HDDL: An Extension to PDDL for Expressing Hierarchical Planning Problems

Daniel Höller,[1] **Gregor Behnke,**[1] **Pascal Bercher,**[1*] **Susanne Biundo,**[1]
Humbert Fiorino,[2] **Damien Pellier,**[2] and **Ron Alford**[3]

[1]Institute of Artificial Intelligence, Ulm University, 89081 Ulm, Germany
[2]University Grenoble Alpes, LIG, F-38000 Grenoble, France
[3]The MITRE Corporation, McLean, Virginia, USA
{daniel.hoeller, gregor.behnke, susanne.biundo}@uni-ulm.de, pascal.bercher@alumni.uni-ulm.de
{humbert.fiorino, damien.pellier}@imag.fr
ralford@mitre.org

Abstract

The research in hierarchical planning has made considerable progress in the last few years. Many recent systems do not rely on hand-tailored advice anymore to find solutions, but are supposed to be domain-independent systems that come with sophisticated solving techniques. In principle, this development would make the comparison between systems easier (because the domains are not tailored to a single system anymore) and – much more important – also the integration into other systems, because the modeling process is less tedious (due to the lack of advice) and there is no (or less) commitment to a certain planning system the model is created for. However, these advantages are destroyed by the lack of a common input language and feature set supported by the different systems. In this paper, we propose an extension to PDDL, the description language used in non-hierarchical planning, to the needs of hierarchical planning systems.

1 Introduction

Much progress has recently been made in hierarchical planning (Bercher, Alford, and Höller 2019). Novel systems based on the traditional, search-based techniques have been introduced (Bit-Monnot, Smith, and Do 2016; Shivashankar, Alford, and Aha 2017; Bercher et al. 2017; Höller et al. 2018; 2019), but also new techniques like the translation to STRIPS/ADL (Alford, Kuter, and Nau 2009; Alford et al. 2016a). Others have been revisited, e.g. the translation to propositional logic (Behnke, Höller, and Biundo 2018a; 2018b; 2019a; 2019b; Schreiber et al. 2019). In contrast to many earlier systems, these systems can be considered to be domain-independent, they do not rely on hand-tailored advice to solve problems, but only on their solving techniques.

Even though the systems share the basic idea of being *hierarchical* planning approaches, the feature set supported by the different systems is manifold. Bit-Monnot, Smith, and Do (2016) focus, e.g., on advanced support for temporal planning, but lack the support for recursion; several systems are restricted to models that do not include partial ordering between tasks (Alford, Kuter, and Nau 2009; Behnke,

Höller, and Biundo 2018a; Schreiber et al. 2019); and some, like the one by Shivashankar, Alford, and Aha (2017) define an entirely new type of hierarchical planning problems.

Even systems restricted to the maybe best-known and most basic hierarchical formalism, *Hierarchical Task Network* (HTN) planning, do not share a common input language, though the differences between the input languages are sometimes rather small. To the best of our knowledge, the hierarchical language introduced for the first International Planning Competition (IPC) by McDermott et al. (1998) is not supported by any recent system.

The lack of a common language has several consequences for the field. First, it makes comparison between systems tedious due to the translation process. It has been identified as main problem for a hierarchical track at the IPC 2020 (Behnke et al. 2019). Second – and even more important – it makes the use of hierarchical planning from a practical perspective laborious, because it is not possible to model a problem at hand and try which system performs best on it. Selecting the system in beforehand requires much insight into the systems. A common input language would make the comparison of systems easier, it could foster a common set of supported features and result in a common benchmark set.

In this paper, we propose the *Hierarchical Domain Definition Language* (HDDL) as common input language for hierarchical planning problems. It is widely based on the input language of PANDA, the framework underlying the planning systems by Bercher et al. (2017), Höller et al. (2018; 2019), and Behnke, Höller, and Biundo (2018a; 2019a; 2019b), i.e. there is a large set of planners based on (1) search in plan space, (2) progression search and (3) a translation to SAT already supporting HDDL. It is defined as an extension of the STRIPS fragment (language level 1) of the PDDL2.1 definition (Fox and Long 2003). To concentrate on a set of features shared by many systems, we restricted the language to basic HTN planning. However, we hope that this is just the starting point for further language extensions similar to the first PDDL version in classical planning.

