Hierarchical planning in BDI agent programming languages: A formal approach

Sardina, Sebastian; De Silva, Lavindra; Padgham, Lin

https://researchrepository.rmit.edu.au/esploro/outputs/9921862538401341/filesAndLinks?institution=61RMIT_INST&index=null

Sardina, S., De Silva, L., & Padgham, L. (2006). Hierarchical planning in BDI agent programming languages: A formal approach. Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS 2006), 1001–1008.

https://researchrepository.rmit.edu.au/discovery/fulldisplay/alma9921862538401341/61RMIT_INST:ResearchRepository

Document Version: Accepted Manuscript

Repository homepage: <https://researchrepository.rmit.edu.au>

© ACM, 2006 This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in 'Proceedings of the fifth international joint conference on autonomous agents and multi agent systems (AAMAS 2006)'

<http://doi.acm.org/10.1145/1160633.1160813>

Downloaded On 2023/04/26 19:18:45 +1000

Hierarchical Planning in BDI Agent Programming Languages: A Formal Approach*

Sebastian Sardina
RMIT University
Melbourne, Australia
ssardina@cs.rmit.edu.au

Lavindra de Silva
RMIT University
Melbourne, Australia
ldesilva@cs.rmit.edu.au

Lin Padgham
RMIT University
Melbourne, Australia
linpa@cs.rmit.edu.au

ABSTRACT

This paper provides a general mechanism and a solid theoretical basis for performing planning within Belief-Desire-Intention (BDI) agents. BDI agent systems have emerged as one of the most widely used approaches to implementing intelligent behavior in complex dynamic domains, in addition to which they have a strong theoretical background. However, these systems either do not include any built-in capacity for “lookahead” type of planning or they do it only at the implementation level without any precise defined semantics. In some situations, the ability to plan ahead is clearly desirable or even mandatory for ensuring success. Also, a precise definition of how planning can be integrated into a BDI system is highly desirable. By building on the underlying similarities between BDI systems and Hierarchical Task Network (HTN) planners, we present a formal semantics for a BDI agent programming language which cleanly incorporates HTN-style planning as a built-in feature. We argue that the resulting integrated agent programming language combines the advantages of both BDI agent systems and hierarchical offline planners.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*Intelligent Agents, Languages and structures*

Keywords

BDI agent-oriented programming, HTN planning

1. INTRODUCTION

The BDI (Belief-Desire-Intention) model is a popular and well-studied architecture of agency for intelligent agents situated in complex and dynamic environments. The model has its roots in philosophy with Bratman’s [2] theory of practical reasoning and Dennett’s theory of intentional systems [9]. There are a number of agent programming languages in the BDI tradition, such as PRS [13], AGENTSPEAK [19], 3APL [12], JACK [3], CAN [24].

*We would like to acknowledge the support of Agent Oriented Software and of the Australian Research Council under the grant “Learning and Planning in BDI Agents” (number LP0560702).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS’06 May 8–12 2006, Hakodate, Hokkaido, Japan.
Copyright 2006 ACM 1-59593-303-4/06/0005 ...\$5.00.

BDI agent-oriented systems are extremely flexible and responsive to the environment, and as a result, well suited for complex applications with real-time reasoning and control requirements. However, a limitation of these systems is that they normally do no *lookahead* or *planning* in the traditional sense; execution is based on a user-provided “plan library” to achieve goals. BDI frameworks rely entirely on context sensitive subgoal expansion, acting as they go. In some circumstances, however, lookahead deliberation (i.e., hypothetical reasoning) about the effects of one choice of expansion over another is clearly desirable, or even mandatory in order to guarantee goal achievability and to avoid undesired situations. In general, this is the case when (a) important resources may be used in taking actions that do not lead to a successful outcome; (b) actions are not always reversible and may lead to states from which there is no successful outcome; (c) execution of actions take substantially longer than “thinking” (or planning); and (d) actions have side effects which are undesirable if they turn out not to be useful.

In this paper, we develop a traditional BDI-style agent programming language that includes an *on-demand planning mechanism* in the style of Hierarchical Task Networks (HTN), whose semantics and implementations are well understood in the planning community [11]. The language we propose, named CANPLAN, provides a flexible approach regarding when to perform full lookahead, and is provably more expressive than either BDI or HTN systems alone. CANPLAN is based on CAN [24] and AGENTSPEAK [19]. One could argue, of course, that it is always possible, in critical situations, to explicitly program lookahead within existing BDI systems. However, such code would generally be domain dependent, can be fairly complex, and would lie outside the infrastructure support provided by the BDI agent platform. Alternatively, there are many frameworks that attempt to interleave BDI-type execution with offline planning (e.g., [1, 23, 10, 17, 14]). Still, these are mostly implemented systems with no precise semantics and with little programmer control over when to plan. Our approach, instead, is to provide a formal specification of planning as a *built-in feature* of the BDI infrastructure that the programmer can use as appropriate.

The contributions of this paper are threefold. Firstly, a precise account of planning within a typical BDI agent programming language is provided. Secondly, the intrinsic relationship between lookahead planning in the context of BDI agents and the HTN approach to planning is formally explored. Lastly, the semantics of CAN given in [24] is substantially improved and simplified.

The rest of the paper is organised as follows. In section 2, we provide a brief overview of BDI agent programming languages and HTN planners; we also provide an informal discussion on their similarities. In section 3, we describe the basic BDI agent language we will use, namely the CAN notation described in [24], but with some modifications to include actions with preconditions and ef-

© ACM, 2006 This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in "Proceedings of the fifth international joint conference on autonomous agents and multi agent systems (AAMAS 2006)" <http://doi.acm.org/10.1145/1160633.1160813>

Citation: Sardina, S, De Silva, L and Padgham, L 2006, 'Hierarchical planning in BDI agent programming languages: A formal approach', in P. Stone & G. Weiss (ed.) Proceedings of the fifth international joint conference on autonomous agents and multi agent systems (AAMAS 2006), Hakodate, Japan, 2006.

DRAFT -- DRAFT -- DRAFT -- DRAFT -- DRAFT --

fects, multiple variable bindings, and a simpler, though equivalent, account of declarative goals. We chose CAN from the numerous available options because it has the desirable features of (a) combining a declarative and procedural view of *goals*, and (b) capturing the semantics of BDI failure recovery and goal persistence. In section 4, we develop CANPLAN, our new integrated account of planning and BDI execution. Besides showing some intuitively expected properties for the combined framework, we prove that, under suitable assumptions, CANPLAN’s planning module reduces to HTN planning. In section 5, a brief discussion on a prototype implementation is given. Section 6 discusses related work. Finally, in section 7, we draw conclusions and outline future lines of research.

2. BDI AND HTN SYSTEMS

There are a number of BDI agent languages and HTN systems. We provide a brief abstract overview of these in order to comment on their similarities, as background for our integrated approach.

2.1 BDI Agent Programming Languages

Generally speaking, BDI agent-oriented programming languages are built around an explicit representation of beliefs, desires, and intentions. A BDI architecture addresses how these components are represented, updated, and processed to determine the agent’s actions. There are a substantial number of implemented BDI systems, as well as a number of formally specified languages.

An agent consists, basically, of a belief base $\mathcal{B}$, a set of recorded pending events (goals), a plan library Π , and an intention base Γ . The *belief base* encodes the agent’s knowledge about the world. The *plan library* contains *plan rules* of the form $e : \psi \leftarrow P$ encoding a *plan-body* or *program* P for handling an event-goal e when *context* condition ψ is believed to hold. The *intention base* contains the current, partially instantiated, plans the agent has already committed to in order to handle or achieve some event-goal.

A BDI system responds to *events*, the inputs to the system, by committing to handle one pending event-goal, selecting a plan rule from the library, and placing its plan-body/program into the intention base. The execution of this program may, in turn, post new sub-goal events to be achieved. If at any point a program fails, then an alternative plan rule is found and its plan-body is placed into the intention base for execution. This process repeats until a plan succeeds completely or until there are no more applicable plans, in which case failure is propagated to the event-goal.

In section 3, we shall discuss in detail one formal BDI language of this sort, namely, CAN [24].

2.2 HTN Planning

Hierarchical Task Network (HTN) planning is an approach to planning based on the decomposition of (high-level) *tasks* in order to accomplish an (initial) *task network*. Two examples of HTN systems include SHOP [15] and its successor SHOP2 [16]. Below, we mostly follow the definitions of HTN-planning from [11].

Tasks can be of two types. A *primitive task* is an action $act(\vec{x})$ that can be directly executed by the agent (e.g., $drive(x_1, x_2)$). A (high-level) *compound task* $e(\vec{x})$ is one that cannot be executed directly (e.g., $build_trip(origin, dest)$). A *task network* $d = [T, \phi]$ is a collection of tasks T that need to be accomplished and a boolean formula of constraints ϕ . Constraints impose restrictions on the ordering of the tasks ($e \prec e'$), on the binding of variables ($x = x'$) and ($x = c$) (c is a constant), and on what literals must be true before or after each task (l, e), (e, l), and (e, l, e'). A *method* (e, ψ, d) encodes a way of decomposing a high-level compound task e into lower-level tasks using task network d when ψ holds. Methods provide the procedural knowledge of the domain.

An HTN *planning domain* $\mathcal{D} = (\Pi, \Lambda)$ consists of a library Π of methods and a library Λ of primitive tasks. Each primitive task in Λ is a STRIPS style action with corresponding preconditions and effects in the form of *add* and *delete* lists. An HTN *planning problem* $\mathbf{P}$ is the triple $\langle d, \mathcal{B}, \mathcal{D} \rangle$ where d is the task network to accomplish, $\mathcal{B}$ is the initial belief state (i.e., a set of all ground atoms that are true in $\mathcal{B}$), and $\mathcal{D}$ is a planning domain. A *plan* σ is a sequence $act_1 \cdot \dots \cdot act_n$ of ground actions (that is, ground primitive tasks).

Given a planning problem instance $\mathbf{P}$, the planning process involves selecting and applying an applicable reduction method from $\mathcal{D}$ to some compound task in d . This results in a new, and typically more “primitive,” task network d' . This reduction process is repeated until only primitive tasks (i.e., actions) remain. If no applicable reduction can be found for a compound task at any stage, the planner “backtracks” and tries an alternative reduction for a compound task previously reduced. If all compound tasks can eventually be reduced, a plan solution σ is obtained.

In [11], a clear operational semantics for HTN planning was given. The set of plans $sol(d, \mathcal{B}, \mathcal{D})$ that solves a planning instance $\mathbf{P} = \langle d, \mathcal{B}, \mathcal{D} \rangle$ is defined as $sol(d, \mathcal{B}, \mathcal{D}) = \bigcup_{n < \omega} sol_n(d, \mathcal{B}, \mathcal{D})$, where $sol_n(d', \mathcal{B}, \mathcal{D})$ is, in turn, defined as follows:

$$\begin{aligned} sol_1(d, \mathcal{B}, \mathcal{D}) &= comp(d, \mathcal{B}, \mathcal{D}), \\ sol_{n+1}(d, \mathcal{B}, \mathcal{D}) &= sol_n(d, \mathcal{B}, \mathcal{D}) \cup \bigcup_{d' \in red(d, \mathcal{B}, \mathcal{D})} sol_n(d', \mathcal{B}, \mathcal{D}). \end{aligned}$$

Intuitively, $comp(d, \mathcal{B}, \mathcal{D})$ is the set of all plan *completion* of a network d containing only primitive tasks (i.e., plans for which the constraint formula ϕ in d is satisfied), and $red(d, \mathcal{B}, \mathcal{D})$ is the set of all *reductions* of d in $\mathcal{B}$ by methods in $\mathcal{D}$. We refer to [11] for more details on HTN and its formal semantics.

2.3 Similarities Between HTN and BDI

As stated in [7], BDI agent programming languages and HTN planners share many similarities despite their different purposes. The similarities come from the knowledge used by both systems as well as from how this knowledge is manipulated to create solutions.

First of all, HTN systems and BDI languages assume an explicit representation of the agent’s knowledge (i.e., the belief base) and a set of primitive tasks or actions that the agent can directly execute in the world. Secondly, procedural knowledge about the domain is available in both HTN and BDI systems in the form of reduction methods and plan rules, respectively. Thirdly, and most importantly, both systems create solutions by reducing higher-level entities into lower-level ones by appealing to a given set of reduction recipes. Whereas a BDI system “reduces” an event into an plan-body/program using a plan rule from the plan library, an HTN planner reduces a compound task into a task network using a reduction method from the method library.

The following table gives an indication of the mapping between HTN and BDI entities.

BDI SYSTEMS	HTN SYSTEMS
belief base	state
plan library	method library
event	compound task
action	primitive task
plan-body/program	network task
plan rule	method
plan rule context	method precondition
test ? l in plan-body	state constraints
sequence in plan-body	ordering constraint $\prec$
parallelism in plan-body	no ordering constraint

The above table is not complete—while some entities have a straightforward mapping some others require a more elaborate translation (we refer to [7, 22] for a more detailed mapping).

BDI agent systems and HTN planners, despite their close relationship, differ fundamentally in their objectives. The former are focused on the *execution* of agent programs where “backtracking” can only happen in the real world. The latter, in contrast, are concerned with *hypothetical reasoning* about actions and their potential interactions within a whole plan for achieving a goal/task.

3. THE BASIC BDI LANGUAGE

The CAN (Conceptual Agent Notation) notation [24] is a high-level plan language in the style of typical agent languages, both in the BDI tradition and elsewhere (e.g., AGENTSPEAK [19], 3APL [12, 21], and even CONGOLOG [5, 6]). Its syntax and semantics attempt to extract the essence of a class of implementable agent platforms and could be considered as a superset of AGENTSPEAK (see [24]). Unlike AGENTSPEAK, though, the semantics for CAN includes both *failure handling* and *declarative goals*—two appealing features for our planning agents.

An agent is created by the specification of a set of base beliefs $\mathcal{B}$ and a set of plans Π . The *belief base* of an agent is a set of formulas from some (knowledge representation) logical language. The programmer may choose any logical language; all that is required is for operations to exist that check whether a condition ϕ —a logical formula over the agent’s beliefs—follows from a belief set (i.e., $\mathcal{B} \models \phi$), and to add and delete a belief b to and from a belief base (i.e., $\mathcal{B} \cup \{b\}$ and $\mathcal{B} \setminus \{b\}$, respectively). In practice, however, the belief base contains ground belief *atoms* in a first-order language.

As explained in section 2, an agent *plan library* Π consists of a collection of plan rules of the form $e : \psi \leftarrow P$, where e is an event and ψ is the context condition which must be true in order for the plan-body P to be applicable.¹ The *plan-body* or *program* P is built from primitive actions *act* that the agent can execute directly, operations to add $+b$ and delete $-b$ beliefs, tests for conditions $?\phi$, and events or (internal) achievement goals $!e$. Complex plans can be specified using sequencing $P_1; P_2$, parallelism $P_1 \parallel P_2$, and declarative goals $\text{Goal}(\phi_s, P, \phi_f)$ (explained later). Hence, the *user language* is described by the following grammar:

$$P ::= \text{act} \mid +b \mid -b \mid ?\phi \mid !e \mid P_1; P_2 \mid P_1 \parallel P_2 \mid \text{Goal}(\phi_s, P_1, \phi_f).$$

There are also a number of auxiliary plan forms which are used internally when assigning semantics to constructs: basic (terminating) program *nil*; and compound plans like $P_1 \triangleright P_2$, which executes P_1 and then executes P_2 only if P_1 failed, and $\langle \psi_1 : P_1, \dots, \psi_n : P_n \rangle$, which is used to encode a set of (relevant) guarded plans. The *full language* is therefore described by the following grammar:

$$P ::= \text{nil} \mid \text{act} \mid ?\phi \mid P_1 \triangleright P_2 \mid \langle \psi_1 : P_1, \dots, \psi_n : P_n \rangle \mid +b \mid -b \mid !e \mid P_1; P_2 \mid P_1 \parallel P_2 \mid \text{Goal}(\phi_s, P_1, \phi_f).$$

In contrast with [19, 24], we take actions as the usual basic means of the agent to change its environment and, hence, actions may have preconditions and effects. One possibility would be to follow [12] and assume that a *partial* function $\mathcal{T}$ specifying the update semantics of basic actions is given: if $\mathcal{T}(\text{act}, \mathcal{B})$ is defined, it yields the new updated belief base $\mathcal{B}'$; otherwise, we say that the action’s precondition is not met in $\mathcal{B}$. However, for simplicity, we shall restrict ourselves to agents that are equipped with a (simple) STRIPS-like *action description library* Λ containing rules of the form $\text{act} : \psi_{\text{act}} \leftarrow \Phi_{\text{act}}^-; \Phi_{\text{act}}^+$, one for each action type in the domain. Formula ψ_{act} corresponds to the action’s precondition, and Φ_{act}^+ and Φ_{act}^- stand for the add and delete lists of atoms, respectively.² For example, action $\text{move}(x, y, z)$, which moves object x

¹An omitted ψ is equivalent to *nil*. Notice that e, ψ , and P may contain free variables; a plan rule is of the form $e(\vec{x}) : \psi(\vec{x}, \vec{y}) \leftarrow P(\vec{x}, \vec{y}, \vec{z})$, where $\vec{x}, \vec{y}$ and $\vec{z}$ are vectors of (distinct) variables.

²Free variables in $\psi_{\text{act}}, \Phi_{\text{act}}^-$ and Φ_{act}^+ are free in *act* too.

from y to z , could be represented in Λ as follows:

$$\text{move}(x, y, z) : \text{Free}(z) \wedge \text{At}(x, y) \leftarrow \{\text{Free}(z), \text{At}(x, y)\}; \{\text{Free}(y), \text{At}(x, z)\}.$$

Next, we show the operational semantics for the above language along the lines of [24]. A *transition relation* $\longrightarrow$ on so-called *configurations* is defined by a set of derivation rules. A transition $C \longrightarrow C'$ specifies that executing configuration C a *single step* yields configuration C' . We write $C \longrightarrow$ to state that there exists C' such that $C \longrightarrow C'$, and $\longrightarrow^*$ to denote the usual reflexive transitive closure of $\longrightarrow$. A derivation rule consists of a, possibly empty, set of premises, which are transitions together with some auxiliary conditions, and a single transition conclusion derivable from these premises. (see [18] for more on operational semantics).

