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid [\pi]\varphi$$
$$\pi \quad ::= \quad a \mid ?\varphi \mid \pi_1 ; \pi_2 \mid \pi_1 \cup \pi_2$$

with the usual abbreviations for $\bot, \vee, \leftrightarrow$ *and* $\langle\pi\rangle$.

For LCC, we read an atomic program $[a]$ as *a believes that*, composition ";" is nested belief, and $\cup$ defines group belief.

LCC extends $\mathcal{L}_{\mathrm{PDL}}$ with modalities for pointed action models $[\mathsf{U}, \mathsf{e}]$. An action model is $\mathsf{U} = (\mathsf{E}, \mathsf{R}, \mathsf{pre}, \mathsf{post})$, with an action $\mathsf{e} \in \mathsf{E}$ being defined by a *precondition* $\mathsf{pre}(\mathsf{e})$, a LCC-formula, and a *postcondition* $\mathsf{post}(\mathsf{e})$, a substitution $\sigma : p \mapsto \varphi$ expressing that after executing e, the truth-value of p becomes that of φ (before the execution). In the present paper, though, we limit to the case $\sigma(p) \in \{\top, p, \bot\}$, studied in [3]. The *accessibility relations* $\mathsf{e}\mathsf{R}(a)\mathsf{f}$ denote actions f that cannot be distinguished from e by a. The skip action is given by the identity substitution. A truthful (resp. lying) communication of p by agent a to a set of (credulous) agents $B \subseteq \mathsf{Ag}$, denoted $p!^a_B$ (resp. $p\dagger^a_B$) has $\mathsf{pre}(p!^a_B) = p$ (resp. $\mathsf{pre}(p\dagger^a_B) = \neg p$).

We further extend the language of LCC with composition $\otimes$ and choice $\cup$ for action models.

Proposition 2.2. *The axioms of LCC [2] plus the next two axioms are a complete axiomatization of* $\mathrm{LCC} + \{\otimes, \cup\}$.

$$[\mathsf{U}^{\cup}, \mathsf{e} \cup \mathsf{e}']\varphi \quad \leftrightarrow \quad [\mathsf{U}, \mathsf{e}]\varphi \wedge [\mathsf{U}, \mathsf{e}']\varphi$$
$$[\mathsf{U} \otimes \mathsf{U}, \mathsf{e} \otimes \mathsf{e}']\varphi \quad \leftrightarrow \quad [\mathsf{U}, \mathsf{e}][\mathsf{U}, \mathsf{e}']\varphi$$

Frame axioms describe the conditions for a formula φ to be preserved after executing e. The presence of ontic actions makes LCC frame axioms $\mathsf{FA}(\mathsf{e}, \varphi)$ non-trivial, see Figure 1. The naive form cannot address the cases $p \vee q$ or $[a]p$.

Proposition 2.3. *The frame axioms* $\mathsf{FA}(\mathsf{e}, \varphi)$ *as in Fig. 1 (Right) can be inductively defined, and are valid in* LCC*: if* $\mathsf{cond}(\mathsf{FA}(\mathsf{e}, \varphi))$ *holds, then* $\models \mathsf{ant}(\mathsf{FA}(\mathsf{e}, \varphi)) \to [\mathsf{U}, \mathsf{e}]\varphi$.

if $\not\models [\mathsf{U},\mathsf{e}]\neg\varphi$ then $\models \varphi \to [\mathsf{U},\mathsf{e}]\varphi$	if $\mathsf{cond}(\mathsf{FA}(\mathsf{e},\varphi))$ then $\models \mathsf{ant}(\mathsf{FA}(\mathsf{e},\varphi)) \to [\mathsf{U},\mathsf{e}]\varphi$

Figure 1: Frame axiom for e,φ: (Left) Naive form. (Right) Correct form.

3. DETERMINISTIC PLANNING.

A planning domain is defined by a set $A \subseteq \mathsf{E}$ of available actions, and a pair (T, G), where $T, G \subseteq \mathcal{L}_{\mathrm{PDL}}$ describe the initial state and goals. Deterministic actions are just some subset $A \subseteq \mathsf{E}$ in LCC + $\{\otimes\}$. Given $\mathsf{e} \in A$, its effects are $\mathsf{X}(\mathsf{e}) = \{\psi \in \mathcal{L}_{\mathrm{PDL}} :\models \mathsf{pre}(\mathsf{e}) \to [\mathsf{U},\mathsf{e}]\psi\}$.