HDDL will be the standard input language for the track on hierarchical planning at the IPC 2020. The corresponding website[1] provides supplementary material like grammar

---

[1]ipc-2020.hierarchical-task.net

definitions or translation tools; but also benchmark sets.

We first introduce a lifted HTN formalism from the literature. Then we introduce the new language, give its syntax and meaning, discuss design choices and the differences and similarities to approaches from the literature, namely PDDL1.2 (McDermott et al. 1998), SHOP(2) (Nau et al. 2003), ANML (Smith, Frank, and Cushing 2008), HPDDL (Alford et al. 2016a), GTOHP (Ramoul et al. 2017), HTN-PDDL (González-Ferrer, Fernández-Olivares, and Castillo 2009), and HATP (de Silva, Lallement, and Alami 2015). Afterwards we give a full EBNF syntax definition based on the definition of PDDL2.1 and discuss every extension and change. We conclude with a short outlook. Before we come to the new language definition, we want to establish a common semantics definition close to the given input language and therefore introduce a lifted HTN planning formalism as known from the literature (Section 2).

2 Lifted HTN Planning

In this section we formally define the problem class that HDDL can describe, i.e., HTN planning as described by Ghallab, Nau, and Traverso (2004). We therefore extend the formalism by Alford, Bercher, and Aha (2015a; 2015b).

Our *lifted* formalism is based on a quantifier-free first-order predicate logic $\mathcal{L} = (P, T, V, C)$. All sets mentioned in the following are finite. T is a set of *type symbols* and V a set of typed variable symbols. P is a set of *predicate symbols*, each having an arity. The arity defines its number of parameter variables (taken from V). C is a set of typed constants, syntactic representations of the objects in the real world. Notice that a single constant can have several types, e.g. *truck* and *vehicle* to support a type hierarchy.

There are two distinct kinds of tasks: primitive tasks (also called actions) and compound tasks (also called abstract tasks). Each task is given by its name followed by a parameter sequence. For instance, a task for driving from a source location $?ls$ to a destination location $?ld$ is given by the first-order atom $drive(?ls, ?ld)$. We use the expressions *task* and *task names* synonymously.

The basic data structure in HTN planning is a *task network*, which represents a partially ordered multi-set of tasks.

Definition 1 (Task Network). *A task network tn over a set of* task *names X (first-order atoms) is a tuple $(I, \prec, \alpha, VC)$ with the following elements:*
1. *I is a (possibly empty) set of* task *identifiers.*
2. *$\prec$ is a strict partial order over I.*
3. *$\alpha : I \to X$ maps task identifiers to task names.*
4. *VC is a set of variable constraints. Each constraint can bind two task parameters to be (non-)equal and it can constrain a task parameter to be (non-)equal to a constant, or to (not) be of a certain type.*

Task identifiers (ids) are arbitrary symbols which serve as place holders for the actual tasks they represent. We use ids because a task can occur multiple times within the same network. We call a task network *ground* if all parameters are bound to (or replaced by) constants from C.

An *action* a is a tuple (*name*, *pre*, *eff*), where *name* is its *task name*, a first-order atom such as $drive(?ls, ?ld)$ consist-

ing of the (actual) name followed by a list of typed parameter variables. *pre* is its *precondition*, a first-order formula over literals over $\mathcal{L}$'s predicates. *eff* is its *effect*, a conjunction of literals over $\mathcal{L}$'s predicates. We define *eff*$^+$ and *eff*$^-$ as the sets of atoms occurring non-negated/negated in *eff*. All variables used in *pre* and *eff* have to be parameters of *name*. We also write *name(a)*, *pre(a)*, and *eff(a)* to refer to these elements. We require that action names *name(a)* are unique.

A *compound task* is simply a task name, i.e., an atom. Its purpose is not to induce state transition, but to reference a pre-defined mapping to one or more task networks by which that compound task can be refined. This mapping is given by a set of *(decomposition) methods M*. A method $m \in M$ is a triple (c, tn, VC) of a compound task name c, a task network tn, and a set of variable constraints VC that allow to (co)designate parameters of c and tn.