Two types of transitions will be used to define the semantics of our agents. The first type defines what it means to execute a single intention and is defined in terms of *basic* configurations. The second type of transition is defined in terms of the first type and defines what it means to execute an agent. A *basic configuration* is a tuple $\langle \mathcal{B}, \mathcal{A}, P \rangle$ consisting of the current belief base $\mathcal{B}$ of the agent, the sequence $\mathcal{A}$ of primitive actions executed so far, and the plan-body P being executed (i.e., the current intention).³

Here are some of the core derivation rules for the language:

$$\begin{array}{c} \Delta = \{\psi_i \theta : P_i \theta \mid e' : \psi_i \leftarrow P_i \in \Pi \wedge \theta = \text{mgu}(e, e')\} \\ \hline \langle \mathcal{B}, \mathcal{A}, !e \rangle \longrightarrow \langle \mathcal{B}, \mathcal{A}, \langle \Delta \rangle \rangle \quad \text{Event} \\ \hline \frac{\psi_i : P_i \in \Delta \quad \mathcal{B} \models \psi_i \theta}{\langle \mathcal{B}, \mathcal{A}, \langle \Delta \rangle \rangle \longrightarrow \langle \mathcal{B}, \mathcal{A}, P_i \theta \triangleright \langle \Delta \setminus P_i \rangle \rangle} \quad \text{Sel} \\ \hline \frac{\langle \mathcal{B}, \mathcal{A}, P_1 \rangle \not\vdash^{\text{di}}}{\langle \mathcal{B}, \mathcal{A}, (P_1 \triangleright P_2) \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}, \mathcal{A}, P_2 \rangle} \triangleright_f \frac{\mathcal{B} \models \phi}{\langle \mathcal{B}, \mathcal{A}, ?\phi \rangle \longrightarrow \langle \mathcal{B}, \mathcal{A}, \text{nil} \rangle} ? \\ \hline \frac{a : \psi \leftarrow \Phi^+; \Phi^- \in \Lambda \quad a\theta = \text{act} \quad \mathcal{B} \models \psi\theta}{\langle \mathcal{B}, \mathcal{A}, \text{act} \rangle \longrightarrow \langle (\mathcal{B} \setminus \Phi^- \theta) \cup \Phi^+ \theta, \mathcal{A} \cdot \text{act}, \text{nil} \rangle} \quad \text{act} \\ \hline \frac{\langle \mathcal{B}, \mathcal{A}, P_1 \rangle \longrightarrow \langle \mathcal{B}', \mathcal{A}', P' \rangle}{\langle \mathcal{B}, \mathcal{A}, (P_1; P_2) \rangle \longrightarrow \langle \mathcal{B}', \mathcal{A}', (P'; P_2) \rangle} \quad \text{Seq} \end{array}$$

Rule *Event* handles achievement goal events by collecting all *relevant* plans for the event in question. Rule *Sel* selects one *applicable* plan from a set of (remaining) relevant plans: program $P \triangleright \langle \Delta \rangle$ states that program P should be tried first, falling back to the remaining alternatives in Δ if required. Notice that plan rules’ context conditions are handled in a lazy manner. Rule $?$ deals with test goals by checking that the condition follows from the current belief base, whereas rule *act* handles the case of primitive actions by using the domain action description library Λ . Rule *Seq* handles sequencing of programs in the usual way. Rule $\triangleright_f$ is used along with rule *Sel* for *failure handling*: if the current plan $P_i \theta$ for a goal fails (i.e., at some point the precondition of an action or a test goal is not met), rule $\triangleright_f$ applies first, and eventually, rule *Sel* may select another *applicable* alternative for the event-goal, if any.

A central distinguishing feature of CAN is its $\text{Goal}(\phi_s, P, \phi_f)$ goal construct, which provides a mechanism for representing both declarative and procedural aspects of goals. Intuitively, a *goal-program* $\text{Goal}(\phi_s, P, \phi_f)$ states that we should achieve the (declarative) goal ϕ_s by using (procedural) plan P ; failing if ϕ_f becomes

³Strictly speaking, the plan and action libraries Π and Λ should also be part of basic configurations. For legibility purposes, we omit them as they are assumed to be static entities. Configurations must also include a variable substitution θ for keeping track of all bindings done so far during the execution of a plan-body. Again, for legibility, we keep substitutions implicit in places where they need to be carried across multiple rules. See [12] on how substitutions are propagated across derivation rules for 3APL.

true. The execution of a goal-program is consistent with some desired properties of declarative goals (namely, persistent, possible, and unachieved). For instance, if P is fully executed but ϕ_s is still not true, P will be re-tried; and if ϕ_s becomes true during P 's execution, the whole goal will succeed immediately.

In order to capture the desired behaviour of goal-programs, a sophisticated operational semantics was given in [24], based on explicit exceptions, a set of conditions being “watched,” and derivation rules with priorities. Here, we provide an alternative much simpler semantics that is equivalent to the original one. The following is the new set of rules for goal-programs:

$$\begin{array}{c}
\frac{P \neq P_1 \triangleright P_2 \quad B \not\models \phi_s \vee \phi_f}{\langle B, A, \text{Goal}(\phi_s, P, \phi_f) \rangle \longrightarrow \langle B, A, \text{Goal}(\phi_s, P \triangleright P, \phi_f) \rangle} G_I \\
\frac{B \models \phi_s}{\langle B, A, \text{Goal}(\phi_s, P, \phi_f) \rangle \longrightarrow \langle B, A, \text{nil} \rangle} G_S \\
\frac{B \models \phi_f}{\langle B, A, \text{Goal}(\phi_s, P, \phi_f) \rangle \longrightarrow \langle B, A, ?\text{false} \rangle} G_f \\
\frac{P = P_1 \triangleright P_2 \quad B \not\models \phi_s \vee \phi_f \quad \langle B, A, P_1 \rangle \longrightarrow \langle B', A', P' \rangle}{\langle B, A, \text{Goal}(\phi_s, P, \phi_f) \rangle \longrightarrow \langle B', A', \text{Goal}(\phi_s, P' \triangleright P_2, \phi_f) \rangle} G_S \\
\frac{P = P_1 \triangleright P_2 \quad B \not\models \phi_s \vee \phi_f \quad \langle B, A, P_1 \rangle \not\longrightarrow}{\langle B, A, \text{Goal}(\phi_s, P, \phi_f) \rangle \longrightarrow \langle B, A, \text{Goal}(\phi_s, P_2 \triangleright P_2, \phi_f) \rangle} G_R
\end{array}$$

When a goal-program is first encountered during execution, rule G_I applies: G_I “initialises” the execution of a goal-program by setting the program in the goal to $P \triangleright P$, where the first P is to be executed and the second P is just used to carry the *original* plan P for (potential) use later on by rule G_R . The second and third rules handle the cases where either the success condition ϕ_s or the failure condition ϕ_f become true. The fourth rule G_S is the one responsible for performing a single step on an already initialised goal-program. Notice that the second part in the pair $P_1 \triangleright P_2$ remains constant. Finally, rule G_R restarts the original program (stored as the second program in pair $P_1 \triangleright P_2$) whenever the current program has finished, but the desired, and still possible, goal has not been achieved yet.

The above semantics for Goal is substantially simpler than the original one in [24] in that we do not appeal to explicit exceptions, “watched” conditions, or special prioritised derivation rules. Although it is not hard to prove that this alternative semantics is equivalent to the original one, due to lack of space, we do not do that here. Finally, we point out that, in the original semantics of CAN, an agent included also a goal base $\mathcal{G}$ to account for the declarative goals the agent has already committed to via goal-programs. Although not done in CAN, the goal base could *potentially* be used to perform (meta)reasoning about goals at the agent level execution, such as goal conflict detection/resolution ([20]). Since we are also not concerned in this paper with this type of reasoning, we completely omit the goal base from our agents.

Agent Level Execution

On top of the above basic rules, we define the evolution of an agent. An *agent configuration*, or just an agent, is a tuple of the form $\langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}, \mathcal{A}, \Gamma \rangle$ where $\mathcal{N}$ is the agent name, Λ is an action description library, Π is a plan library, $\mathcal{B}$ is a belief base, $\mathcal{A}$ is the sequence of actions already performed by the agent, and Γ is the set of current intentions (i.e., plan-bodies). Transitions between agent configurations are dictated by the following three rules:

$$\begin{array}{c}
\frac{P \in \Gamma \quad \langle B, A, P \rangle \longrightarrow \langle B', A', P' \rangle}{\langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}, \mathcal{A}, \Gamma \rangle \Longrightarrow \langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}, \mathcal{A}', (\Gamma \setminus \{P\}) \cup \{P'\} \rangle} A_{step} \\
\frac{e \text{ is a new external event}}{\langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}, \mathcal{A}, \Gamma \rangle \Longrightarrow \langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}, \mathcal{A}, \Gamma \cup \{!e\} \rangle} A_{event} \\
\frac{P \in \Gamma \quad \langle B, A, P \rangle \not\longrightarrow}{\langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}, \mathcal{A}, \Gamma \rangle \Longrightarrow \langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}, \mathcal{A}, \Gamma \setminus \{P\} \rangle} A_{clean}
\end{array}$$

The first rule performs a single step in one intention; the second rule creates a new intention from an external event; and the last rule removes a completed intention from the intention base (i.e., an intention *nil* or one that is blocked and cannot make a transition).

Next, we define the meaning of an agent execution and two related notions that will be used later in the paper.

DEFINITION 1 (BDI EXECUTION). A *BDI execution* E of an agent $C_0 = \langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}_0, \mathcal{A}_0, \Gamma_0 \rangle$ is a, possibly infinite, sequence of agent configurations $C_0 \cdot C_1 \cdot \dots \cdot C_n \cdot \dots$ such that $C_i \Longrightarrow C_{i+1}$, for every $i \geq 0$. A *terminating execution* is a finite execution $C_0 \cdot \dots \cdot C_n$ where $C_n = \langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}_n, \mathcal{A}_n, \{\} \rangle$. An *environment-free execution* is one in which rule A_{event} has not been used.

Sometimes we will be only interested in those steps of an execution where changes occur in either the executed actions or the belief of the agent—agent steps where the belief base and the executed actions remain unchanged can be disregarded. So, if $E = C_0 \cdot \dots \cdot C_n$ is a (finite) execution, then the *derived execution* $\bar{E}$ is the sequence of configurations obtained from E by deleting all configurations C_j of the sequence such that $\mathcal{B}_j = \mathcal{B}_{j+1}$ and $\mathcal{A}_j = \mathcal{A}_{j+1}$.

In addition, we give the following notation to track an intention during an execution. If $C_0 \cdot \dots \cdot C_n$ is a normal or derived execution and P is an intention in C_0 (i.e., $P \in \Gamma_0$), then the sequence $P_0 = P, P_1, \dots, P_n$ denotes P 's evolution within the execution and either (i) $P_i \in \Gamma_i$; or (ii) $P_i = \epsilon$, if the intention has already been removed from the intention base at some C_j , where $j \leq i$.

DEFINITION 2. Two, possibly derived, agent executions $C_0 \cdot \dots \cdot C_n$ and $C'_0 \cdot \dots \cdot C'_n$ are *equivalent modulo intentions* iff $C'_i = \langle \mathcal{N}_i, \Lambda_i, \Pi_i, \mathcal{B}_i, \mathcal{A}_i, \Gamma'_i \rangle$, for every $0 \leq i \leq n$. Also, the two executions are *equivalent modulo intentions* $P_0 \in \Gamma_0$ and $P'_0 \in \Gamma'_0$ if they are equivalent modulo intentions and for every $0 \leq i \leq n$, $(\Gamma'_i \setminus \{P'_i\}) = (\Gamma_i \setminus \{P_i\})$ (where P_i (P'_i) is P_0 's (P'_0 's) evolution in configuration C_i (C'_i)).

Lastly, we define what we mean by the execution of an intention and by a program (weakly) simulating another program.

DEFINITION 3 (INTENTION EXECUTION). Let E be a BDI execution $C_0 \cdot C_1 \cdot \dots \cdot C_n$ for an agent $C_0 = \langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}_0, \mathcal{A}_0, \Gamma_0 \rangle$, where $\Gamma_0 = \Gamma'_0 \cup \{P_0\}$. Intention P_0 in C_0 has been *fully executed* in E if $P_n = \epsilon$; otherwise P_0 is *currently executing* in E . In addition, intention P_0 in C_0 has been *successfully executed* in E if $P_i = \text{nil}$, for some $i \leq n$; intention P_0 has *failed* in E if it has been fully but not successfully executed in E .

DEFINITION 4 (PROGRAM SIMULATION). Let E be an execution of $C = \langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}, \mathcal{A}, \Gamma \cup \{P\} \rangle$. Program P' *simulates* program P in execution E iff there is an execution E' of configuration $C' = \langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}, \mathcal{A}, \Gamma \cup \{P'\} \rangle$ such that (a) $\bar{E}$ and $\bar{E}'$ are equivalent modulo P and P' ; and (b) if P has been successfully executed in E , so has P' in E' . We say that P' *simulates* P iff P' simulates P in every execution of any configuration.

We have, so far, defined the necessary technical foundations for adding HTN-style planning into the CAN BDI agent language, including substantially polishing and simplifying the original CAN's semantics from [24], incorporating extra representation for actions, and providing the necessary definitions of agent execution that were not addressed in [24]. Let us now move on to the core of the paper.

4. PLANNING IN BDI SYSTEMS

In this section, we shall integrate hierarchical planning into the BDI architecture of section 3. To do so, several issues need to

be addressed. Firstly, we want to keep the language as *uniform* as possible. Secondly, we allow control over when and on what planning is to be performed within the BDI architecture. Thirdly, we need to decide what domain information the planner will use—we want the planner to re-use as much information as possible from an existing BDI specification. Lastly, the result of the planning process ought to be carried on, and possibly monitored, within the BDI execution cycle in a uniform manner.

To address the above issues, we extend the CAN language by introducing a new language construct **Plan** for offline lookahead planning, so that $\text{Plan}(P)$, where P is a plan-body, means “*plan for P offline, searching for a complete hierarchical decomposition.*” In this way, the BDI agent on **Plan** does a full lookahead search before committing to even the first step.

As with other constructs in the language, we need to provide the operational rules for the **Plan** construct. To do this, we shall distinguish, from now on, between two types of transitions on basic configurations, namely, *bdi* and *plan* (labelled) transitions. We write $C \xrightarrow{t} C'$ to specify a single step transition of type t (when no label is stated, both types apply). Intuitively, *bdi*-type steps will be used to model the normal BDI execution cycle, whereas *plan*-type transitions will be used to model (internal) deliberation steps within a *planning* context.

Following [6], the main operational rule states that configuration $\langle \mathcal{B}, \mathcal{A}, \text{Plan}(P) \rangle$ can evolve to $\langle \mathcal{B}', \mathcal{A}', \text{Plan}(P') \rangle$ provided that $\langle \mathcal{B}, \mathcal{A}, P \rangle$ can evolve to $\langle \mathcal{B}', \mathcal{A}', P' \rangle$ from where it is possible to reach a *final* configuration in a finite number of *planning* steps:

$$\frac{\langle \mathcal{B}, \mathcal{A}, P \rangle \xrightarrow{\text{plan}} \langle \mathcal{B}', \mathcal{A}', P' \rangle \quad \langle \mathcal{B}', \mathcal{A}', P' \rangle \xrightarrow{\text{plan}^*} \langle \mathcal{B}'', \mathcal{A}'', \text{nil} \rangle}{\langle \mathcal{B}, \mathcal{A}, \text{Plan}(P) \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}', \mathcal{A}', \text{Plan}(P') \rangle} \text{Plan}$$

There are also three extra simpler rules associated with construct **Plan** that are shown in Figure 1. Rule Plan_f handles the case where no planning solution can be found; rule Plan_t deals with the trivial case of planning on program *nil*; and, lastly, Plan_p handles the **Plan** construct within a planning context.

In addition to these three derivation rules for **Plan**, we need to restrict the two derivation rules $\triangleright_f$ and $\mathbf{G}_R$ from section 3 to the *bdi* context only. This is because failure handling and goal restarting should not be made available during planning—they are features of the BDI execution cycle only. Hence, planning is *not* merely doing lookahead on the BDI execution cycle. We refer to the new versions of the rules as $\triangleright_f^{\text{bdi}}$ and $\mathbf{G}_R^{\text{bdi}}$, respectively. Also, since we now have two types of transition for basic configurations, we need to slightly modify the top-level agent rules A_{step} and A_{clean} to be defined in terms of *bdi*-type transitions. We only show here rules $\triangleright_f^{\text{bdi}}$ and A_{step} (rules $\mathbf{G}_R^{\text{bdi}}$ and A_{clean} should be obvious):

$$\frac{\langle \mathcal{B}, \mathcal{A}, P_1 \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}, \mathcal{A}, P_2 \rangle}{\langle \mathcal{B}, \mathcal{A}, (P_1 \triangleright P_2) \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}, \mathcal{A}, P_2 \rangle} \triangleright_f^{\text{bdi}}$$

$$\frac{P \in \Gamma \quad \langle \mathcal{B}, \mathcal{A}, P \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}', \mathcal{A}', P' \rangle}{\langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}, \mathcal{A}, \Gamma \rangle \Rightarrow \langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}', \mathcal{A}', (\Gamma \setminus \{P\}) \cup \{P'\} \rangle} A_{\text{step}}$$

Observe that, with the alternative rule $\triangleright_f^{\text{bdi}}$, only the BDI execution cycle would be allowed to re-try alternative plans for an event upon the failure of some failed alternative. Indeed, a program of the form $(?false \triangleright (\Delta))$ has no transition within a *plan* context, whereas program (Δ) would be tried within a *bdi* context.