Definition 3.1. *Given a planning domain (T, G), actions A, and a program π, we say π is a* solution *for (T, G) in A iff (1) $\vdash \bigwedge T \to \langle\pi\rangle\top$, and (2) $\vdash \bigwedge T \to [\pi] \bigwedge G$.*

A solution must (1) be executable in T, and (2) lead to G. To solve a planning domain (T, G), we adopt the Breadth First Search (BFS) for incremental backward planning: starting with the empty plan for G, at each step $\pi_k = (\kappa_0, \dots, \kappa_k)$ we add a step $\pi_{k+1} = (\kappa_0, \dots, \kappa_k, \kappa_{k+1})$, delete the open goals of π_k enforced by κ_{k+1}, and add as new open goals $\mathsf{pre}(\kappa_{k+1})$. This step κ can be an action step $\mathsf{e} \in A$, or a proof step $\mathcal{A}$. Proof steps split complex goals, e.g. $\varphi\wedge\psi$, into simpler goals φ, ψ each of which can directly be enforced by some action $\mathsf{e} \in A$. This is done by means of a planned LCC-proof $\mathcal{A} = \langle\varphi, \psi, \dots, \varphi \wedge \psi\rangle$, where $\mathsf{pre}(\mathcal{A}) = \{\varphi, \psi\}$ denotes the (non-tautological) premisses of $\mathcal{A}$ and $\mathsf{X}(\mathcal{A}) = \varphi\wedge\psi$ its conclusion. Action steps must respect the frame axioms $\mathsf{FA}(\mathsf{e}, \varphi)$ for each goals φ in π_k unaddressed by π_{k+1}. That is, for e_{k+1} to refine π_k into a *plan* π_{k+1}, the condition $\mathsf{cond}(\mathsf{FA}(\mathsf{e}_{k+1}, \varphi))$ must be true, and $\mathsf{ant}(\mathsf{FA}(\mathsf{e}_{k+1}, \varphi))$ must be added as an open goal of π_{k+1}. Finally, the set of open goals of π_{k+1} must also be consistent. Similar conditions apply to proof steps $\mathcal{A}$, to make π_{k+1} a plan. Note the plan $\pi = (\mathsf{e}_0, \mathcal{A}_0, \dots, \mathcal{A}_n, \mathsf{e}_n)$ translates into logical form $[\mathsf{U}, \mathsf{e}_n]$,..., $[\mathsf{U}, \mathsf{e}_0]$, with action steps in inverse order, and where proof steps are omitted (LCC will enforce them anyway).

Theorem 3.2. *Let $(\mathsf{e}_0, \dots, \mathsf{e}_n)$ be an output of the BFS algorithm for (T, G) in A. Then $[\mathsf{U}, \mathsf{e}_n] \dots [\mathsf{U}, \mathsf{e}_0]$ is a solution for (T, G). Conversely, suppose some deterministic solution $[\mathsf{U}, \mathsf{e}_n] \dots [\mathsf{U}, \mathsf{e}_0]$ exists for (T, G) in A. Then the BFS algorithm terminates with a solution for (T, G) in A.*

Planning in others' shoes For multi-agent scenarios, we can define an algorithm that computes the reactions to one's plan by other planner agents. Then, a plan is called *stable* if these reactions do not lead to a state where G is not satisfied.

Example 3.3. *(Cont'd) Recall Example 1.1.* b*'s goals are* $(\mathsf{beer} \vee \mathsf{jazz}) \to @\mathsf{party}(\mathsf{b})$ *as well as* $\{[\mathsf{b}]\varphi \to [\mathsf{c}]\varphi\}_{\varphi\in\mathcal{L}_{\mathrm{PDL}}}$*; and* c *has goal* $\mathsf{beer} \to @\mathsf{party}(\mathsf{b})$*. It can be seen that the naive solution* $\mathsf{beer}!^{\mathsf{a}}_{\mathsf{b}}$ *is not stable: agents' reactions lead to the output* $(\mathsf{beer}!^{\mathsf{a}}_{\mathsf{b}} \otimes \mathsf{beer}!^{\mathsf{b}}_{\mathsf{c}} \otimes \mathsf{go.party}(\mathsf{b}) \otimes \mathsf{go.party}(\mathsf{c}))$*, which makes* $\neg@\mathsf{party}(\mathsf{c})$ *false. In contrast,* $\mathsf{jazz}!^{\mathsf{a}}_{\mathsf{b}}$ *is stable: the output is* $\langle(\mathsf{jazz}!^{\mathsf{a}}_{\mathsf{b}} \otimes \mathsf{jazz}!^{\mathsf{b}}_{\mathsf{c}} \otimes \mathsf{go.party}(\mathsf{b})\rangle$.