Definition 2 (Planning Domain). *A planning domain $\mathcal{D}$ is a tuple $(\mathcal{L}, T_P, T_C, M)$ defined as follows:*
- *$\mathcal{L}$ is the underlying predicate logic.*
- *T_P and T_C are sets of primitive and compound tasks.*
- *M is a set of decomposition methods with compound tasks from T_C and task networks over the names $T_P \cup T_C$.*

The domain implicitly defines the set of all states S, being defined over all subsets of all ground predicates.

Definition 3 (Planning Problem). *A planning problem $\mathcal{P}$ is a tuple $(\mathcal{D}, s_I, tn_I, g)$, where:*
- *$\mathcal{D}$ is a planning domain.*
- *$s_I \in S$ is the initial state, a ground conjunction of positive literals over the predicates.*
- *tn_I is the initial task network (not necessarily ground).*
- *g is the goal description, a first-order formula over the predicates (not necessarily ground).*

HTN planning is *not* about achieving a certain state-based goal, so it makes perfect sense to specify no goal formula at all. We added it to be closer to the PDDL specification. Supporting it is more convenient in case it is needed, it has clearly defined semantics, and causes no problems to systems that do not support it directly (since it can be compiled away, see Geier and Bercher, 2011).

Solutions in HTN planning are executable, ground, primitive task networks that can be obtained from the problem's initial task network via applying methods, adding ordering constraints, and grounding.

Lifted problems are a compact representation of their ground instantiations that are up to exponentially smaller (Alford, Bercher, and Aha 2015a; 2015b). We define the semantics of a lifted problem (i.e. the set of solutions) in terms of the standard semantics of its ground instantiation. For details on the grounding process, we refer the reader to Alford, Bercher, and Aha (2015a). There are currently only two publications devoted to grounding in more detail (see Ramoul et al., 2017, and Behnke et al., 2020). We now give the following definitions based on a *ground* problem. Note that we do not need to represent variable constraints anymore since their constraints are represented in the grounding.

Given ground problems we define *executability* of task networks as follows. Let A be the set of ground actions obtained from T_P. An action $a \in A$ is called executable in

a state $s \in S$ if and only if $s \models pre(a)$. The state transition function $\gamma : S \times A \to S$ is defined as follows: If a is executable in s, then $\gamma(s, a) = (s \setminus eff^-(a)) \cup eff^+(a)$, otherwise $\gamma(s, a)$ is undefined. The extension of γ to action sequences, $\gamma^* : S \times A^* \to S$ is defined straightforwardly.

Definition 4 (Executability). *A task network $tn = (I, \prec, \alpha)$ is called executable if and only if there is a sequence $i_1, \ldots, i_n$ of its task identifiers with $n = |I|$, such that $\alpha(i_1), \ldots, \alpha(i_n)$ is executable in s_I.*

Decomposition is used to transform one task network into another. When a task is decomposed using a method, it is removed from the network, the method's subtasks are added, and they inherit the ordering relations that held for the abstract task. Formally it is defined as:

Definition 5 (Decomposition). *Let $m = (c, (I_m, \prec_m, \alpha_m))$ be a decomposition method, $tn_1 = (I_1, \prec_1, \alpha_1)$ a task network, and $I_m \cap I_1 = \emptyset$ (the latter can be achieved by renaming). Then, m decomposes a task identifier $i \in I_1$ into a task network $tn_2 = (I_2, \prec_2, \alpha_2)$ if and only if $\alpha_1(i) = c$ and*

$$I_2 = (I_1 \setminus \{i\}) \cup I_m$$
$$\prec_2 = (\prec_1 \cup \prec_m \cup \ \{(i_1, i_2) \in I_1 \times I_m \mid (i_1, i) \in \prec_1\} \cup$$
$$\{(i_1, i_2) \in I_m \times I_1 \mid (i, i_2) \in \prec_1\})$$
$$\setminus \{(i', i'') \in I_1 \times I_1 \mid i' = i \text{ or } i'' = i\}$$
$$\alpha_2 = (\alpha_1 \cup \alpha_m) \setminus \{(i, c)\}$$

Now we can formally define the solution criteria.