In [24], it was required that the success and failure conditions in a goal-program be mutually exclusive. There is also another sensible restriction on goal-programs, namely, that the program P provided as a method for achieving a (declarative) goal ϕ_s does not make the failure condition ϕ_f true by itself.

DEFINITION 5. A goal-program $\text{Goal}(\phi_s, P, \phi_f)$ is *coherent* (relative to a plan library and an action library) if for every belief bases $\mathcal{B}, \mathcal{B}', \mathcal{B}''$ and sequences of actions $\mathcal{A}, \mathcal{A}', \mathcal{A}''$ such that $\langle \mathcal{B}, \mathcal{A}, P \rangle \xrightarrow{\text{bdi}^*} \langle \mathcal{B}', \mathcal{A}', P' \rangle \xrightarrow{\text{bdi}^*} \langle \mathcal{B}'', \mathcal{A}'', \text{nil} \rangle$, it is the case that $\mathcal{B}' \not\models \phi_f$. An agent is *coherent* if every goal-program mentioned in its plan library is coherent.

From now on, we assume that agents are *coherent*—only the environment or other concurrent intentions may make the failure condition of a goal-program true.⁴ As expected, if the agent’s only intention $\text{Plan}(P)$ is able to start executing, then there is at least one full successful BDI execution for such intention, provided there is no intervention from the outside environment. Equally important, under the same provisions, no execution of the agent will end up failing the intention.

THEOREM 1. Let $C = \langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}, \mathcal{A}, \{\text{Plan}(P)\} \rangle$ such that $\langle \mathcal{B}, \mathcal{A}, \text{Plan}(P) \rangle \xrightarrow{\text{bdi}}$. If E is an environment-free agent execution of C , then intention $\text{Plan}(P)$ is either executing or has been successfully executed in E . Moreover, there is an execution E^s of C in which intention $\text{Plan}(P)$ has been successfully executed in E^s .

PROOF. This relies on the following lemma: if $\langle \mathcal{B}, \mathcal{A}, P \rangle \xrightarrow{\text{plan}^*} \langle \mathcal{B}_f, \mathcal{A}_f, \text{nil} \rangle$, then $\langle \mathcal{B}, \mathcal{A}, \text{Plan}(P) \rangle \xrightarrow{\text{bdi}^*} \langle \mathcal{B}_f, \mathcal{A}_f, \text{nil} \rangle$.

On the contrary, suppose there is an environment-free execution E of the form $C_0 = C \cdot \dots \cdot C_k$ such $\langle \mathcal{B}_k, \mathcal{A}_k, \text{Plan}(P_k) \rangle \not\xrightarrow{\text{bdi}}$. Observe, though, that $\langle \mathcal{B}, \mathcal{A}, \text{Plan}(P) \rangle \xrightarrow{\text{bdi}_k} \langle \mathcal{B}_k, \mathcal{A}_k, \text{Plan}(P_k) \rangle$. By the rule **Plan**, $\langle \mathcal{B}, \mathcal{A}, P \rangle \xrightarrow{\text{plan}_k} \langle \mathcal{B}_k, \mathcal{A}_k, P_k \rangle \xrightarrow{\text{plan}^*} \langle \mathcal{B}_f, \mathcal{A}_f, \text{nil} \rangle$ and $\langle \mathcal{B}, \mathcal{A}, P \rangle \xrightarrow{\text{plan}^*} \langle \mathcal{B}_f, \mathcal{A}_f, \text{nil} \rangle$ applies. By using the above lemma, we get that $\langle \mathcal{B}, \mathcal{A}, \text{Plan}(P) \rangle \xrightarrow{\text{bdi}^*} \langle \mathcal{B}_f, \mathcal{A}_f, \text{nil} \rangle$. Next, since $\langle \mathcal{B}, \mathcal{A}, \text{Plan}(P) \rangle \not\xrightarrow{\text{bdi}} \langle \mathcal{B}_f, \mathcal{A}_f, \text{nil} \rangle$, there exist $\mathcal{B}'', \mathcal{A}'', P''$ such that $\langle \mathcal{B}, \mathcal{A}, \text{Plan}(P) \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}'', \mathcal{A}'', P'' \rangle \xrightarrow{\text{bdi}^*} \langle \mathcal{B}_f, \mathcal{A}_f, \text{nil} \rangle$. Thus, $\langle \mathcal{B}, \mathcal{A}, \text{Plan}(P) \rangle \xrightarrow{\text{bdi}}$ and the above E cannot exist.

The second part follows easily from the fact that $\xrightarrow{\text{plan}^*}$ stands for a finite chain of transitions: if the agent follows those exact transitions, P will eventually terminate successfully. $\square$

Thus, by using the new lookahead construct $\text{Plan}(P)$, the programmer can make sure—to some extent—that *failing* executions of program P will be avoided. This contrasts with the usual (default) BDI execution of P which may potentially *fail* program P due to wrong decisions at choice points. Nonetheless, it should be clear that the proposed deliberation module is *local* in the sense that it does not take into account the potential interactions with the external environment and other concurrent intentions.

Let us now focus on the relationship between our planning construct **Plan** and existing HTN planners. To that end, we say that a CANPLAN agent is a *bounded agent* if its belief base and all belief conditions are defined in a language which follows the same constraints as those imposed by HTN planners [11] (e.g., first-order atoms, finite domains, close world assumption). It is worth pointing out that, in practice, most existing BDI programming language implementations do actualise such constraints and deal only with bounded agents. We also assume, without loss of generality, that bounded agents do not make use of $+b$ and $-b$ statements in their plans—only primitive actions can change the belief base. ($+b$ and $-b$ statements can always be represented via special BDI actions.)

⁴This definition is a bit too strong in that it requires a goal-program to be “sound” w.r.t. the failure condition for *every* possible belief base and *every* chain of *bdi*-transitions, including failed recovered executions. Even though a weaker version could be obtained with a more involved definition, we stick, for simplicity, to the above one.

$$\begin{array}{c}
\frac{\langle \mathcal{B}, \mathcal{A}, (nil \triangleright P') \rangle \longrightarrow \langle \mathcal{B}, \mathcal{A}, nil \rangle \triangleright_t \quad \langle \mathcal{B}, \mathcal{A}, +b \rangle \longrightarrow \langle \mathcal{B} \cup \{b\}, \mathcal{A}, nil \rangle +b \quad \langle \mathcal{B}, \mathcal{A}, -b \rangle \longrightarrow \langle \mathcal{B} \setminus \{b\}, \mathcal{A}, nil \rangle -b}{\langle \mathcal{B}, \mathcal{A}, P_1 \rangle \longrightarrow \langle \mathcal{B}', \mathcal{A}', P' \rangle} \triangleright \\
\frac{\langle \mathcal{B}, \mathcal{A}, (P_1 \triangleright P_2) \rangle \longrightarrow \langle \mathcal{B}', \mathcal{A}', (P' \triangleright P_2) \rangle \triangleright \quad \langle \mathcal{B}, \mathcal{A}, (nil ; P) \rangle \longrightarrow \langle \mathcal{B}, \mathcal{A}, P \rangle Seq_t \quad \langle \mathcal{B}, \mathcal{A}, (P \parallel nil) \rangle \longrightarrow \langle \mathcal{B}, \mathcal{A}, P \rangle \parallel_{t_2}}{\langle \mathcal{B}, \mathcal{A}, P_1 \rangle \longrightarrow \langle \mathcal{B}', \mathcal{A}', P' \rangle} \triangleright \\
\frac{\langle \mathcal{B}, \mathcal{A}, (P_1 \parallel P_2) \rangle \longrightarrow \langle \mathcal{B}', \mathcal{A}', (P' \parallel P_2) \rangle \parallel_1 \quad \langle \mathcal{B}, \mathcal{A}, (P_1 \parallel P_2) \rangle \longrightarrow \langle \mathcal{B}', \mathcal{A}', (P_1 \parallel P') \rangle \parallel_2 \quad \langle \mathcal{B}, \mathcal{A}, (nil \parallel P) \rangle \longrightarrow \langle \mathcal{B}, \mathcal{A}, P \rangle \parallel_{t_1}}{\langle \mathcal{B}, \mathcal{A}, Plan(nil) \rangle \longrightarrow \langle \mathcal{B}, \mathcal{A}, nil \rangle Plan_t} \\
\frac{\langle \mathcal{B}, \mathcal{A}, P \rangle \xrightarrow{plan} \langle \mathcal{B}', \mathcal{A}', P' \rangle}{\langle \mathcal{B}, \mathcal{A}, Plan(P) \rangle \xrightarrow{plan} \langle \mathcal{B}', \mathcal{A}', Plan(P') \rangle} Plan_p
\end{array}$$

$$\Delta = \{A_{step}, A_{clean}, A_{event}\} \cup \{Event, Sel, +b, -b, act, ?, Seq, Seq_t, \triangleright, \triangleright_t, \triangleright_f^{bdi}, \parallel_1, \parallel_2, \parallel_{t_1}, \parallel_{t_2}, G_I, G_s, G_S, G_R^{bdi}, Plan, Plan_t, Plan_p\}.$$

Figure 1: CANPLAN's complete set of rules Δ is built from the rules described in the text plus the ones shown here.

The next theorem establishes, formally, the link between the Plan construct and HTN planning. First, we prove that the new construct Plan could indeed be seen as an HTN planner. Second, we show that executions of program Plan(P) encode HTN plan solutions. Lastly, and not so surprisingly, we demonstrate that a straight-line HTN plan solution could be successfully executed by the BDI execution cycle. For clarity, we keep the translation between the BDI domain knowledge (i.e., libraries Π and Λ , and program) and the HTN procedural knowledge (i.e., planning domain, "task network" P) implicit. (the theorem's proof is based on the relationship between the BDI's and HTN's entities as discussed in section 2.3.)

THEOREM 2. *For any bounded agent,*

1. $\langle \mathcal{B}, \mathcal{A}, Plan(P) \rangle \xrightarrow{bdi} \text{iff } sol(P, \mathcal{B}, \Pi \cup \Lambda) \neq \emptyset.$
2. $\langle \mathcal{B}, \mathcal{A}, Plan(P) \rangle \xrightarrow{bdi_*} \langle \mathcal{B}', \mathcal{A} \cdot act_1 \cdot \dots \cdot act_k, Plan(P') \rangle$
with $k \geq 1$ iff there exists a plan $\sigma \in sol(P, \mathcal{B}, \Pi \cup \Lambda)$, such that $\sigma = act_1 \cdot \dots \cdot act_k \cdot \dots \cdot act_n$, for some $n \geq k$.
3. If there exists a plan $\sigma = act_1 \cdot \dots \cdot act_n \in sol(P, \mathcal{B}, \Pi \cup \Lambda)$, then $\langle \mathcal{B}, \mathcal{A}, (act_1; \dots; act_n) \rangle \xrightarrow{bdi_*} \langle \mathcal{B}', \mathcal{A} \cdot \sigma, nil \rangle.$

Therefore, provided we restrict to the language of HTN [11], our deliberator construct Plan provides a built-in HTN planner within the whole BDI framework. The above theorem is an important practical result as it gives us the rationale for using existing HTN planner systems, such as SHOP [15] and SHOP2 [16], within current BDI implementations (e.g., AGENTSPEAK [19], JACK [3]).

4.1 Planning for Declarative Goals

So far we have seen how lookahead planning can be done on (procedural) programs. Let us now discuss how (classical) planning for a declarative goal ϕ_s using a procedural program P can be done. There are a few choices for this and the following five properties that we may be interested in satisfying:

- (A) P is used towards the eventual satisfaction of goal ϕ_s .
- (B) P may execute partially if goal ϕ_s is achieved before P completion. That is, P need not be executed completely.
- (C) There is a commitment to the goal ϕ_s so that P is reinstantiated and retried until the goal in question is established.
- (D) There exists a mechanism for dropping the goal when a failure condition ϕ_f becomes true.
- (E) At planning time, P is solved up to the point where the goal is met. That is, it may not be required to solve P completely.

The different alternatives that we shall consider together with the properties satisfied by each one are described in the following table:

ALTERNATIVES	A	B	C	D	E
Plan($P; ?\phi_s$)	✓				
Plan(Goal(ϕ_s, P, ϕ_f))	✓	✓		✓	
Goal($\phi_s, Plan(P), \phi_f$)		✓	✓	✓	
Goal($\phi_s, Plan(P; ?\phi_s), \phi_f$)	✓	✓	✓	✓	
Goal($\phi_s, Plan(Goal(\phi_s, P, \phi_f)), \phi_f$)	✓	✓	✓	✓	✓

Interestingly, a *first-principles* account of planning can easily be obtained by using the first alternative Plan($P; ?\phi$) by taking $P = !seqActions$ the special event *seqActions* can be solved with any sequence of primitive actions.

Notice that the last four alternatives make use of the the special Goal construct available in CANPLAN to handle declarative goals within the BDI execution cycle. Observe also that the last option is the only one satisfying all five properties combining then the advantages from the BDI execution cycle and the planning module. Consequently, it is sensible to define a new language construct Plan(ϕ_s, P, ϕ_f) in the following way:

$$Plan(\phi_s, P, \phi_f) \stackrel{\text{def}}{=} Goal(\phi_s, Plan(Goal(\phi_s, P, \phi_f)), \phi_f).$$

Among other results, it can be shown that Plan(ϕ_s, P, ϕ_f) subsumes all the executions of Plan(Goal(ϕ_s, P, ϕ_f)).

THEOREM 3. *For every ϕ_s, ϕ_f and P , program Plan(ϕ_s, P, ϕ_f) simulates program Plan(Goal(ϕ_s, P, ϕ_f)).*

To recap: combinations of the Plan and Goal constructs suggest an interesting range of programs for declarative goals. We believe that Plan(ϕ_s, P, ϕ_f) provides a convenient mechanism for dealing with declarative goals at both planning and execution time.

4.2 Planning vs BDI Execution

We conclude this section by exploring the differences between the execution of a planning program and the normal BDI execution. A CANPLAN⁻ agent is a CANPLAN agent whose plan language does not include the $\parallel$ and Goal constructs. This restriction corresponds to classical BDI agent programming languages like AGENTSPEAK and to *total-order* HTN planners like SHOP; neither system include concurrency and goals natively. Under such restricted CANPLAN agents, the planning module is no more than a lookahead mechanism on top of the BDI execution cycle.

THEOREM 4. *Program P simulates program Plan(P) in every CANPLAN⁻ agent.*

On the other hand, when concurrency or goal-programs are considered, performing planning may result in extra executions. In fact, it can be shown that executing Plan(Plan(P_1) $\parallel$ P₂) is equivalent to executing Plan(P_1 $\parallel$ P₂), which in turn, is very different

from executing $\text{Plan}(P_1 \parallel P_2)$.⁵ A similar situation arises with program $\text{Plan}(\text{Goal}(\phi_s, \text{Plan}(P), \phi_f))$. The reason, technically, is that a **Plan** construct is ignored within the context of another **Plan** construct—there is no notion of planning within planning.

Surprisingly, also, the BDI execution engine may obtain successful executions that the planner cannot produce.

THEOREM 5. *There exists an agent configuration C of the form $\langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}, \mathcal{A}, \Gamma \cup \{P\} \rangle$ for which there is an execution where P is successfully executed, but such that no execution of $C' = \langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}, \mathcal{A}, \Gamma \cup \{\text{Plan}(P)\} \rangle$ can successfully execute $\text{Plan}(P)$.*

PROOF. *Let us build a counter-example. Suppose that all actions are possible and that action act_1 just makes p true, that p and q are both false initially, and that there are only two plan rules in the plan library Π for handling event e : (i) $e : \text{true} \leftarrow act_1; ?q; act_2$; and (ii) $e : p \leftarrow act_3; act_2$. There is no solution for $\text{Plan}(!e)$, but a BDI execution that would successfully execute $!e$ can be obtained by partially executing plan rule (i) (action act_1) and then, upon failure, fully executing plan rule (ii). $\square$*

As one can observe, the proof’s counter-example relies on both the plan failure handling mechanism built into the BDI execution cycle and the programmer not having provided a full set of plans. In fact, if the plan library in the above proof’s counter-example had included a third rule of the form $e : \text{true} \leftarrow act_1; ?p; act_3; act_2$, then the planner would have found a full execution. Still, as agent programs are often developed incrementally and in modules, the above situation could well arise.

It follows then that the combined framework of (default) BDI execution plus local hierarchical planning is strictly more general than hierarchical planning alone. Furthermore, as discussed after Theorem 1, by using the new local planning mechanism the programmer can rule out BDI executions that are bound to fail.

5. IMPLEMENTATION ISSUES

In earlier work [8], we presented an implementation that combined BDI reasoning with HTN planning. We used JACK⁶ BDI system and JSHOP⁷ HTN planner, a Java version of SHOP [15]. Although the integrated framework does not fully realise the operational semantics presented here, it does incorporate some important concepts from it. In particular, it allows the programmer to specify from within a JACK program the points at which JSHOP should be called, in a manner similar to the **Plan** construct. Consistent with the semantics of **Plan**, JSHOP uses the same domain representation as JACK does (i.e., the plan library Π and belief base $\mathcal{B}$). In fact, the framework builds at runtime a JSHOP planning problem representation automatically from the JACK domain knowledge.

Some differences in the implementation arise from the nature of the systems chosen for the implementation. Since JSHOP is a total-order HTN planner, it does not use the $\parallel$ construct defined in P . However, since parallelism has benefits, the integrated framework converts JSHOP’s total-order solutions into partial-order solutions so that JACK can exploit possible parallelism at execution time. Some other differences exist between the implementation and semantics for the sake of simplicity. For example, we exclude the $\text{Goal}(\phi_s, P, \phi_f)$ construct in our system, as this construct does not have a direct matching concept in JACK or JSHOP. Including this

⁵A framework where $\text{Plan}(\text{Plan}(P_1 \parallel P_2))$ is not equivalent to $\text{Plan}(P_1 \parallel P_2)$ would require an account of HTN planning within an HTN planner. This framework can be obtained by dropping rule **Plan_p** and making rule **Plan** also available within the plan context.