4. NON-DETERMINISTIC PLANNING

For planning involving actions with disjunctive effects $\varphi_1 \vee \varphi_2$, one first stipulates actions e_i with $\varphi_i \in \mathsf{X}(\mathsf{e}_i)$. While $\mathsf{e}_1 \cup \mathsf{e}_2 \in A$ is available, individual actions e_i are not: $\mathsf{e}_i \in \mathsf{E} \smallsetminus A$.

A plan is now a 4-tuple: *plan* = (*sequence of actions*, *open goals*, *initial state*, *original goals*). We reduce the problem of building a non-deterministic plan into that of solving a sequence of deterministic planning problems. To do so, we define a plan set Π as a sequence of plans $\Pi = \langle\pi^{\xi}, \pi^{\xi'}, \dots\rangle$ enumerated by sequences $\xi \in \{\emptyset\} \cup \{1, 2\}^{<\omega}$ and ordered lexicographically, e.g. $\emptyset <_{\mathsf{lex}} 1 <_{\mathsf{lex}} 11 <_{\mathsf{lex}} 12 <_{\mathsf{lex}} 2$. This ordering represents the order in which plans are solved. See Figure 2 for an illustration of the algorithm in Example 1.2. Non-deterministic planning is done by a series of BFS searches.

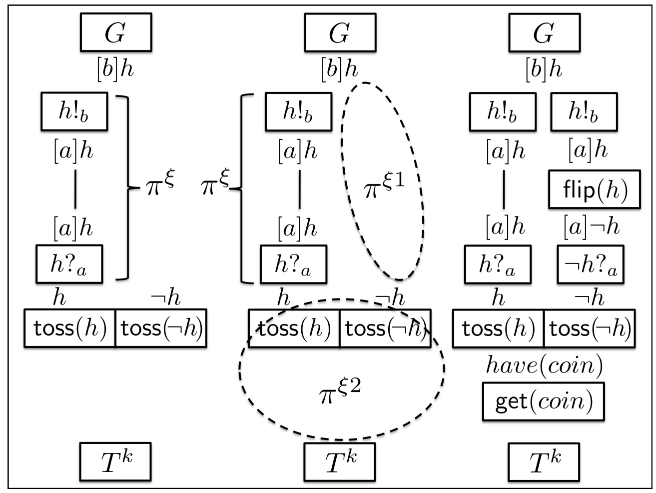

Figure 2: Refining π^{ξ} with $\mathsf{toss}(h) \cup \mathsf{toss}(\neg h)$ splits π^{ξ} into three deterministic plan search problems $\pi^{\xi}, \pi^{\xi 1}, \pi^{\xi 2}$.

To make sure that a plan π with $\mathsf{e}_1 \cup \mathsf{e}_2$ is logically acceptable, a reset action ρ might be needed to harmonize effects: $\mathsf{e}_1 \cup (\mathsf{e}_2 \otimes \rho)$. These ρ do not contribute to the success of π.

Theorem 4.1. *Let π be an output of the BFS for (T, G). Then π is a solution for (T, G) in A. Conversely, if a non-deterministic solution exists, then the BFS output for (T, G) in A is a solution.*

Conclusions and Future Work

We studied backward deterministic and strong planning for LCC logics. Several directions seem interesting: belief revision, the $*$ operator for strong cyclic planning and decidability/complexity issues, among others.

Acknowledgements: This work has been funded by projects AT (CSD 2007-022); LoMoReVI (FFI2008-03126-E/FILO FP006); ARINF (TIN2009-14704-C03-03); and the GenCat 2009-SGR-1434 and EPSRC grant EP/J002607/1.

5. REFERENCES

[1] T. Bolander and M. Birkegaard Andersen *Epistemic planning for single- and multi-agent systems* Journal of Applied Non-Classical Logics, (in press)

[2] J. van Benthem, J. van Eijck and B. Kooi *Logics of Communication and Change* Information and Computation, vol 204 pp. 1620–1662 (2006)

[3] H. van Ditmarsch, B. Kooi *Semantic results for ontic and epistemic change*, in Bonanno, van der Hoek and Wooldridge (eds.), LOFT 7, pp. 87–117 (2008)