Definition 6 (Solutions). *Let $\mathcal{P} = (\mathcal{D}, s_I, tn_I, g)$ be a planning problem with $\mathcal{D} = (\mathcal{L}, T_P, T_C, M)$ and $tn_S = (I_S, \prec_S, \alpha_S)$. tn_S is a solution to an HTN planning problem P if and only if the following conditions hold:*

- *There is a sequence of decompositions from tn_I to $tn = (I, \prec, \alpha)$, such that $I = I_S$, $\prec \subseteq \prec_S$, and $\alpha = \alpha_S$ and*

- *tn_S is primitive and has an executable action linearization leading to a state $s \models g$.*

3 The Hierarchical Domain Definition Language

In this section we explain our extensions to the PDDL definition based on a transport domain. To keep the example simple, the domain includes only a single transporter that has to deliver one or more packages. For each new language element we introduce its syntax and meaning and discuss the way it is modeled in other input languages.

The predicate and type definition is the same as in PDDL:

```
1 (define (domain transport)
2   (:types location package - object)
3   (:predicates
4       (road ?l1 ?l2 - location)
5       ...)
```

All other languages except for HATP (de Silva, Lallement, and Alami 2015) use the same theoretical model of objects and predicates as PDDL. HATP models its objects in an object-oriented way instead and further allows for SAS+ variables (Bäckström and Nebel 1995) in the input language.
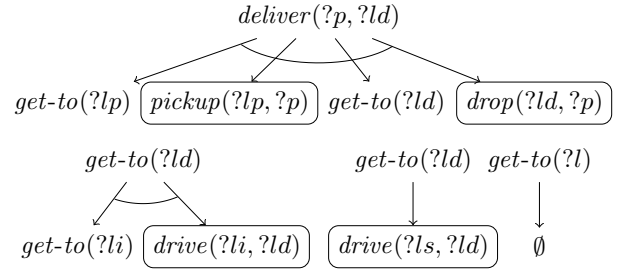

Figure 1: The method set of a simple transport domain. Actions are given as boxed nodes, abstract tasks are unboxed. All methods are totally ordered. There exists a smaller, equivalent model. However, the model has been created this way to illustrate the different language features. Two methods contain additional constraints that are not given in the figure but are discussed in the respective sections.

The full method set of the domain is illustrated in Figure 1. Each method will be discussed in this section, but first we introduce how abstract tasks are defined.

3.1 Abstract Task Definition

The domain contains two abstract tasks *deliver* and *get-to*. We propose to include an explicit definition of abstract tasks as it is the case for actions. HPDDL (Alford et al. 2016a) also defines abstract tasks explicitly, albeit with a slightly different syntax. Both ANML (Smith, Frank, and Cushing 2008) and HTN-PDDL (González-Ferrer, Fernández-Olivares, and Castillo 2009) require an explicit declaration of abstract tasks and their parameter types as well, but here the declaration is not separated from other elements of the domain as both declare methods together with their abstract tasks.

Some description languages for HTN problems define abstract tasks only in an implicit way by their use in methods. This includes the language used by SHOP and SHOP2 (Nau et al. 2003), PDDL1.2 (McDermott et al. 1998), HATP, as well as GTOHP (Ramoul et al. 2017). Implicit definition of the compound task set has also been chosen in some formal definitions of hierarchical problem classes (Alford, Bercher, and Aha 2015a; 2015b). However, this can be very cumbersome when debugging domains. If the modeler forgot to define a specific primitive task, the domain will still be valid, as it would be interpreted as an abstract task.

Another problem with such a definition is that the argument types are defined implicitly, namely as those with which the task can be instantiated via any method. The language of GTOHP further does not allow for using different types (that share a common ancestor in the type hierarchy) to be used for the same task. For example, there might be different methods for the *deliver* task, depending on the type of transported package, e.g. *regularPackage* and *valuablePackage* (the latter requiring an armored transporter). We assume that *regularPackage* and *valuablePackage* are disjunct types, sharing a common super-type *package*, which would be the correct parameter type for the *deliver* task. If the type is not declared explicitly,

the planner can either reject the domain, as GTOHP does, or infer possible types of the arguments of an abstract task.

Declaring abstract tasks and their parameter types explicitly is also in line with the design choices of PDDL. Similar to abstract tasks, PDDL could omit the explicit definition of predicates as their types could be inferred from their usages. This is however discouraged from a modeling point of view.