⁶www.agent-software.com.au

⁷www.cs.umd.edu/projects/shop/description.html

goal construct and using SHOP2 [16] to accommodate parallel execution of sub-goals natively are left for future work.

The main difference, however, is that the implementation *does not re-plan* at every step, as indicated by the **Plan** rule defined in the semantics. This would clearly be unnecessarily inefficient. Instead, JSHOP was modified to return the relevant methods and bindings (rather than simply the actions); the BDI execution engine then follows step-by-step the decomposition suggested by the planner. Relevant environmental changes are detected by virtue of a step in the returned plan no longer being applicable within the BDI cycle. At that point, the planner is then called once again to provide an updated plan, and if none is available failure will occur in the BDI system. A disadvantage of this is that environmental changes leading to failure may be detected later in the implemented version than in the semantic rules. However, this drawback is offset by the much greater efficiency in what can be expected to be the majority of cases. This approach also has the benefit that an intention produced by a call to **Plan** will, in fact, *terminate*—successfully if there is no environmental interference. This is stronger than what Theorem 1 states, in which we needed to account for the strange, but potentially possible, situation where the **Plan** module continually returns a new and different plan prior to termination.

6. RELATED WORK

Except for INDIGOLOG [6], which is not *per se* a typical BDI agent programming language (see below), we are not aware of any other *formal* BDI-style agent programming language (e.g., AGENTSPEAK [19], 3APL [12], PRS [13]) providing a clean account of planning as we do here with CANPLAN. There are however a number of (implemented) *systems* or *frameworks* which do, in some way or another, mix planning and BDI-style execution. Some of these are *planners*, such as IPeM [1] and SAGE [14], that allow for the interleaving of action execution during the planning process. Others are *agent architectures*, such as RETSINA [17], CYPRESS and CPEF [23], and PROPICE-PLAN [10]; they are able to do lookahead planning. CYPRESS is based on the ACT [22] formalism that provides a uniform representation framework for BDI execution systems and hierarchical planners, hence supporting the type of mapping we have proposed in section 2.2. PROPICE-PLAN is perhaps the most similar system to ours, in that it is a typical BDI agent system that is able to call a planning module to find a solution for a particular problem. Like CANPLAN, a unified representation is used by both the planner and the BDI system. The work done in this paper differs, at least, in two ways from the above systems. First of all, we are particularly concerned with the *formal specification* of a BDI agent with built-in planning capabilities as well as with the formal relation between BDI systems and HTN planners. To our knowledge, none of the above systems come with a precise formal semantics. In some sense, however, our work was much motivated by the existence of these systems in an analogous way to how AGENTSPEAK [19] was motivated by systems like PRS [13]. Secondly, CANPLAN provides a mechanism for local deliberation *on-demand* that the programmer can use, as opposed to a fixed integration of planning within the execution engine (e.g., planning always [17] or just on (goal) failure [10, 23]).

Our work is possibly most related to that of De Giacomo and Levesque [6] in which INDIGOLOG, an incremental version of CONGOLOG [5] with a local deliberation module Σ , is proposed in the context of the situation calculus. Several ideas are taken from that work and applied to the BDI context. Our work is however different in that (i) INDIGOLOG is a cognitive agent language, with no explicit notions of events, plan library, plan selection, failure handling, intention base, etc., whereas our approach is linked to a

whole family of typical BDI languages and systems; (ii) our planning mechanism is provably linked to a well understood approach in the planning community, namely HTN planning, whereas, as far as we know, the INDIGOLOG deliberator module is very general and is not directly related to any planning system; (iii) the integration of the planning module with the notion of declarative goals in CANPLAN has no counterpart in INDIGOLOG. In some ways, our approach has a more practical orientation than that of INDIGOLOG. It would be interesting to investigate the relations between INDIGOLOG and CANPLAN (e.g., identify the BDI subclass of agents that could be written and executed in INDIGOLOG).

7. CONCLUSION AND FUTURE WORK

We have proposed a mechanism for planning within BDI systems based on the intrinsic requirements of the BDI architecture. To do so, we provided an operational semantics that substantially simplifies and extends that presented in [24] to incorporate a new *planning* construct **Plan**. The new construct offers power and flexibility to the BDI programmer for specifying lookahead planning points in programs. We described results showing that the integration between the planning module and the whole BDI execution is the one intuitively expected, and proved that, under suitable assumptions, the planning task reduces to HTN planning. Lastly, we showed that the combined system allows a larger set of “good” executions than the planning module alone and discussed an implementation that incorporates many of the concepts from the semantics. We believe the work presented here is a significant step towards incorporating lookahead deliberation into BDI-type agent systems in a principled manner. More importantly, it provides a firm foundation for a range of interesting further work.

The fact that we have chosen to provide planning via a new construct is very much in the spirit of BDI systems, namely, allowing for direct encoding of programmer or domain expertise. In this case, the knowledge about when planning would be beneficial. It may be argued, though, that an intelligent agent should (also) make its own decisions as to when to plan. It is therefore worth investigating a more general account in which the agent could itself take the initiative to plan; for example, when all plans for a goal fail or when there is substantial spare time. *Re-planning* following failure of a plan produced by the planner module is also a topic we have not explored here and which deserves further work (see [23]).

We have already started exploring how to accommodate extensions to classical HTN planning within our formal framework. For instance, decoupling the hierarchical structure of BDI plans and using a planning account more akin to first principles would allow for potential discovery of new plan structures. This, in turn, could provide the basis for the agent to “learn” new plans that could be added to the plan library. Also, it can be useful to plan only to a certain level of abstraction or detail, leaving further remaining decompositions to execution time or until absolutely necessary as done in [4]. Both above generalisations are likely to require (or benefit from) extra representation of effects for high-level plans. Such extra representation would also provide support for reasoning about interactions of the plan being explored with other goals and intentions of the agent [4, 20]. In particular, we would like to extend our current *local* lookahead mechanism so that the agent considers, at least, all of its own active intentions when performing planning.

The framework presented here provides a basis for exploring the interaction between declarative goals [24, 21]—preliminary results were given in section 4.1 but further investigation is needed.

Lastly, it would be interesting to develop *resource-bounded* version of our planning module. To that end, we are considering developing an *anytime* or *incremental* version of the **Plan** construct.

8. REFERENCES

- [1] J. Ambros-Ingerson. *IPEM: Integrated Planning, Execution, and Monitoring*. PhD thesis, Dept. of Computer Science, University of Essex, U.K., 1987.
- [2] M. Bratman. *Intentions, Plans, and Practical Reason*. Harvard University Press, 1987.
- [3] P. Busetta, R. Rönquist, A. Hodgson, and A. Lucas. JACK Intelligent Agents - Components for Intelligent Agents in Java, AgentLink News Letter, Agent Oriented Software Pty. Ltd., Melbourne, January 1999.
- [4] B. J. Clement and E. H. Durfee. Theory for Coordinating Concurrent Hierarchical Planning Agents Using Summary Information. In *Proc. of AAAI-99*, pages 495–502, 1999.
- [5] G. De Giacomo, Y. Lespérance, and H. Levesque. ConGolog, A Concurrent Programming Language Based on the Situation Calculus. *Artificial Intelligence*, 121(1–2):109–169, 2000.
- [6] G. De Giacomo and H. Levesque. An Incremental Interpreter for High-Level Programs with Sensing. In H. Levesque and F. Pirri, editors, *Logical Foundation for Cognitive Agents: contr. in honor of Ray Reiter*, pages 86–102. Springer, 1999.
- [7] L. P. de Silva and L. Padgham. A Comparison of BDI Based Real-Time Reasoning and HTN Based Planning. In *Proc. of Australian Joint Conference on AI*, pages 1167–1173, 2004.
- [8] L. P. de Silva and L. Padgham. Planning on Demand in BDI Systems. In *Proc. of ICAPS-05 (Poster)*, 2005.
- [9] D. Dennett. *The Intentional Stance*. MIT Press, 1987.
- [10] O. Despouys and F. F. Ingrand. Propice-Plan: Toward a Unified Framework for Planning and Execution. In *Proc. of European Conference on Planning*, pages 278–293, 1999.
- [11] K. Erol, J. Hendler, and D. S. Nau. HTN Planning: Complexity and Expressivity. In *Proc. of AAAI-94*, pages 1123–1228, 1994.
- [12] K. V. Hindriks, F. S. de Boer, W. van der Hoek, and J.-J. Ch. Meyer. Agent Programming in 3APL. *Autonomous Agents and Multi-Agent Systems*, 2(4):357–401, 1999.
- [13] F. F. Ingrand, M. P. Georgeff, and A. S. Rao. An Architecture for Real-Time Reasoning and System Control. *IEEE Expert: Intelligent Systems and Their Applications*, 7(6):34–44, 1992.
- [14] C. A. Knoblock. Planning, Executing, Sensing, and Replanning for Information Gathering. In *Proc. of IJCAI-95*, pages 1686–1693, 1995.
- [15] D. S. Nau, Y. Cao, A. Lotem, and H. Muñoz-Avila. SHOP: Simple Hierarchical Ordered Planner. In *Proc. of IJCAI-99*, pages 968–973, 1999.
- [16] D. S. Nau, H. Muñoz-Avila, Y. Cao, A. Lotem, and S. Mitchell. Total-Order Planning with Partially Ordered Subtasks. In *Proc. of IJCAI-01*, pages 425–430, 2001.
- [17] P. Paolucci, O. Shehory, K. P. Sycara, K. P. Kalp, and A. Pannu. A Planning Component for RETSINA Agents. In *Proc. of ATAL-99*, pages 147–161, 1999.
- [18] G. Plotkin. A Structural Approach to Operational Semantics. Technical Report DAIMI-FN-19, Dept. of Computer Science Department, Aarhus University, Denmark, 1981.
- [19] A. S. Rao. AgentSpeak(L): BDI Agents Speak Out in a Logical Computable Language. In W. V. Velde and J. W. Perram, editors, *Agents Breaking Away (LNAI)*, volume 1038 of *LNAI*, pages 42–55. Springer-Verlag, 1996.
- [20] J. Thangarajah, L. Padgham, and M. Winikoff. Detecting & Exploiting Positive Goal Interaction in Intelligent Agents. In *Proc. of AAMAS-03*, pages 401–408, 2003.
- [21] M. B. van Riemsdijk, M. Dastani, and J.-J. Ch. Meyer. Semantics of Declarative Goals in Agent Programming. In *Proc. of AAMAS-05*, pages 133–140, 2005.
- [22] D. E. Wilkins and K. L. Myers. A Common Knowledge Representation for Plan Generation and Reactive Execution. *Journal of Logic and Computation*, 5(6):731–761, 1995.
- [23] D. E. Wilkins and K. L. Myers. A Multiagent Planning Architecture. In *Proc. of AIPS-98*, pages 154–162, 1998.
- [24] M. Winikoff, L. Padgham, J. Harland, and J. Thangarajah. Declarative & Procedural Goals in Intelligent Agent Systems. In *Proc. of KR-02*, pages 470–481, 2002.

A. COMPLETE OPERATIONAL SEMANTICS FOR CANPLAN

$$\begin{array}{c}
\frac{\langle \mathcal{B}, \mathcal{A}, P \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}'', \mathcal{A}', P' \rangle}{\langle \mathcal{N}, \Pi, \Lambda, \mathcal{B}, \mathcal{A}, \Gamma \rangle \Rightarrow \langle \mathcal{N}, \Pi, \Lambda, \mathcal{B}'', \mathcal{A}', (\Gamma \setminus \{P\}) \cup \{P'\} \rangle} A_{step} \quad \frac{e \text{ is a new external event}}{\langle \mathcal{N}, \Pi, \Lambda, \mathcal{B}, \mathcal{A}, \Gamma \rangle \Rightarrow \langle \mathcal{N}, \Pi, \Lambda, \mathcal{B}, \mathcal{A}, \Gamma \cup \{!e\} \rangle} A_{event} \\
\\
\frac{P \in \Gamma \quad P = \text{nil}}{\langle \mathcal{N}, \Pi, \Lambda, \mathcal{B}, \mathcal{A}, \Gamma \rangle \Rightarrow \langle \mathcal{N}, \Pi, \Lambda, \mathcal{B}, \mathcal{A}, \Gamma \setminus \{P\} \rangle} A_{clean1} \quad \frac{P \in \Gamma \quad \langle \mathcal{B}, \mathcal{A}, P \rangle \xrightarrow{\text{bdi}}}{\langle \mathcal{N}, \Pi, \Lambda, \mathcal{B}, \mathcal{A}, \Gamma \rangle \Rightarrow \langle \mathcal{N}, \Pi, \Lambda, \mathcal{B}, \mathcal{A}, \Gamma \setminus \{P\} \rangle} A_{clean2} \\
\\
\hline
\frac{\Delta = \{\psi_i \theta : P_i \theta \mid e' : \psi_i \leftarrow P_i \in \Pi \wedge \theta = \text{mgu}(e, e')\}}{\langle \mathcal{B}, \mathcal{A}, !e \rangle \longrightarrow \langle \mathcal{B}, \mathcal{A}, \langle \Delta \rangle \rangle} Event \quad \frac{\psi_i : P_i \in \Delta \quad \mathcal{B} \models \psi_i \theta}{\langle \mathcal{B}, \mathcal{A}, \langle \Delta \rangle \rangle \longrightarrow \langle \mathcal{B}, \mathcal{A}, P_i \theta \triangleright \langle \Delta \setminus P_i \rangle \rangle} Sel \\
\\
\frac{}{\langle \mathcal{B}, \mathcal{A}, +b \rangle \longrightarrow \langle \mathcal{B} \cup \{b\}, \mathcal{A}, \text{nil} \rangle} +b \quad \frac{}{\langle \mathcal{B}, \mathcal{A}, -b \rangle \longrightarrow \langle \mathcal{B} \setminus \{b\}, \mathcal{A}, \text{nil} \rangle} -b \\
\\
\frac{a : \psi \leftarrow \Phi^+; \Phi^- \in \Lambda \quad a\theta = \text{act} \quad \mathcal{B} \models \psi\theta}{\langle \mathcal{B}, \mathcal{A}, \text{act} \rangle \longrightarrow \langle (\mathcal{B} \setminus \Phi^- \theta) \cup \Phi^+ \theta, \mathcal{A} \cdot \text{act}, \text{nil} \rangle} act \quad \frac{\mathcal{B} \models \phi\theta}{\langle \mathcal{B}, \mathcal{A}, ?\phi \rangle \longrightarrow \langle \mathcal{B}, \mathcal{A}, \text{nil} \rangle} ? \\
\\
\frac{\langle \mathcal{B}, \mathcal{A}, P_1 \rangle \longrightarrow \langle \mathcal{B}'', \mathcal{A}', P' \rangle}{\langle \mathcal{B}, \mathcal{A}, (P_1; P_2) \rangle \longrightarrow \langle \mathcal{B}'', \mathcal{A}', (P'; P_2) \rangle} Seq \quad \frac{}{\langle \mathcal{B}, \mathcal{A}, (\text{nil}; P) \rangle \longrightarrow \langle \mathcal{B}, \mathcal{A}, P \rangle} Seq_t \\
\\
\frac{\langle \mathcal{B}, \mathcal{A}, P_1 \rangle \longrightarrow \langle \mathcal{B}'', \mathcal{A}', P' \rangle}{\langle \mathcal{B}, \mathcal{A}, (P_1 \parallel P_2) \rangle \longrightarrow \langle \mathcal{B}'', \mathcal{A}', (P' \parallel P_2) \rangle} \parallel_1 \quad \frac{\langle \mathcal{B}, \mathcal{A}, P_2 \rangle \longrightarrow \langle \mathcal{B}'', \mathcal{A}', P' \rangle}{\langle \mathcal{B}, \mathcal{A}, (P_1 \parallel P_2) \rangle \longrightarrow \langle \mathcal{B}'', \mathcal{A}', (P_1 \parallel P') \rangle} \parallel_2 \\
\\
\frac{}{\langle \mathcal{B}, \mathcal{A}, (\text{nil} \parallel P) \rangle \longrightarrow \langle \mathcal{B}, \mathcal{A}, P \rangle} \parallel_{t1} \quad \frac{}{\langle \mathcal{B}, \mathcal{A}, (P \parallel \text{nil}) \rangle \longrightarrow \langle \mathcal{B}, \mathcal{A}, P \rangle} \parallel_{t2} \\
\\
\frac{\langle \mathcal{B}, \mathcal{A}, P_1 \rangle \longrightarrow \langle \mathcal{B}'', \mathcal{A}', P' \rangle}{\langle \mathcal{B}, \mathcal{A}, (P_1 \triangleright P_2) \rangle \longrightarrow \langle \mathcal{B}'', \mathcal{A}', (P' \triangleright P_2) \rangle} \triangleright \quad \frac{}{\langle \mathcal{B}, \mathcal{A}, (\text{nil} \triangleright P) \rangle \longrightarrow \langle \mathcal{B}, \mathcal{A}, \text{nil} \rangle} \triangleright_t \quad \frac{\langle \mathcal{B}, \mathcal{A}, P_1 \rangle \xrightarrow{\text{bdi}}}{\langle \mathcal{B}, \mathcal{A}, (P_1 \triangleright P_2) \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}, \mathcal{A}, P_2 \rangle} \triangleright_f^{bdi} \\
\\
\frac{P \neq P_1 \triangleright P_2 \quad \mathcal{B} \not\models \phi_s \vee \phi_f}{\langle \mathcal{B}, \mathcal{A}, \text{Goal}(\phi_s, P, \phi_f) \rangle \longrightarrow \langle \mathcal{B}, \mathcal{A}, \text{Goal}(\phi_s, P \triangleright P, \phi_f) \rangle} G_I \\
\\
\frac{\mathcal{B} \models \phi_s}{\langle \mathcal{B}, \mathcal{A}, \text{Goal}(\phi_s, P, \phi_f) \rangle \longrightarrow \langle \mathcal{B}, \mathcal{A}, \text{nil} \rangle} G_s \quad \frac{\mathcal{B} \models \phi_f}{\langle \mathcal{B}, \mathcal{A}, \text{Goal}(\phi_s, P, \phi_f) \rangle \longrightarrow \langle \mathcal{B}, \mathcal{A}, ?\text{false} \rangle} G_f \\
\\
\frac{\mathcal{B} \not\models \phi_s \vee \phi_f \quad \langle \mathcal{B}, \mathcal{A}, P_1 \rangle \longrightarrow \langle \mathcal{B}', \mathcal{A}', P' \rangle}{\langle \mathcal{B}, \mathcal{A}, \text{Goal}(\phi_s, P_1 \triangleright P_2, \phi_f) \rangle \longrightarrow \langle \mathcal{B}', \mathcal{A}', \text{Goal}(\phi_s, P' \triangleright P_2, \phi_f) \rangle} G_S \\
\\
\frac{\mathcal{B} \not\models \phi_s \vee \phi_f \quad \langle \mathcal{B}, \mathcal{A}, P_1 \rangle \xrightarrow{\text{bdi}}}{\langle \mathcal{B}, \mathcal{A}, \text{Goal}(\phi_s, P_1 \triangleright P_2, \phi_f) \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}, \mathcal{A}, \text{Goal}(\phi_s, P_2 \triangleright P_2, \phi_f) \rangle} G_R^{bdi} \\
\\
\frac{\langle \mathcal{B}, \mathcal{A}, P \rangle \xrightarrow{\text{plan}} \langle \mathcal{B}', \mathcal{A}', P' \rangle \quad \langle \mathcal{B}', \mathcal{A}', P' \rangle \xrightarrow{\text{plan}_*} \langle \mathcal{B}'', \mathcal{A}'', \text{nil} \rangle}{\langle \mathcal{B}, \mathcal{A}, \text{Plan}(P) \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}', \mathcal{A}', \text{Plan}(P') \rangle} Plan \\
\\
\frac{\langle \mathcal{B}, \mathcal{A}, P \rangle \xrightarrow{\text{plan}} \langle \mathcal{B}', \mathcal{A}', P' \rangle}{\langle \mathcal{B}, \mathcal{A}, \text{Plan}(P) \rangle \xrightarrow{\text{plan}} \langle \mathcal{B}', \mathcal{A}', \text{Plan}(P') \rangle} Plan_p \quad \frac{}{\langle \mathcal{B}, \mathcal{A}, \text{Plan}(\text{nil}) \rangle \longrightarrow \langle \mathcal{B}, \mathcal{A}, \text{nil} \rangle} Plan_t
\end{array}$$

B. AUXILIARY DEFINITIONS

DEFINITION 6 (REMOVE PLAN TRANSFORMATION $RP(P_1, P_2)$ AND $RPA(P_1)$). Let $RP(P_1, P_2)$ be a relation that holds iff P_2 is exactly like P_1 with $0 \leq m \leq n$ occurrences of the **Plan** construct in P_1 removed, where n is the number of occurrences of the **Plan** constructs in P_1 . Finally, let $RPA(P_1)$ be the program obtained from deleting all **Plan** constructs in P_1 (clearly $RPA(P_1, RPA(P_1))$).

B.1 Execution Simulation

DEFINITION 7 (EXECUTION SIMULATION). Let C and C' be two configurations, probably for different agents. Let E be an execution of C . Configuration C' simulates configuration C in execution E iff there is an execution E' of C' such that (a) $\overline{E}$ and $\overline{E'}$ are equivalent, and (b) if E is terminating, so is E' . Configuration C' simulates C iff C' simulates C in every execution of C .

DEFINITION 8 (DERIVED TRANSITIONS). An internal step transition relation $\rightarrow_i$ is obtained by restricting $\rightarrow$: $R \rightarrow_i R'$ iff $R \rightarrow R'$ and $\mathcal{O}(R) = \mathcal{O}(R')$. The visible transition relation $\rightarrow_v$ is defined by: $R \rightarrow_v R'$ iff there exists R^* such that $R \xrightarrow{*}_i R^*$ and $R^* \rightarrow R'$ but $R^* \not\rightarrow_i R'$ (i.e., $\mathcal{O}(R^*) \neq \mathcal{O}(R')$).

DEFINITION 9 (SIMULATION RELATION). A binary relation Sim between programs is a simulation relation w.r.t. transition $\rightarrow$ iff for every P_1 and P_2 such that $Sim(P_1, P_2)$, it is the case that for every $\mathcal{B}, \mathcal{A}$:

1. if $\langle \mathcal{B}, \mathcal{A}, P_2 \rangle \xrightarrow{*}_i \langle \mathcal{B}', \mathcal{A}', nil \rangle$, then $\langle \mathcal{B}, \mathcal{A}, P_1 \rangle \xrightarrow{*}_i \langle \mathcal{B}', \mathcal{A}', nil \rangle$.
2. if $\langle \mathcal{B}, \mathcal{A}, P_2 \rangle \rightarrow \langle \mathcal{B}', \mathcal{A}', P'_2 \rangle$, there exists R', R'', P'_1 such that (a) $\langle \mathcal{B}, \mathcal{A}, P_1 \rangle \xrightarrow{*}_i R', R' \rightarrow R''$ or $R' = R''$, and $R'' \xrightarrow{*}_i \langle \mathcal{B}', \mathcal{A}', P'_1 \rangle$; and (b) $Sim(P'_1, P'_2)$.

C. AUXILIARY RESULTS

The following Lemma states that htn-type transitions completely ignore the second plan in programs of the form $P_1 \triangleright P_2$.

LEMMA 6 (PROPERTIES OF $\triangleright$ IN plan-TYPE TRANSITIONS). For every $\mathcal{B}, \mathcal{A}$ and P_1, P_2 :

1. $\langle \mathcal{B}, \mathcal{A}, P_1 \triangleright P_2 \rangle \xrightarrow{plan*} \langle \mathcal{B}'', \mathcal{A}', nil \rangle$ iff $\langle \mathcal{B}, \mathcal{A}, P_1 \rangle \xrightarrow{plan*} \langle \mathcal{B}'', \mathcal{A}', nil \rangle$;
2. If $\langle \mathcal{B}, \mathcal{A}, P_1 \rangle \not\xrightarrow{plan}$, then $\langle \mathcal{B}, \mathcal{A}, fail \triangleright P_2 \rangle \not\xrightarrow{plan}$;
3. $\langle \mathcal{B}, \mathcal{A}, Goal(\phi_s, P_1 \triangleright P_2, \phi_f) \rangle \xrightarrow{plan*} \langle \mathcal{B}'', \mathcal{A}', nil \rangle$ iff $\langle \mathcal{B}, \mathcal{A}, Goal(\phi_s, P_1, \phi_f) \rangle \xrightarrow{plan*} \langle \mathcal{B}'', \mathcal{A}', nil \rangle$;
4. If $\langle \mathcal{B}, \mathcal{A}, P_1 \rangle \not\xrightarrow{plan}$ and $\mathcal{B} \not\models \phi_s \vee \phi_f$, then $\langle \mathcal{B}, \mathcal{A}, Goal(\phi_s, P_1 \triangleright P_2, \phi_f) \rangle \not\xrightarrow{plan}$

PROOF. Direct from the fact that rules $\triangleright_f^{bdi}$ and G_R^{bdi} are not available within the plan context. $\square$

LEMMA 7. If $\langle \mathcal{B}, \mathcal{A}, P \rangle \xrightarrow{plan} \langle \mathcal{B}', \mathcal{A}', P' \rangle$ without using the **Plan_P** derivation rule, then $\langle \mathcal{B}, \mathcal{A}, P \rangle \xrightarrow{bdi} \langle \mathcal{B}', \mathcal{A}', P' \rangle$.

PROOF. Direct from the fact that any plan-type derivation rule except for **Plan_P** is also available as a bdi-type. $\square$

LEMMA 8. Let P_1 and $P_2 = RPA(P_1)$. Then, $\langle \mathcal{B}, \mathcal{A}, Plan(P_1) \rangle \xrightarrow{bdi} \langle \mathcal{B}', \mathcal{A}', Plan(P'_1) \rangle$ iff $\mathcal{A}' = \mathcal{A}$ and $P_2 = RPA(P'_1)$ or (ii) $\langle \mathcal{B}, \mathcal{A}, Plan(P_2) \rangle \xrightarrow{bdi} \langle \mathcal{B}', \mathcal{A}', Plan(RPA(P'_1)) \rangle$.

PROOF. By assumption we know that (a) $\langle \mathcal{B}, \mathcal{A}, P_1 \rangle \xrightarrow{plan} \langle \mathcal{B}', \mathcal{A}', P'_1 \rangle$, and (b) $\langle \mathcal{B}', \mathcal{A}', P'_1 \rangle \xrightarrow{plan} \langle \mathcal{B}_f, \mathcal{A}_f, nil \rangle$. Let us consider the following two cases on (a). If transition (a) is supported by derivation rule **Plan_t**, then clearly $\mathcal{B}' = \mathcal{B}$, $\mathcal{A}' = \mathcal{A}$ and $RP(P_1, P'_1)$ (that is, some subprogram **Plan**(nil) in P_1 was reduced to just nil). In this case, we it is obvious that $P_2 = RPA(P'_1)$ and (i) applies. Now suppose transition (a) is not supported by rule **Plan_t** and let us take $P'_2 = RPA(P'_1)$. Because any **Plan** construct in P_1 is completely ignored due to rule **Plan_P**, it follows that $\langle \mathcal{B}, \mathcal{A}, Plan(P_2) \rangle \xrightarrow{bdi} \langle \mathcal{B}', \mathcal{A}', Plan(P'_2) \rangle$ $\square$

C.1 Properties of Plan

The following Lemma states that the planning module on a program of the form $P_1 \triangleright P_2$ is fully committed to the first program P_1 only.

LEMMA 9 ($\triangleright$ IN Plan). For every $\mathcal{B}, \mathcal{A}, P_1, P_2$: $\langle \mathcal{B}, \mathcal{A}, Plan(P_1 \triangleright P_2) \rangle \xrightarrow{bdi*} \langle \mathcal{B}'', \mathcal{A}', nil \rangle$ iff $\langle \mathcal{B}, \mathcal{A}, Plan(P_1) \rangle \xrightarrow{bdi*} \langle \mathcal{B}'', \mathcal{A}', nil \rangle$.

PROOF. This is a direct consequence of point (1) in Lemma 6. $\square$

THEOREM 10 (NESTED Plan). The following are some axiularly properties of **Plan**:

1. $\langle \mathcal{B}, \mathcal{A}, Plan(Plan(P)) \rangle \xrightarrow{bdi} \langle \mathcal{B}', \mathcal{A}', Plan(Plan(P')) \rangle$ iff $\langle \mathcal{B}, \mathcal{A}, Plan(P) \rangle \xrightarrow{bdi} \langle \mathcal{B}', \mathcal{A}', Plan(P') \rangle$.
2. $\langle \mathcal{B}, \mathcal{A}, Plan(Plan(P)) \rangle \xrightarrow{bdi} \langle \mathcal{B}', \mathcal{A}', Plan(nil) \rangle$ iff $\langle \mathcal{B}, \mathcal{A}, Plan(P) \rangle \xrightarrow{bdi} \langle \mathcal{B}', \mathcal{A}', nil \rangle$.

PROOF. For the first part we have:

$$\begin{aligned} & \langle \mathcal{B}, \mathcal{A}, Plan(Plan(P)) \rangle \xrightarrow{bdi} \langle \mathcal{B}', \mathcal{A}', Plan(Plan(P')) \rangle \\ & \quad \text{iff} \\ & \langle \mathcal{B}, \mathcal{A}, Plan(P) \rangle \xrightarrow{plan} \langle \mathcal{B}', \mathcal{A}', Plan(P') \rangle \text{ and } \langle \mathcal{B}', \mathcal{A}', Plan(P') \rangle \xrightarrow{plan*} \langle \mathcal{B}'', \mathcal{A}'', nil \rangle \\ & \quad \text{iff} \\ & \langle \mathcal{B}, \mathcal{A}, P \rangle \xrightarrow{plan} \langle \mathcal{B}', \mathcal{A}', P' \rangle \text{ and } \langle \mathcal{B}', \mathcal{A}', P' \rangle \xrightarrow{plan*} \langle \mathcal{B}'', \mathcal{A}'', nil \rangle \\ & \quad \text{iff} \\ & \langle \mathcal{B}, \mathcal{A}, Plan(P) \rangle \xrightarrow{bdi} \langle \mathcal{B}', \mathcal{A}', Plan(P') \rangle \end{aligned}$$

For the second part we have:

$$\langle \mathcal{B}, \mathcal{A}, Plan(Plan(P)) \rangle \xrightarrow{bdi} \langle \mathcal{B}, \mathcal{A}, Plan(nil) \rangle \text{ iff } \langle \mathcal{B}, \mathcal{A}, Plan(P) \rangle \xrightarrow{plan} \langle \mathcal{B}', \mathcal{A}', nil \rangle \text{ iff } P = nil \text{ iff } \langle \mathcal{B}, \mathcal{A}, Plan(P) \rangle \xrightarrow{bdi} \langle \mathcal{B}, \mathcal{A}, nil \rangle. \quad \square$$

C.1.1 Interaction between Plan and Goal

THEOREM 11. Let $R = \langle \mathcal{B}, \mathcal{A}, \text{Plan}(\text{Goal}(\phi_s, P, \phi_f)) \rangle$ such that $\mathcal{B} \not\models \phi_s \wedge \phi_f$. Then,

- if $\mathcal{B} \models \phi_s$, then there exists $R' = \langle \mathcal{B}, \mathcal{A}, \text{Plan}(\text{nil}) \rangle$ such that $R \xrightarrow{\text{bdi}} R'$, and for every R^* such that $R \xrightarrow{\text{bdi}} R^*$, $R^* = R'$.
- if $\mathcal{B} \models \phi_f$, then $R \not\xrightarrow{\text{bdi}}$.

PROOF. If $\mathcal{B} \models \phi_s$, then the only applicable derivation rule is $\mathbf{G}_s$ and $R' = \langle \mathcal{B}, \mathcal{A}, \text{Plan}(\text{nil}) \rangle$. Similarly, if $\mathcal{B} \models \phi_f$, then for every R' such that $\langle \mathcal{B}, \mathcal{A}, \text{Goal}(\phi_s, P, \phi_f) \rangle \xrightarrow{\text{plan}} R'$, it is the case that $R' = \langle \mathcal{B}, \mathcal{A}, \text{?false} \rangle$. $\square$

THEOREM 12 (PROPERTIES OF PROGRAM $\text{Plan}(\phi_s, P, \phi_f)$). Let $C_0 = \langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}, \mathcal{A}, \Gamma_0 \cup \{P_0\} \rangle$ be an agent configuration with $P_0 = \text{Plan}(\phi_s, P, \phi_f)$ (P is a user program). Let $E = C_0 \cdot \dots \cdot C_k$ be an execution where P_0 is currently executing and let $P^\# = \text{Plan}(\text{Goal}(\phi_s, P, \phi_f))$. Then, $P_k \in \{P_0, \text{Goal}(\phi_s, \text{Plan}(\text{Goal}(\phi_s, P^* \triangleright P, \phi_f)) \triangleright P^\#, \phi_f), \text{nil}, \text{?false}\}$.

If $P_k^b = \text{Goal}(\phi_s, \text{Plan}(\text{Goal}(\phi_s, P^* \triangleright P, \phi_f)) \triangleright P^\#, \phi_f)$, then $\langle \mathcal{B}_k, \mathcal{A}_k, P_k^b \rangle \xrightarrow{\text{bdi}} C_k^b \implies$, and the following hold:
 Suppose $\langle \mathcal{B}_k, \mathcal{A}_k, P_k^b \rangle \xrightarrow{\text{bdi}} R$. If $\mathcal{B} \models \phi_s$, then $R = \langle \mathcal{B}_k, \mathcal{A}_k, \text{nil} \rangle$. If $\mathcal{B} \models \phi_f$, then $R = \langle \mathcal{B}_k, \mathcal{A}_k, \text{?false} \rangle$.

Suppose $\mathcal{B} \not\models \phi_s \vee \phi_f$. Then:

$$\begin{aligned}
 & \langle \mathcal{B}_k, \mathcal{A}_k, P^* \rangle \xrightarrow{\text{plan}} \langle \mathcal{B}', \mathcal{A}', P' \rangle \xrightarrow[\text{iff}]{\text{plan}_*} \langle \mathcal{B}'', \mathcal{A}'', P'' \rangle, \mathcal{B}'' \models \phi_s \\
 & \quad \langle \mathcal{B}_k, \mathcal{A}_k, P_k^b \rangle \xrightarrow{\text{plan}} \\
 & \langle \mathcal{B}', \mathcal{A}', \text{Goal}(\phi_s, \text{Plan}(\text{Goal}(\phi_s, P' \triangleright P, \phi_f)) \triangleright P^\#, \phi_f) \rangle \\
 & \quad \text{there is no } R \text{ such that} \\
 & \langle \mathcal{B}_k, \mathcal{A}_k, P^* \rangle \xrightarrow{\text{plan}} R \xrightarrow[\text{iff}]{\text{plan}_*} \langle \mathcal{B}'', \mathcal{A}'', P'' \rangle, \mathcal{B}'' \models \phi_s \\
 & \langle \mathcal{B}_k, \mathcal{A}_k, P_k \rangle \xrightarrow{\text{plan}} \langle \mathcal{B}_k, \mathcal{A}_k, \text{Goal}(\phi_s, P^\# \triangleright P^\#, \phi_f) \rangle.
 \end{aligned}$$

D. PROOF OF THEOREM 1

LEMMA 13. For every $\mathcal{B}, \mathcal{A}$ and P , if $\langle \mathcal{B}, \mathcal{A}, P \rangle \xrightarrow{\text{plan}*} \langle \mathcal{B}_f, \mathcal{A}_f, \text{nil} \rangle$, then $\langle \mathcal{B}, \mathcal{A}, \text{Plan}(P) \rangle \xrightarrow{\text{bdi}*} \langle \mathcal{B}_f, \mathcal{A}_f, \text{nil} \rangle$.