Omitting the distinct definition of tasks and methods would also mean a significant deviation from the contemporary theoretical work on HTN planning. It could hinder further language extensions like annotating abstract tasks with constraints, e.g. preconditions and effects, as done by a couple of systems (see e.g. the survey by Bercher et al., 2016).

Here is the abstract task definition for the example:

```
6  (:task deliver :parameters (?p - package
       ?l - location))
7  (:task get-to :parameters (?l - location))
```

3.2 Decomposition Method Definition

There is only a single method in the model to decompose deliver tasks (given at the top of Figure 1). It decomposes the task into four ordered subtasks: getting to the package, picking it up, getting to its final position, and dropping the package. The definition in HDDL could look like this:

```
8   (:method m-deliver
9     :parameters (?p - package
         ?lp ?ld - location)
10    :task (deliver ?p ?ld)
11    :ordered-subtasks (and
12      (get-to ?lp)
13      (pick-up ?ld ?p)
14      (get-to ?ld)
15      (drop ?ld ?p)))
```

The method definition starts with the method's name. We decided to give parameters explicitly (line 9). This allows e.g. to restrict the types used in the subtasks and the decomposed task to subtypes of the original task parameters.

We assume these parameters to be a superset of all parameters used in the entire method definition. The parameter definition is followed by the specification of the abstract task decomposed by the method and its parameters (line 10).

The same syntactical structure is used by HPDDL. In contrast, ANML, PDDL1.2, HATP and HTN-PDDL aggregate all decomposition methods belonging to a certain abstract task and define them as part of the definition of the task. As such, the variables that are declared as the arguments of an abstract task are automatically variables in a methods' task network. All of them type variables in methods explicitly.

In GTOHP's language, methods don't have names but are identified via the abstract task they refine.

In SHOP, all variables inside a method are defined implicitly by their usage as parameters. The definition of a SHOP method starts e.g. with :method followed by an abstract task and its parameters – which if they are variables are automatically declared as new (untyped) variables. The same holds for variables that only occur as parameters of a method's subtasks. GTOHP and HTN-PDDL follow this pattern, but enforce that the parameters of the abstract task

are typed, i.e. declared explicitly. Their languages however do not allow specifying the types of variables that occur in the method that are not parameters of the abstract task. Declaring the variables is, again, in line with the PDDL standard and e.g. done the same way in actions. We think it less error-prone. When the modeler explicitly defines the variables and their types, the system can check the compatibility of types and warn the modeler when undeclared variables are used.

Subtasks and Ordering The subtasks of the decomposition method are given afterwards (starting in line 11). We decided to have two keywords to start the definition :ordered-subtasks (as given here) and :subtasks (which we will show in the next method definition). When the :ordered-subtasks keyword is used, the given list of subtasks is supposed to be totally ordered. HPDDL uses the keyword :tasks, which might cause errors if mixed up with the :task keyword. Since GTOHP does only support totally-ordered HTN planning problems, its language only allows for specifying sequences of actions with the keyword :expansion.

In the subtask section, all abstract tasks and actions defined in the domain can be used as subtasks (and only these). Variables defined in the method's parameters and constants from the domain can be used as parameters.

The *get-to* task from our example domain is again abstract and can be decomposed by using one of the three methods given at the bottom of Figure 1. We start with the left one that is used when there is no direct road connection. Then the transporter needs to go to the final location $?ld$ via some intermediate location $?li$. Therefore the method decomposes the task into another abstract *get-to* task, followed by a *drive* action with the destination location $?ld$.

```
16   (:method m-drive-to-via
17     :parameters (?li ?ld - location)
18     :task (get-to ?ld)
19     :subtasks (and
20       (t1 (get-to ?li))
21       (t2 (drive ?li ?ld)))
22     :ordering (and
23       (< t1 t2)))
```

Line 19 shows the mentioned :subtask definition that allows for partially ordered tasks. The task definition contains ids (named t1 and t2) that can be used to define ordering constraints (line 22). They consist of a list of ordering constraints between subtasks. However, in the given example the resulting ordering is, again, a total order (and is just defined that way to demonstrate this kind of definition).