PROOF. We prove this by induction on the length n of the plan-type derivation.

Base case: Suppose $n = 0$. Then $P = \text{nil}$, $\mathcal{B}_f = \mathcal{B}$, and $\mathcal{A}_f = \mathcal{A}$. By using derivation rule Plan_t , we get $\langle \mathcal{B}, \mathcal{A}, \text{Plan}(P) \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}_f, \mathcal{A}_f, \text{nil} \rangle$.

Inductive Case: Suppose $n = k + 1$. Then, there exists $R^* = \langle \mathcal{B}^*, \mathcal{A}^*, P^* \rangle$ such that (a) $\langle \mathcal{B}, \mathcal{A}, P \rangle \xrightarrow{\text{plan}} R^*$ and (b) $R^* \xrightarrow{\text{plan}_k^*} \langle \mathcal{B}_f, \mathcal{A}_f, \text{nil} \rangle$. Using (a) and (b), we can use derivation rule Plan to obtain $\langle \mathcal{B}, \mathcal{A}, \text{Plan}(P) \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}^*, \mathcal{A}^*, \text{Plan}(P^*) \rangle$. Moreover, by (b) and the induction hypothesis, $\langle \mathcal{B}^*, \mathcal{A}^*, \text{Plan}(P^*) \rangle \xrightarrow{\text{bdi}*} \langle \mathcal{B}_f, \mathcal{A}_f, \text{nil} \rangle$. Thus, $\langle \mathcal{B}, \mathcal{A}, \text{Plan}(P) \rangle \xrightarrow{\text{bdi}*} \langle \mathcal{B}_f, \mathcal{A}_f, \text{nil} \rangle$ follows. $\square$

Using the above auxiliary lemma, we now prove the main result.

PROOF OF THEOREM 1: Let $C = \langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}, \mathcal{A}, \{\text{Plan}(P)\} \rangle$ such that $\langle \mathcal{B}, \mathcal{A}, \text{Plan}(P) \rangle \xrightarrow{\text{bdi}}$. If E is an environment-free agent execution of C , then intention $\text{Plan}(P)$ is either executing or has been successfully executed in E . Moreover, there is an execution E^s of C in which intention $\text{Plan}(P)$ has been successfully executed in E^s .

First Claim: On the contrary, assume that $\text{Plan}(P)$ failed in some execution. Then, there is an environment-free execution E of the form $C_0 = C, C_1, \dots, C_k$, such that $k \geq 1$, $C_k = \langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}_k, \mathcal{A}_k, \{\text{Plan}(P_k)\} \rangle$, and $\langle \mathcal{B}_k, \mathcal{A}_k, \text{Plan}(P_k) \rangle \not\xrightarrow{\text{bdi}}$. That is, the execution of C ended up in configuration C_k where the original intention $\text{Plan}(P)$ is stuck.

We know that $\langle \mathcal{B}_{k-1}, \mathcal{A}_{k-1}, \text{Plan}(P_{k-1}) \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}_k, \mathcal{A}_k, \text{Plan}(P_k) \rangle$ which means that (a) $\langle \mathcal{B}_{k-1}, \mathcal{A}_{k-1}, P_{k-1} \rangle \xrightarrow{\text{plan}} \langle \mathcal{B}_k, \mathcal{A}_k, P_k \rangle$; and (b) there exists $\mathcal{B}_f, \mathcal{A}_f$ such that $\langle \mathcal{B}_k, \mathcal{A}_k, P_k \rangle \xrightarrow{\text{plan}*} \langle \mathcal{B}_f, \mathcal{A}_f, \text{nil} \rangle$. By Lemma 13, $\langle \mathcal{B}_k, \mathcal{A}_k, \text{Plan}(P_k) \rangle \xrightarrow{\text{bdi}*} \langle \mathcal{B}_f, \mathcal{A}_f, \text{nil} \rangle$. Because $\text{Plan}(P_k) \neq \text{nil}$, $\langle \mathcal{B}_k, \mathcal{A}_k, \text{Plan}(P_k) \rangle \xrightarrow{\text{bdi}}$ follows (i.e., there is at least one possible next transition).

Second Claim: By assumption, there exist $\mathcal{B}_f$ and $\mathcal{A}_f$ such that $\langle \mathcal{B}, \mathcal{A}, P \rangle \xrightarrow{\text{plan}*} \langle \mathcal{B}_f, \mathcal{A}_f, \text{nil} \rangle$ hold. By Lemma 13, we know that $\langle \mathcal{B}, \mathcal{A}, \text{Plan}(P) \rangle \xrightarrow{\text{bdi}*} \langle \mathcal{B}_f, \mathcal{A}_f, \text{nil} \rangle$. It is not hard to see that we can use this basic bdi-type derivation to construct an environment-free agent execution E_s for C such that the original intention $\text{Plan}(P)$ is successfully executed in it. $\square$

An immediate consequence of this theorem is the following corollary.

COROLLARY 14. Let E be a terminating environment-free execution of agent $\langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}, \mathcal{A}, \{\text{Plan}(P)\} \rangle$. Furthermore, suppose that $\langle \mathcal{B}, \mathcal{A}, \text{Plan}(P) \rangle \xrightarrow{\text{bdi}}$. Then, intention $\text{Plan}(P)$ has been successfully executed in E .

PROOF. From Theorem 1 and the fact that if E is terminating, then it is the case that intention $\text{Plan}(P)$ has either been successfully executed in E or it has failed in E . $\square$

E. TRANSLATION OF BDI LIBRARIES TO HTN DOMAIN KNOWLEDGE: THEOREM 2

Here, we show how to translate a CANPLAN plan library Π and action description library Λ into a planning domain $\mathcal{D} = \langle Op(\Lambda), Me(\Pi) \rangle$.

DEFINITION 10 (BOUNDED BDI AGENTS). We say that a CANPLAN agent is a bounded agent if its belief base and all belief conditions are defined in a language which follows the same constraints as those imposed by HTN planners [11] (e.g., first-order atoms, finite domains, close world assumption). Moreover, a plan-body in a bounded agent cannot include explicit addition or deletion of belief statements (i.e., $\vdash b$ and $\neg b$ statements).

From now on, we assume that CANPLAN agents are bounded. In what follows, we will show two theorems which are specific version of Theorem 2. Theorem 15 is the most trivial version where the translation from the BDI language to the HTN one is almost trivial. However, such theorem holds only for agents that do not make use of the **Goal** construct. Theorem 17 extends Theorem 15 to accommodate **Goal**-programs and relies on a complex transformation for such specific programs. Finally, we discuss how such transformation can be avoided all together if we slightly change the semantics of HTN planning.

E.1 Converting BDI Belief Conditions into HTN Constraints

Given a formula ϕ , we define ϕ^* and $(\phi, n)^*$ inductively as follows:

1. if $\phi = l$, then $\phi^* = l$ and $(\phi, n)^* = (l, n)$.
2. if $\phi = l_1 \wedge l_2$, then $\phi^* = l_1 \wedge l_2$ and $(\phi, n)^* = (l_1, n) \wedge (l_2, n)$.
3. if $\phi = l_1 \vee l_2$, then $\phi^* = l_1 \vee l_2$ and $(\phi, n)^* = (l_1, n) \vee (l_2, n)$.
4. if $\phi = \neg \phi_1$, then $\phi^* = \phi_1^*$ and $(\phi, n)^* = (\phi_1, n)^*$.
5. if $\phi = \neg(\phi_1 \wedge \phi_2)$, then $\phi^* = (\neg \phi_1 \vee \neg \phi_2)^*$ and $(\phi, n)^* = (\neg \phi_1 \vee \neg \phi_2, n)^*$.
6. if $\phi = \neg(\phi_1 \vee \phi_2)$, then $\phi^* = (\neg \phi_1 \wedge \neg \phi_2)^*$ and $(\phi, n)^* = (\neg \phi_1 \wedge \neg \phi_2, n)^*$.

The definitions of $(n, \phi)^*$ and $(n_1, \phi, n_2)^*$ are analogous.

E.2 Converting BDI Action Description Libraries into HTN Operators: $Op(\Lambda)$

Suppose that the (bounded) CANPLAN agent contains an action description library Λ with actions of the following form:

$$act(\vec{x}) : l_1(\vec{x}) \wedge \dots \wedge l_n(\vec{x}) \leftarrow \{a_1^-(\vec{x}), \dots, a_m^-(\vec{x})\}^- ; \{a_1^+(\vec{x}), \dots, a_p^+(\vec{x})\}^+,$$

where l_i are literals, and a_i^-, a_i^+ are atoms. Given a library of actions Λ , we define the corresponding set of HTN primitive tasks/operators $Op(\Lambda)$ as follows:

$$Op(\Lambda) = \{[act(\vec{x}), (pre : l_1(\vec{x}), \dots, l_n(\vec{x})), (post : \neg a_1^-(\vec{x}), \dots, \neg a_m^-(\vec{x}), a_1^+(\vec{x}), \dots, a_p^+(\vec{x}))] \mid act(\vec{x}) : l_1(\vec{x}) \wedge \dots \wedge l_n(\vec{x}) \leftarrow \{a_1^-(\vec{x}), \dots, a_m^-(\vec{x})\}^- ; \{a_1^+(\vec{x}), \dots, a_p^+(\vec{x})\}^+ \in \Lambda\}.$$

E.3 Converting Goal-free BDI Plan Libraries into BDI Method Libraries: $Me(\Pi)$

In this section, we show how a BDI plan library Π that contains no **Goal**-program can be mapped into an HTN method library $Me(\Pi)$. We shall later consider libraries that make use of **Goal**-programs. In the simple case, the mapping is quite straightforward.

Below, $|P(\vec{x})|$ refers to the size of program $P(\vec{x})$ measured as the number of complex constructs in it. We shall also use a special operator (or primitive task) *dummyTask*, which is always possible and has no effects whatsoever when executed. When σ is an HTN plan, *clean*(σ) stands for plan σ with all actions “*dummyTask*” operators (primitive tasks) deleted.

Given a bounded BDI CANPLAN library Π mentioning no goal-programs, we define its corresponding HTN method-library $Me(\Pi)$ as follows:

$$Me(\Pi) = \bigcup_{(e(\vec{x}) : \psi(\vec{x}, \vec{y}) \leftarrow P(\vec{x}, \vec{y}, \vec{z})) \in \Pi} \{(e(\vec{x}), \psi(\vec{x}, \vec{y})^*, \mathcal{T}_1(P(\vec{x}, \vec{y}, \vec{z}), 0))\} \cup \mathcal{T}_2(P(\vec{x}, \vec{y}, \vec{z}), 0).$$

Function $\mathcal{T}(P, n)$ maps a plan-body P and a natural number n into a pair $\langle T, M \rangle$, where T is an HTN task-network and M is a set of methods. We write $\mathcal{T}_1(P, n)$ to refer to the first argument T of the pair, and $\mathcal{T}_2(P, n)$ to refer to the second argument M —that is, if $\mathcal{T}(P, n) = \langle T, M \rangle$ then $\mathcal{T}_1(P, n) = T$ and $\mathcal{T}_2(P, n) = M$. Function $\mathcal{T}(P, n)$ is defined inductively on the structure of P as follows:

Base Cases: In this case, the program can be a primitive action, a test condition, an event, or the trivial program *nil*:

- If $P(\vec{x}) = act(\vec{x})$, then $\mathcal{T}(P(\vec{x}), n) = \langle [\{n : act(\vec{x})\}, nil], \emptyset \rangle$.
- If $P(\vec{x}) = ?\phi(\vec{x})$, then $\mathcal{T}(P(\vec{x}), n) = \langle [\{n : dummyTask\}, (\phi(\vec{x}), n)^*], \emptyset \rangle$.
- If $P(\vec{x}) = e(\vec{x})$, then $\mathcal{T}(P(\vec{x}), n) = \langle [\{n : e(\vec{x})\}, nil], \emptyset \rangle$.
- If $P(\vec{x}) = nil$, then $\mathcal{T}(P(\vec{x}), n) = \langle [\{n : dummyTask\}, nil], \emptyset \rangle$.

Inductive Cases: Suppose $P_b(\vec{x})$ is any of the programs in the above base cases, that is, $P_b(\vec{x}) = act(\vec{x}) \mid ?\phi(\vec{x}) \mid e(\vec{x}) \mid nil$. Then,

- If $P(\vec{x}) = (P_b(\vec{x}); P'(\vec{x}))$, then

$$\mathcal{T}(P(\vec{x}), n) = \langle [T_b \cup T', (n < n + 1) \wedge C_b \wedge C'], \emptyset \rangle,$$

where $\mathcal{T}(P_b(\vec{x}), n) = \langle T_b, C_b \rangle$ and $\mathcal{T}(P'(\vec{x}), n + 1) = \langle T', C' \rangle$.

- If $P(\vec{x}) = (P_1(\vec{x}) \parallel P_2(\vec{x}))$, then

$$\mathcal{T}(P(\vec{x}), n) = \langle [T_1 \cup T_2, C_1 \wedge C_2], \emptyset \rangle,$$

where $\mathcal{T}(P_1(\vec{x}), n) = \langle T_1, C_1 \rangle$ and $\mathcal{T}(P_2(\vec{x}), n + |P_1(\vec{x})| + 1) = \langle T_2, C_2 \rangle$.

- If $P(\vec{x}) = (P_1(\vec{x}) \triangleright P_2(\vec{x}))$, then $\mathcal{T}(P(\vec{x}), n) = \mathcal{T}(P_1(\vec{x}), n)$.
- If $P(\vec{x}) = (\psi_1(\vec{x}) : P_1(\vec{x}), \dots, \psi_k(\vec{x}) : P_k(\vec{x}))$, then

$$\mathcal{T}(P(\vec{x}), n) = \langle [n : choice^P(\vec{x}), nil], \bigcup_{i=1}^k \{(choice^P(\vec{x}), \psi_i(\vec{x})^*, \mathcal{T}_1(P_i(\vec{x}), 0))\} \cup \mathcal{T}_2(P_i(\vec{x}), 0) \rangle.$$

THEOREM 15 (Plan and HTN-PLANNING IN Goal-FREE PROGRAMS). Let $\langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}, \mathcal{A}, \Gamma \cup \{\text{Plan}(P)\} \rangle$ be a bounded agent, where Π is a goal-free plan library and P is a goal-free plan-body. Let $\mathcal{D} = \langle \text{Op}(\Lambda) \cup \{\text{dummyTask}\}, \text{Me}(\Pi) \rangle$ be the corresponding HTN problem domain and let $T_P = T_1(P, 0)$ be the network task obtained from P (see as defined in Appendixes E.2 and E.3). Then,

1. $\langle \mathcal{B}, \mathcal{A}, \text{Plan}(P) \rangle \xrightarrow{\text{bdi}} \text{iff } \text{sol}(T_P, \mathcal{B}, \mathcal{D}) \neq \emptyset$.
2. $\langle \mathcal{B}, \mathcal{A}, \text{Plan}(P) \rangle \xrightarrow{\text{bdi}^*} \langle \mathcal{B}', \mathcal{A} \cdot \text{act}_1 \cdot \dots \cdot \text{act}_k, \text{Plan}(P') \rangle$ with $k \geq 1$ iff there exists a plan $\sigma \in \text{sol}(T_P, \mathcal{B}, \mathcal{D})$, such that $\text{clean}(\sigma) = \text{act}_1 \cdot \dots \cdot \text{act}_k \cdot \dots \cdot \text{act}_n$, for some $n \geq k$.
3. $\langle \mathcal{B}, \mathcal{A}, (\text{act}_1; \dots; \text{act}_n) \rangle \xrightarrow{\text{bdi}^*} \langle \mathcal{B}', \mathcal{A} \cdot \sigma, \text{nil} \rangle$, for every $\sigma \in \text{sol}(T_P, \mathcal{B}, \mathcal{D})$ such that $\text{clean}(\sigma) = \text{act}_1 \cdot \dots \cdot \text{act}_n$.

PROOF. This is a laborious proof showing that plan-type transitions perform no more than the task decomposition done by HTN planners. The proof is based on the relationship between BDI's and HTN's entities as discussed in section 2.3. $\square$

E.4 Converting Full BDI Plan Libraries into BDI Method Libraries: $\text{Me}^+(\Pi)$

We shall now consider BDI plan libraries and plan-bodies that may mention Goal-programs. To that end, we will extend the transformation given in Appendix E.3 to include Goal-programs; this transformation is a bit more involved given that HTN does not accommodate naturally a construct of that sort. We start by converting a BDI plan library into one that is Goal-free.

E.4.1 Mapping Goal-programs to Goal-free Programs

Now, let us consider goal-programs of the form $\text{Goal}(\phi_s, P, \phi_f)$. Informally, we will construct a library Π^* from an original library Π in such a way that executing goal-program $\text{Goal}(\phi_s, P, \phi_f)$ w.r.t. library Π is equivalent to executing program P w.r.t. library Π^* . We observe this transformation does not complain with the notion of *eliminability* (see [?]) as it substantially changes the structure of the agent.

DEFINITION 11 (PROGRAM AND LIBRARY TRANSFORMATIONS P^γ AND Π^γ). Let P be a plan-body and let $\gamma(\vec{w})$ be a belief condition. The plan body P^γ is the plan body obtained from P as follows:

1. Every event $e(\vec{t})$ mentioned in P is replaced with the (new) event $e^\gamma(\vec{t}, \vec{w})$ in P^γ .
2. Every test condition $\phi(\vec{t})$ mentioned in P is replaced with the test condition $\phi(\vec{t}) \vee \gamma(\vec{w})$ in P^γ .
3. Every action $\text{act}(\vec{t})$ mentioned in P is replaced with the (new) event $e_{\text{act}}^\gamma(\vec{t}, \vec{w})$ in P^γ .

When Π is a plan library, we build the new library Π^γ from Π , as follows:

1. Every plan-rule of the form $e(\vec{x}) : \psi(\vec{x}, \vec{y}) \leftarrow P(\vec{x}, \vec{y}, \vec{z})$ is replaced with the following two plan-rules:

$$\begin{aligned} e^\gamma(\vec{x}, \vec{w}) : \neg\gamma(\vec{w}) \wedge \psi(\vec{x}, \vec{y}) &\leftarrow P^\gamma, \\ e^\gamma(\vec{x}, \vec{w}) : \gamma(\vec{w}) &\leftarrow \text{nil}. \end{aligned}$$

2. For every action $\text{act}(\vec{x})$ in the domain, the following two plan-rules are included:

$$\begin{aligned} e_{\text{act}}^\gamma(\vec{x}, \vec{w}) : \neg\gamma(\vec{w}) &\leftarrow \text{act}(\vec{x}), \\ e_{\text{act}}^\gamma(\vec{x}, \vec{w}) : \gamma(\vec{w}) &\leftarrow \text{nil}. \end{aligned}$$

Relation $\xrightarrow{*}_\phi$ stands for the reflexive transitive closure of $\xrightarrow{*}$ where ϕ holds at every configuration. More concretely, $C_1 \xrightarrow{*}_\phi C_n$ ($n \geq 2$) iff there exists $C_2, \dots, C_{n-1}$ such that $C_i \xrightarrow{*} C_{i+1}$ and $\mathcal{B}_i \models \phi$, for $1 \leq i \leq n-1$.

THEOREM 16. For every $\mathcal{B}, \mathcal{A}, \Pi$, and program $\text{Goal}(\phi_s, P, \phi_f)$,

1. $\langle \Pi, \mathcal{B}, \mathcal{A}, \text{Goal}(\phi_s, P, \phi_f) \rangle \xrightarrow{\text{bdi}^*} \langle \Pi, \mathcal{B}', \mathcal{A}', \text{nil} \rangle$ iff $\langle \Pi \cup \Pi^{\phi_s}, \mathcal{B}, \mathcal{A}, P^{\phi_s} \rangle \xrightarrow{\text{bdi}^*}_{\neg\phi_f} \langle \Pi \cup \Pi^{\phi_s}, \mathcal{B}', \mathcal{A}', \text{nil} \rangle$
2. $\langle \Pi \cup \Pi^{\phi_s}, \mathcal{B}, \mathcal{A}, P^{\phi_s} \rangle \xrightarrow{\text{bdi}^*}_{\neg\phi_f} \langle \Pi \cup \Pi^{\phi_s}, \mathcal{B}', \mathcal{A}', \text{nil} \rangle$ iff $\langle \Pi^{\phi_s}, \mathcal{B}, \mathcal{A}, P^{\phi_s} \rangle \xrightarrow{\text{bdi}^*}_{\neg\phi_f} \langle \Pi^{\phi_s}, \mathcal{B}', \mathcal{A}', \text{nil} \rangle$

PROOF. Point (2) follows directly from the fact that there is no plan rule for any of the events mentioned in program P^ϕ or library Π^ϕ . Let us now prove then point (1):

($\Rightarrow$)
($\Leftarrow$) $\square$

Once again, this theorem is a very weak notion of eliminability for the Goal-construct. However, we only need that for our purposes, namely, that

E.4.2 Obtaining an HTN Method Library from a BDI Library with Goal-programs

We now extend the transformation $\text{Me}(\Pi)$ given in Appendix E.3 to accommodate plan-bodies mentioning Goal-programs. To that end, we define the transformation $\text{Me}^+(\Pi)$ to be exactly like $\text{Me}(\Pi)$ with the following extra inductive case in the definition of $\mathcal{T}$:

- If $P(\vec{x}) = \text{Goal}(\phi_s, P, \phi_f)$, then

$$\mathcal{T}(P(\vec{x}), n) = \{ \{ \{ n : \text{dummyTask}, n+1 : \text{achieve}_{\mathcal{P}}^{\phi_s}, n+2 : \text{dummyTask} \}, C \}, \text{Me}(\{ \text{achieve}_{\mathcal{P}}^{\phi_s} : \text{nil} \leftarrow P^{\phi_s} \} \cup \Pi'^{\phi_s}) \},$$

where $C = (n < n+1) \wedge (n+1 < n+2) \wedge (n, \neg\phi_f, n+2)^*$, and Π' is Π with all $\text{Goal}(\phi_s, P, \phi_f)$ programs replaced with P .

Next, we generalize Theorem 15 for any kind of BDI plan library.

THEOREM 17 (Plan and HTN-PLANNING IN FULL PROGRAMS). Let $\langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}, \mathcal{A}, \Gamma \cup \{\text{Plan}(P)\} \rangle$ be a bounded agent, where Π is a goal-free plan library and P is a goal-free plan-body. Let $\mathcal{D} = \langle \text{Op}_\Lambda \cup \{\text{dummyTask}\}, \text{Me}(\Pi) \cup M_P \rangle$ be the corresponding HTN problem domain where $\mathcal{T}(P, 0) = \langle T_P, M_P \rangle$ (see Appendixes E.2 and E.3, and the extension of $\mathcal{T}$ above). Then,

1. $\langle \mathcal{B}, \mathcal{A}, \text{Plan}(P) \rangle \xrightarrow{\text{bdi}} \text{iff } \text{sol}(T_P, \mathcal{B}, \mathcal{D}) \neq \emptyset$.
2. $\langle \mathcal{B}, \mathcal{A}, \text{Plan}(P) \rangle \xrightarrow{\text{bdi}^*} \langle \mathcal{B}', \mathcal{A} \cdot \text{act}_1 \cdot \dots \cdot \text{act}_k, \text{Plan}(P') \rangle$ with $k \geq 1$ iff there exists a plan $\sigma \in \text{sol}(T_P, \mathcal{B}, \mathcal{D})$, such that $\text{clean}(\sigma) = \text{act}_1 \cdot \dots \cdot \text{act}_k \cdot \dots \cdot \text{act}_n$, for some $n \geq k$.
3. $\langle \mathcal{B}, \mathcal{A}, (\text{act}_1; \dots; \text{act}_n) \rangle \xrightarrow{\text{bdi}^*} \langle \mathcal{B}', \mathcal{A} \cdot \sigma, \text{nil} \rangle$, for every $\sigma \in \text{sol}(T_P, \mathcal{B}, \mathcal{D})$ such that $\text{clean}(\sigma) = \text{act}_1 \cdot \dots \cdot \text{act}_n$.

PROOF. Follows from Theorems 15 and 16. $\square$

E.5 An Extension of HTN Semantics for Partial Plans

Here we explain that is actually cleaner and more practical to slightly change the semantics of HTN planning to accommodate a success condition.

F. PROOF OF THEOREM 3

THEOREM 18. Let Sim be a simulation relation w.r.t. relation $\xrightarrow{\text{bdi}}$ such that $\text{Sim}(P^*, P)$. Then, program P^* simulates program P .

PROOF. Informally, P^* simulates P because: (i) whenever P can be terminate with internal steps, P^* can also be terminated in the same way; and (ii) whenever P performs a bdi-step, P^* can perform the same step result, possibly by also doing some extra internal steps.

Let $E = C_1 \cdot \dots \cdot C_n$, $n \geq 1$, be an execution of a configuration $C_1 = \langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}_1, \mathcal{A}_1, \Gamma'_1 \cup \{P\} \rangle$. We shall obtain an execution E^* of $C_1^* = \langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}_1, \mathcal{A}_1, \Gamma'_1 \cup \{P\} \rangle$ such that $\overline{E}$ and $\overline{E}^*$ are equivalent modulo P and P^* , and if P has been successfully executed in E so has P^* in E^* .

Suppose $|E| = 1$ and hence $E = C_1$. If P has not been successfully executed yet in E , then we take $E^* = C_1^*$ which is trivially equivalent to E modulo P and P^* and the thesis follows. If P has been executed in E , then $P = \text{nil}$ and $\langle \mathcal{B}, \mathcal{A}, P \rangle \xrightarrow{\text{bdi}_*} \langle \mathcal{B}', \mathcal{A}', \text{nil} \rangle$. By point (1) in Definition 9, $\langle \mathcal{B}, \mathcal{A}, P^* \rangle \xrightarrow{\text{bdi}_*} \langle \mathcal{B}', \mathcal{A}', \text{nil} \rangle$. We can then use these internal transitions to build an execution $E^* = C_1^* \cdot \dots \cdot C_k^*$, $k \geq 1$, such that $P_k^* = \text{nil}$ and $\mathcal{O}(C_i^*) = \mathcal{O}(C_1^*)$, for $i = 1..k$. Clearly, $\overline{E}^* = C_k^*$ is equivalent to E and P^* has been executed in E^* .

Next suppose then $|E| = n = k + 1$, $k \geq 0$. Then, there exist an execution $E' = C_2 \cdot \dots$ such that $|E'| = k$ and $C_1 \implies C_2$. We consider then the following two cases:

1. If agent transition $C_1 \implies C_2$ is not a transition on intention P itself, then it is clear that the same agent transition can be performed from C_1^* to obtain C_2^* . That is, if $C_2 = \langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}_2, \mathcal{A}_2, \Gamma'_2 \cup \{P\} \rangle$ then $C_2^* = \langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}_2, \mathcal{A}_2, \Gamma'_2 \cup \{P^*\} \rangle$ such that $C_1^* \implies C_2^*$. By the induction hypothesis, there exists an execution E'^* of C_2^* such that $\overline{E'}$ and $\overline{E'^*}$ are equivalent modulo P and P^* , and if P has been successfully executed in E' so has P^* in E'^* . Then, it follows that $\overline{C_1 \cdot E'}$ and $\overline{C_1^* \cdot E'^*}$ are equivalent modulo P and P^* , and if P has been successfully executed in $C_1 \cdot E'$ so has P^* in $C_1^* \cdot E'^*$.
2. Suppose the agent transition $C_1 \implies C_2$ is a transition on intention P itself. Then, $C_2 = \langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}_2, \mathcal{A}_2, \Gamma'_1 \cup \{P'\} \rangle$ such that $\langle \mathcal{B}_1, \mathcal{A}_1, P \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}_2, \mathcal{A}_2, P' \rangle$. Then, there has to exist R', R'', P'_1 such that (a) $\langle \mathcal{B}_1, \mathcal{A}_1, P^* \rangle \xrightarrow{\text{bdi}_*} R'$, $R' \xrightarrow{\text{bdi}} R''$ or $R' = R''$, and $R'' \xrightarrow{\text{bdi}_*} \langle \mathcal{B}_2, \mathcal{A}_2, P^{*'} \rangle$; and (b) $\text{Sim}(P_1^{*'}, P')$. We use all the basic transitions from (a) to obtain an execution $E_a^* = C_1^* \cdot C_1^i \cdot \dots \cdot C_j^i \cdot C_2^*$ of C_1^* , where $j \geq 0$ and $C_2^* = \langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}_2, \mathcal{A}_2, \Gamma'_1 \cup \{P^{*'}\} \rangle$. Because, $\text{Sim}(P_1^{*'}, P')$ and the induction hypothesis, there exists an execution E'^* of C_2^* such that $\overline{E'}$ and $\overline{E'^*}$ are equivalent modulo P' and $P_1^{*'}$, and if P' has been successfully executed in E' so has $P_1^{*'}$ in E'^* . Then, it follows that $\overline{C_1 \cdot E'}$ and $\overline{C_1^* \cdot C_1^i \cdot \dots \cdot C_j^i \cdot E'^*}$ are equivalent modulo P and P^* , and if P has been successfully executed in $C_1 \cdot E'$ so has P^* in $C_1^* \cdot C_1^i \cdot \dots \cdot C_j^i \cdot E'^*$.

□

PROOF OF THEOREM 3: For every ϕ_s, ϕ_f and P , program $\text{Plan}(\phi_s, P, \phi_f)$ simulates program $\text{Plan}(\text{Goal}(\phi_s, P, \phi_f))$.

Let $\text{Sim}(-, -)$ be any relation satisfying, at least, the following conditions for every P, P', ϕ_s, ϕ_f :

1. $\text{Sim}(\text{Goal}(\phi_s, \text{Plan}(\text{Goal}(\phi_s, P, \phi_f))), \phi_f), \text{Plan}(\text{Goal}(\phi_s, P, \phi_f)))$.
2. $\text{Sim}(\text{Goal}(\phi_s, \text{Plan}(\text{Goal}(\phi_s, P \triangleright P, \phi_f))) \triangleright \text{Plan}(\text{Goal}(\phi_s, P, \phi_f)), \phi_f), \text{Plan}(\text{Goal}(\phi_s, P \triangleright P, \phi_f)))$.
3. $\text{Sim}(\text{Goal}(\phi_s, \text{Plan}(\text{Goal}(\phi_s, P' \triangleright P, \phi_f))) \triangleright \text{Plan}(\text{Goal}(\phi_s, P, \phi_f)), \phi_f), \text{Plan}(\text{Goal}(\phi_s, P' \triangleright P, \phi_f)))$.
4. $\text{Sim}(\text{nil}, \text{Plan}(\text{nil}))$.
5. $\text{Sim}(\text{nil}, \text{nil})$.

It is not hard to verify that the above relation Sim is indeed a simulation relation w.r.t. relation $\xrightarrow{\text{bdi}}$ as defined in Definition 9. By Theorem 18, the thesis follows. □

LEMMA 19. For every $\mathcal{B}, \mathcal{A}$ and P_1, P_2 such that $RP(P_1, P_2)$:

1. $\langle \mathcal{B}, \mathcal{A}, \text{Plan}(P_1) \rangle \xrightarrow{\text{plan}} \langle \mathcal{B}', \mathcal{A}', P'_1 \rangle$ iff $\langle \mathcal{B}, \mathcal{A}, P_1 \rangle \xrightarrow{\text{plan}} \langle \mathcal{B}', \mathcal{A}', P'_1 \rangle$.
2. If $\langle \mathcal{B}, \mathcal{A}, \text{Plan}(P_1) \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}', \mathcal{A}', P'_1 \rangle$, then there exists P'_2 such that $\langle \mathcal{B}, \mathcal{A}, \text{Plan}(P_2) \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}', \mathcal{A}', P'_2 \rangle$ and $RP(P'_1, P'_2)$.
3. If $\langle \mathcal{B}, \mathcal{A}, \text{Plan}(P_2) \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}', \mathcal{A}', P'_2 \rangle$, then there exists P'_1 such that $\langle \mathcal{B}, \mathcal{A}, \text{Plan}(P_1) \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}', \mathcal{A}', P'_1 \rangle$ and $RP(P'_1, P'_2)$.

G. PROOF OF THEOREM 4

LEMMA 20. *For every agent free of $\parallel$ and Goal, if $\langle \mathcal{B}, \mathcal{A}, \text{Plan}(P) \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}', \mathcal{A}', \text{Plan}(P') \rangle$, $\langle \mathcal{B}, \mathcal{A}, P \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}', \mathcal{A}', P' \rangle$.*

PROOF. By assumption, we know that rule **Plan** was used and : (a) $\langle \mathcal{B}, \mathcal{A}, P \rangle \xrightarrow{\text{plan}} \langle \mathcal{B}', \mathcal{A}', P' \rangle$; and (b) $\langle \mathcal{B}', \mathcal{A}', P' \rangle \xrightarrow{\text{plan}^*} \langle \mathcal{B}_f, \mathcal{A}_f, \text{nil} \rangle$. Given (a) and (b), we shall prove, by induction on the number of **plan**-type derivation rules involved in (a), that $\langle \mathcal{B}, \mathcal{A}, P \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}', \mathcal{A}', P' \rangle$. Suppose then that only one derivation rule is required to prove (a). Then, we have the following cases:

- Suppose $P = \text{act}$, $P = ?\phi$, $P = +b$ or $P = -b$. In this case, $P' = \text{nil}$, and $\langle \mathcal{B}, \mathcal{A}, P \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}', \mathcal{A}', \text{nil} \rangle$ follows trivially.
- Suppose $P = (\text{nil}; P^a)$. In this case, $P' = P^a$ and $\langle \mathcal{B}, \mathcal{A}, P \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}', \mathcal{A}', P^a \rangle$ follows trivially.
- Suppose $P = (\text{nil} \triangleright P^a)$. In this case, $P' = \text{nil}$ and $\langle \mathcal{B}, \mathcal{A}, P \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}', \mathcal{A}', \text{nil} \rangle$ follows trivially.
- Suppose $P = !e$. In this case, $P' = \langle \Delta \rangle$, and $\langle \mathcal{B}, \mathcal{A}, P \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}', \mathcal{A}', \langle \Delta \rangle \rangle$ follows trivially.
- Suppose $P = \langle \Delta \rangle$. In this case, $P' = P_i \theta \triangleright \langle \Delta \setminus P_i \rangle$ due to rule **Sel** and $\langle \mathcal{B}, \mathcal{A}, P \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}', \mathcal{A}', P_i \theta \triangleright \langle \Delta \setminus P_i \rangle \rangle$ holds too.
- Suppose $P = \text{Plan}(\text{nil})$. In this case, $P' = \text{nil}$ and $\langle \mathcal{B}, \mathcal{A}, P \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}', \mathcal{A}', \text{nil} \rangle$ follows directly from derivation rule **Plan_t**.