HPDDL uses the same keyword, but with a slightly different syntax. The format omits the and and the < signs. We would argue that our notation is more readable. As stated above, GTOHP cannot specify partial orders. ANML is primarily designed for temporal domains and uses a temporal syntax, e.g. end(t1) < start(t2). SHOP2 and HTN-PDDL use a different approach to represent the task ordering. Instead of specifying a set of constraints they require specifying the order as a single expression. This is a nested definition that contains

two constructors: `ordered` and `unordered`. In SHOP2, e.g. `((:unordered (t1 t2) t3) t4)` corresponds to the ordering constraints `(< t1 t2)`, `(< t2 t4)`, and `(< t3 t4)`. This construction cannot express all possible orderings. Consider an ordering over five task identifiers `t1`, ..., `t5`, where `(< t1 t4)`, `(< t2 t4)`, `(< t2 t5)`, and `(< t3 t5)`. This ordering cannot be expressed with SHOP's nested ordered/unordered constructs. PDDL1.2 also uses this mode as a default, but does also allow for an order specification as we and HPDDL do. Notably PDDL1.2 intertwines the definition of a method's subtasks and the definition of their order. The syntax of PDDL1.2 to specify the contents of methods and the order of tasks in them is somewhat convoluted and not easily readable, so we have not adopted their syntax.

HATP uses a programming-language-style syntax for the encoding of methods. It provides explicit means to determine the order in which groundings of methods should be explored during progression search. HATP's syntax for methods allows for specifying partial order, but its semantics is different from standard HTN planning. A HATP method containing partial order is interpreted as multiple totally-ordered method, one for each linearization of the given partial order. This allows for a more compact representation, but prohibits task interleaving.

HDDL – like HPDDL, SHOP2, HTN-PDDL, and ANML – only allows to specify a fixed *set* of ordering constraints. Notably, the HTN planner UMCP (Erol, Hendler, and Nau 1994) allows for *arbitrary formulae* that specify the orderings. E.g. they allow to specify an ordering $(t_1 \prec t_2) \Rightarrow (t_3 \prec t_4)$. We have not included such a generic means to formulate ordering constraints as they do not seem to be used and supported by any current HTN planning system.

Method Preconditions   A common feature of many HTN planning systems is the possibility of specifying state-based preconditions for methods as supported by the SHOP2 system. The feature is somewhat problematic: First, because it is (at least from our experience) usually used to guide the search and thus often breaks with the philosophy of PDDL to specify a model that does not include advice.

The second problem is the question how it is realized in the planning systems. Some systems introduce a new primitive task that holds the method's preconditions. It is added to the method and placed before all other tasks in the subtask network. Consider a totally ordered domain (i.e., the subtasks of all methods and the initial task network are totally ordered): here, the action is executed directly before the other subtasks of the method and the position where the preconditions are checked is fine. But consider a partially ordered domain: here, the newly introduced action is not necessarily placed directly before the other subtasks, but we just know that it is placed somewhere before, i.e., the condition did hold at some point before the other tasks are executed, but may have changed meanwhile.

Other systems check the precondition exactly before the first action resulting from a subtask of the method. Here, it is fully specified when the precondition is checked, but the system needs to (natively) support this feature, because a compilation is not easily possible.

Though we are aware of these problems, the feature is often used and thus we integrated it. However, to make the burden of supporting it as small as possible, we assume the compilation semantics as given above and consider the other definition a feature for future extensions.

The preconditions are defined as follows:

```
24    (:method m-already-there
25       :parameters (?l - location)
26       :task (get-to ?l)
27       :precondition (tAt ?l)
28       :subtasks ())
```

Here the method can be applied in a state where the transporter is already located at its destination. The given method has therefore no subtasks, but still has to assure that the transporter is at its destination.

Method preconditions are typically featured in languages expressing HTNs. HPDDL uses the same syntax we are proposing. GTOHP uses, as noted above, a separate `:constraints` section, where the method precondition has to be specified as a `before` constraint. This is (presumably) to allow for other state constraints later on. PDDL1.2 also features method preconditions, but they are specified as part of the task network. In ANML, there is no explicit means for writing down method preconditions, but they can be encoded into the state constraints allowed by ANML.