Next, suppose that $k + 1$ derivation rules are used for (a). We then have the following cases:

- Suppose $P = (P_1; P_2)$. In this case, $P \neq \text{nil}$ and $P' = (P'_1; P_2)$ such that $\langle \mathcal{B}, \mathcal{A}, P_1 \rangle \xrightarrow{\text{plan}} \langle \mathcal{B}', \mathcal{A}', P'_1 \rangle \xrightarrow{\text{plan}^*} \langle \mathcal{B}'', \mathcal{A}'', \text{nil} \rangle$. By the induction hypothesis, $\langle \mathcal{B}, \mathcal{A}, P_1 \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}', \mathcal{A}', P'_1 \rangle$ and $\langle \mathcal{B}, \mathcal{A}, P \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}', \mathcal{A}', P' \rangle$ holds due to rule **Seq**.
- Suppose $P = (P_1 \triangleright P_2)$. In this case, $P \neq \text{nil}$ and $P' = (P'_1 \triangleright P_2)$ such that $\langle \mathcal{B}, \mathcal{A}, P_1 \rangle \xrightarrow{\text{plan}} \langle \mathcal{B}', \mathcal{A}', P'_1 \rangle \xrightarrow{\text{plan}^*} \langle \mathcal{B}'', \mathcal{A}'', \text{nil} \rangle$. By the induction hypothesis, $\langle \mathcal{B}, \mathcal{A}, P_1 \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}', \mathcal{A}', P'_1 \rangle$ and $\langle \mathcal{B}, \mathcal{A}, P \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}', \mathcal{A}', P' \rangle$ holds due to rule $\triangleright$.
- Suppose $P = \text{Plan}(P_1)$. In this case, $P' = P'_1$ such that $\langle \mathcal{B}, \mathcal{A}, P_1 \rangle \xrightarrow{\text{plan}} \langle \mathcal{B}', \mathcal{A}', P'_1 \rangle \xrightarrow{\text{plan}^*} \langle \mathcal{B}'', \mathcal{A}'', \text{nil} \rangle$. By directly applying rule **Plan**, we know that $\langle \mathcal{B}, \mathcal{A}, \text{Plan}(P_1) \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}', \mathcal{A}', P'_1 \rangle$ and thus $\langle \mathcal{B}, \mathcal{A}, P \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}', \mathcal{A}', P' \rangle$.

This concludes the proof of the Theorem. $\square$

LEMMA 21. *Let $\text{Sim}(_, _)$ be any relation satisfying, at least, the following conditions for every program P that does not mention construct $\parallel$:*

1. $\text{Sim}(P, \text{Plan}(P))$.
2. $\text{Sim}(\text{nil}, \text{Plan}(\text{nil}))$.
3. $\text{Sim}(\text{nil}, \text{nil})$.

Then, Sim is a simulation relation w.r.t. transition $\xrightarrow{\text{bdi}}$ for any agent that is free of concurrency.

PROOF. Take a program $P \neq \text{nil}$. $\square$

PROOF OF THEOREM 4: Program P simulates program $\text{Plan}(P)$ in every CANPLAN^- agent.

Let $\text{Sim}(_, _)$ be any relation satisfying, at least, the following conditions for every program P that does not mention constructs $\parallel$ and **Goal**:

1. $\text{Sim}(P, \text{Plan}(P))$.
2. $\text{Sim}(\text{nil}, \text{Plan}(\text{nil}))$.
3. $\text{Sim}(\text{nil}, \text{nil})$.

We also prove the following result. $\square$

THEOREM 22. *Let $C = \langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}, \mathcal{A}, \Gamma \cup \{\text{Plan}(P)\} \rangle$ be a CANPLAN agent. Let $C' = \langle \mathcal{N}, \Lambda', \Pi, \mathcal{B}, \mathcal{A}, \Gamma' \cup \{P'\} \rangle$ where Λ', Γ' and P' are obtained from Λ, Γ and P by deleting all **Plan** constructs from plan-bodies. Then, configuration C' simulates configuration C .*

H. OTHER STUFF

LEMMA 23 (EQUAL INITIALIZATION). Let $P_a = \text{Plan}(\text{Goal}(\phi_s, P, \phi_f))$, $P_b = \text{Goal}(\phi_s, \text{Plan}(P), \phi_f)$, and $P^* = \text{Goal}(\phi_s, P_a, \phi_f)$, where P is a user program. Suppose that $\langle \mathcal{B}, \mathcal{A}, P_a \rangle \xrightarrow{\text{bdi}} C_a$, $\langle \mathcal{B}, \mathcal{A}, P_b \rangle \xrightarrow{\text{bdi}} C_b$, and $\langle \mathcal{B}, \mathcal{A}, P^* \rangle \xrightarrow{\text{bdi}} C^*$. Then,

1. $C_a = \langle \mathcal{B}', \mathcal{A}', \text{nil} \rangle$ iff $C_b = C^* = \langle \mathcal{B}', \mathcal{A}', \text{nil} \rangle$ (here, $\mathcal{B} \models \phi_s$).
2. C_a does not actually exist iff $C_b = C^* = \langle \mathcal{B}, \mathcal{A}, ?\text{false} \rangle$ (here, $\mathcal{B} \models \phi_f$).
3. $C_a = \langle \mathcal{B}', \mathcal{A}', P'_a \rangle$, where $P'_a = \text{Plan}(\text{Goal}(\phi_s, P \triangleright P, \phi_f))$ iff $C_b = \langle \mathcal{B}', \mathcal{A}', \text{Goal}(\phi_s, \text{Plan}(P) \triangleright \text{Plan}(P_o), \phi_f) \rangle$ iff $C^* = \langle \mathcal{B}', \mathcal{A}', \text{Goal}(\phi_s, P_a \triangleright P_a, \phi_f) \rangle$ and for every C^{**} such that $C^* \xrightarrow{\text{bdi}} C^{**}$, $C^{**} = \langle \mathcal{B}', \mathcal{A}', \text{Goal}(\phi_s, P'_a \triangleright P_a, \phi_f) \rangle$
4. One of the above three cases must apply for C_a .

LEMMA 24. Let $P_a = \text{Plan}(\text{Goal}(\phi_s, P \triangleright P_o, \phi_f))$ and $P^* = \text{Goal}(\phi_s, P_a \triangleright P_o, \phi_f)$ be two programs. Then, for every $\mathcal{B}, \mathcal{A}$:

1. $\langle \mathcal{B}, \mathcal{A}, P_a \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}, \mathcal{A}, \text{Plan}(\text{nil}) \rangle$ iff $\langle \mathcal{B}, \mathcal{A}, P^* \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}, \mathcal{A}, \text{nil} \rangle$.
2. If $\langle \mathcal{B}, \mathcal{A}, P_a \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}, \mathcal{A}, P'_a \rangle$, where $P'_a = \text{Plan}(\text{Goal}(\phi_s, P' \triangleright P_o, \phi_f))$, then $\langle \mathcal{B}, \mathcal{A}, P^* \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}, \mathcal{A}, \text{Goal}(\phi_s, P'_a \triangleright P_o, \phi_f) \rangle$.
3. If $\langle \mathcal{B}, \mathcal{A}, P_a \rangle \xrightarrow{\text{bdi}} C$ and $\langle \mathcal{B}, \mathcal{A}, P^* \rangle \xrightarrow{\text{bdi}} C$, then $C = \langle \mathcal{B}, \mathcal{A}, \text{Goal}(\phi_s, P_o \triangleright P_o, \phi_f) \rangle$ or $C = \langle \mathcal{B}, \mathcal{A}, ?\text{false} \rangle$ and $\mathcal{B} \models \phi_f$.
4. If $\langle \mathcal{B}, \mathcal{A}, P_a \rangle \xrightarrow{\text{bdi}} C$, then $C = \langle \mathcal{B}, \mathcal{A}, \text{Plan}(\text{nil}) \rangle$ or $C = \langle \mathcal{B}', \mathcal{A}', \text{Plan}(\text{Goal}(\phi_s, P' \triangleright P_o, \phi_f)) \rangle$ for some program P' .

PROOF. (1) Follows directly from the fact that if $\langle \mathcal{B}, \mathcal{A}, P_a \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}, \mathcal{A}, \text{Plan}(\text{nil}) \rangle$ or $\langle \mathcal{B}, \mathcal{A}, P^* \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}, \mathcal{A}, \text{nil} \rangle$, then $\mathcal{B} \models \phi_s$.

(2) Suppose that $\langle \mathcal{B}, \mathcal{A}, P_a \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}, \mathcal{A}, P'_a \rangle$ where $P'_a = \text{Plan}(\text{Goal}(\phi_s, P' \triangleright P_o, \phi_f))$. Then, $\langle \mathcal{B}, \mathcal{A}, \text{Goal}(\phi_s, P \triangleright P_o, \phi_f) \rangle \xrightarrow{\text{plan}} \langle \mathcal{B}, \mathcal{A}, \text{Goal}(\phi_s, P' \triangleright P_o, \phi_f) \rangle$ due to rule G_S and therefore $\mathcal{B} \not\models \phi_s \vee \phi_f$. Then, by applying rule G_S , $\langle \mathcal{B}, \mathcal{A}, P^* \rangle \xrightarrow{\text{bdi}} \langle \mathcal{B}, \mathcal{A}, \text{Goal}(\phi_s, P'_a \triangleright P_o, \phi_f) \rangle$ follows.

(3) Suppose that $\langle \mathcal{B}, \mathcal{A}, P_a \rangle \not\xrightarrow{\text{bdi}}$ and that there exists C such that $\langle \mathcal{B}, \mathcal{A}, P^* \rangle \xrightarrow{\text{bdi}} C$. The transition to C may be due to rules $\text{G}_s, \text{G}_f, \text{G}_S$ or G_R^{bdi} . Because $\langle \mathcal{B}, \mathcal{A}, P_a \rangle \not\xrightarrow{\text{bdi}}$, there are no $\mathcal{B}'$ and $\mathcal{A}'$ such that $\langle \mathcal{B}, \mathcal{A}, \text{Goal}(\phi_s, P \triangleright P_o, \phi_f) \rangle \xrightarrow{\text{plan}^*} \langle \mathcal{B}', \mathcal{A}', \text{nil} \rangle$. Clearly, $\mathcal{B} \not\models \phi_s$ or otherwise, $\langle \mathcal{B}, \mathcal{A}, \text{Goal}(\phi_s, P \triangleright P_o, \phi_f) \rangle \xrightarrow{\text{plan}} \langle \mathcal{B}, \mathcal{A}, \text{nil} \rangle$ would hold. As a result, rule G_s may not apply. Similarly, rule G_S may not apply either or otherwise $\langle \mathcal{B}, \mathcal{A}, P_a \rangle \xrightarrow{\text{bdi}}$.

Now, if $\mathcal{B} \models \phi_f$, then rule G_f applies and the thesis follows trivially. Otherwise, if $\mathcal{B} \not\models \phi_f$, then the only rule that may apply is G_R^{bdi} and the thesis follows as well.

(4) Trivial since the only derivation rules that may apply for making a Plan-transition on program P_a are $\text{G}_s, \text{G}_R^{\text{bdi}}$ or G_S . $\square$

H.1 Goal and Plan Text

A central feature of CAN and CANPLAN is the Goal construct for handling declarative goals. It is important then to understand the interaction of goal-programs with the new planning construct Plan. That is, we want to explore the differences and similarities between “having the goal to plan” (i.e., $\text{Goal}(\phi_s, \text{Plan}(P), \phi_f)$) and “planning for a declarative goal” (i.e., $\text{Plan}(\text{Goal}(\phi_s, P, \phi_f))$).

To begin with, it is easy to prove that both programs are “initialised” in equivalent ways: the first goal-plan-program is initialised to program $\text{Goal}(\phi_s, \text{Plan}(P) \triangleright \text{Plan}(P), \phi_f)$ iff the second plan-goal program is initialised to program $\text{Plan}(\text{Goal}(\phi_s, P \triangleright P, \phi_f))$.

Now, a goal-plan program of the form $\text{Goal}(\phi_s, \text{Plan}(P), \phi_f)$ has the same meaning as the one provided in CAN: execute $\text{Plan}(P)$ until the goal ϕ_s is achieved or the failure condition ϕ_f applies. Note that nothing precludes solving P completely without achieving ϕ_s . Still, due to the *persistence* property of Goal, the whole goal will be retried if $\text{Plan}(P)$ fails or does not achieve ϕ_s . In addition, P needs to be executed completely, that is, a partial execution of P that would achieve ϕ_s would not work.

Consider next its plan-goal counterpart $\text{Plan}(\text{Goal}(\phi_s, P, \phi_f))$. Such a program is closer to a *classical planning problem*: plan for achieving the declarative goal ϕ_s (by using the given program P). If P cannot be solved or it is not able to achieve ϕ_s , then the whole plan-goal will fail. Thus, we are not merely interested in solving program P , but also doing it in a way that will in fact achieve the desired goal ϕ_s . In addition, program P need not be executed completely but only to the point where ϕ_s is realized. This type of behaviour on (declarative) goals does not exist in CAN. It is thus sensible, for convenience, to define a new language construct $\text{Plan}(\phi, P) \stackrel{\text{def}}{=} \text{Plan}(\text{Goal}(\phi, P, \text{false}))$: plan for achieving ϕ using P . Notice that $\text{Plan}(\phi, P)$ is close but more expressive than program $(P; ?\phi)$ —the latter requires P to be *completely* executed.

We summarize the differences between goal-plan and plan-goal programs as follows:

$\text{Goal}(\phi_s, \text{Plan}(P), \phi_f)$	$\text{Plan}(\text{Goal}(\phi_s, P, \phi_f))$
BDI execution driven $\checkmark$	classical planning driven $\checkmark$
P may achieve ϕ_s	P must achieve ϕ_s $\checkmark$
P may be re-tried $\checkmark$	P is never re-tired
P must execute fully	P may execute partially $\checkmark$
Fails if ϕ_f holds	Fails if ϕ_f holds

What we would like is a program that combines the advantages of both versions. That is, we may want to deliberate and act towards a clear goal ϕ_s using a program P and be committed to the goal until it is either achieved or impossible. Interestingly, program $\text{PlanGoal}(\phi_s, P, \phi_f) \stackrel{\text{def}}{=} \text{Goal}(\phi_s, \text{Plan}(\text{Goal}(\phi_s, P, \phi_f)), \phi_f)$ has such properties and it will (a) always execute towards goal ϕ_s ; (b) terminate when the goal is achieved or is impossible; and (c) re-try and insist on program P if necessary.

THEOREM 25. Let $C_0 = \langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}, \mathcal{A}, \Gamma_0 \cup \{P_0^X\} \rangle$ be an agent configuration with $X \in \{a, b\}$, where (P is a user program)

$$P_0^a = \text{Plan}(\text{Goal}(\phi_s, P, \phi_f)), \quad P_0^b = \text{Goal}(\phi_s, \text{Plan}(P), \phi_f),$$

Let $E = C_0 \dots C_n$ be an execution of C_0 such that P_0^X is currently executing in E (i.e., $P_n^X \neq \epsilon$). Next, let us consider $C_0^* = \langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}, \mathcal{A}, \Gamma_0 \cup \{P_0^*\} \rangle$, where $P_0^* = \text{Goal}(\phi_s, P_0^a, \phi_f)$.

Then, there exists an execution $E^* = C_0^* \dots C_m^*$ of C_0^* such that (a) $|E|$ and $|E^*|$ are equivalent modulo intentions P_0^X and P_0^* ; and (b) intention P_0^X has been successfully executed in E iff intention P_0^* has been successfully executed in E^* .

PROOF. Let $P^a = \text{Plan}(\text{Goal}(\phi_s, P, \phi_f))$ and let $P^b = \text{Goal}(\phi_s, \text{Plan}(P), \phi_f)$. Suppose $E = C_0, \dots, C_n$ is an execution of $C_0 = \langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}, \mathcal{A}, \Gamma \cup \{P^a\} \rangle$ such that P^a is currently executing in E . Then, there exists an execution $E' = C'_0, \dots, C'_n$ of $C'_0 = \langle \mathcal{N}, \Lambda, \Pi, \mathcal{B}, \mathcal{A}, \Gamma \cup \{P^b\} \rangle$ such that E' and E are equivalent and

1. $\Gamma_n = \Gamma^* \cup \{P_n^a\}$ iff $\Gamma'_n = \Gamma^* \cup \{P_n^b\}$;
2. if $P_n^a = \text{Plan}(\text{Goal}(\phi_s, P' \triangleright P, \phi_f))$, then $P_n^b = \text{Goal}(\phi_s, \text{Plan}(P') \triangleright \text{Plan}(P), \phi_f)$;
3. if $P_n^a = \text{Plan}(\text{nil})$, then $P_n^b = \text{nil}$;
4. if $\langle \mathcal{B}_n, \mathcal{A}_n, P_n^a \rangle \not\stackrel{\text{bdi}}{\longrightarrow}$ and $P_n^a \neq \text{nil}$, then $P_n^a = \text{Plan}(\text{Goal}(\phi_s, P' \triangleright P, \phi_f))$ and $\langle \mathcal{B}_n, \mathcal{A}_n, \text{Plan}(P') \rangle \not\stackrel{\text{bdi}}{\longrightarrow}$.

The proof relies on a few lemmas; the most important of them states that for any $\mathcal{B}, \mathcal{A}, P_o, \phi_s, \phi_f$, and $P \neq \text{nil}$: if $\langle \mathcal{B}, \mathcal{A}, \text{Plan}(\text{Goal}(\phi_s, P \triangleright P_o, \phi_f)) \rangle$ can make a bdi-type transition to $\langle \mathcal{B}', \mathcal{A}', \text{Plan}(\text{Goal}(\phi_s, P' \triangleright P_o, \phi_f)) \rangle$, then it is the case that $\langle \mathcal{B}, \mathcal{A}, \text{Goal}(\phi_s, \text{Plan}(P) \triangleright \text{Plan}(P_o), \phi_f) \rangle$ can make a bdi-type transition to $\langle \mathcal{B}', \mathcal{A}', \text{Goal}(\phi_s, \text{Plan}(P') \triangleright \text{Plan}(P_o), \phi_f) \rangle$. $\square$

The reader may wonder if it is possible to combine the features of both type of programs. That is, one may want to deliberate on, possibly partially, doing P in a way that will indeed achieve ϕ_s while allowing for the possibility of re-trying P if an external event interferes with the solution found at planning time. Interestingly, the program we are looking for is $\text{Goal}(\phi_s, \text{Plan}(\text{Goal}(\phi_s, P, \phi_f)), \phi_f)$. To show this, we first need to abstract away from some irrelevant internal steps in an agent execution. If $E = C_0 \cdot \dots \cdot C_n$, $n \geq 0$, is an agent execution, then $\overline{C_0 \cdot \dots \cdot C_n}$ is the sequence obtained from E by deleting all elements C_j of the sequence such that the $C_j \implies C_{j+1}$ step used in E is due to the agent derivation rule A_{step} and either of the basic derivation rules G_I or Plan_t .