There is a strong contrast between what can be expressed in SHOP[2] and all other HTN formats. In SHOP, several methods for the same abstract task can be arranged in a single method declaration, each featuring its own method precondition. For the i^{th} method to be usable, it is not sufficient that its precondition is satisfied. In addition, the preconditions of all previous methods have to be not satisfied as well. Thus SHOP's method preconditions are in essence a chain of if-else constructs. This structure can be compiled into several individual methods with preconditions. In case one of the preconditions contains an existential quantifier (or in SHOP's case a free variable) this leads to universal quantified preconditions in the methods after it. Nevertheless we propose to drop the ability to use such if-else chains, most notably since none of the newer languages supports it. Further, this kind of if-else is essentially a means to guide a depth-first search planner in an efficient way.

In addition to method preconditions, HPDDL features method effects, which are modeled after SHOP2's assert and retract functionality. Method effects are executed in the state in which the method preconditions are evaluated. As far as we know, their formal semantics is not defined in any publication. We propose to drop this feature (at least for the given definition intended to be the core language), since it is not commonly used. Note that even without method effects in the description language, they can still be simulated with additional actions in the methods' definitions (at least in the compilation-based definition as given above).

---

[2]This also applies to HTN-PDDL, it uses a similar syntax. The description is unfortunately not explicit on the critical point in semantics (González-Ferrer, Fernández-Olivares, and Castillo 2009).

General Constraints   Sometimes it might be useful to define constraints in a method, e.g. on its variables or sorts. This is demonstrated in the following example where the transporter's source position must be different from its destination.

```
29    (:method m-direct
30      :parameters (?ls ?ld - location)
31      :task (get-to ?ld)
32      :constraints
33        (not (= ?li ?ld))
34      :subtasks (drive ?ls ?ld))
```

We are aware that PDDL allows for variable constraints in the precondition of actions. Due to consistency we also allow this when method preconditions are specified. However, many HTN models are defined without methods that have preconditions and we think it is not intuitive to specify a `precondition` section solely to define variable constraints. Furthermore, we think that other constraints apart from simple variable constraints might be added to the standard, e.g. state constants like used by Erol, Hendler, and Nau (1994). Therefore we integrated a constraint section to the method definition (line 32f.) though our current definition only allows for equality and inequality constraints.

HPDDL places the variable constraints of a decomposition method into its preconditions. In addition to equality and inequality it features type constraints, where e.g. `(valuablePackage ?p)` is the constraint that `?p` belongs to the type `valuablePackage`. GTOHP allows for equality and inequality constraints that are also within the `:constraints` section, but are located in a separate `before` block. In SHOP's syntax, variable constraints have to be compiled into method preconditions referring to predicates for the individual types and an explicitly declared `equals` predicate. ANML also allows for variable constraints that can be declared freely inside a method.

3.3   Action Definition

We left the action definition unchanged compared to the PDDL standard we build on.

```
35    (:action drive
36      :parameters (?l1 ?l2 - location)
37      :precondition (and
38        (tAt ?l1)
39        (road ?l1 ?l2))
40      :effect (and
41        (not (tAt ?l1))
42        (tAt ?l2)))
43    ...)
```

3.4   Problem Definition

The problem definition is slightly adapted to represent the additional elements necessary for HTN planning (line 6).

```
1 (define (problem p)
2 (:domain transport)
3 (:objects
4  city-loc-0 city-loc-1 city-loc-2 -
       location
5  package-0 package-1 - package)
```

```
6  (:htn
7   :tasks (and
8    (deliver package-0 city-loc-0)
9    (deliver package-1 city-loc-2))
10   :ordering ()
11   :constraints ())
12  (:init
13   (road city-loc-0 city-loc-1)
14   (road city-loc-1 city-loc-0)
15   (road city-loc-1 city-loc-2)
16   (road city-loc-2 city-loc-1)
17   (at package-0 city-loc-1)
18   (at package-1 city-loc-1)))
```

The section starts with a keyword that specifies the problem class, here it starts with `:htn`. However, there are several other problem classes in hierarchical planning, e.g. HTN planning with task insertion. An overview of hierarchical problem classes can be found in the survey by Bercher, Alford, and Höller (2019). Some problem classes are even syntactically equivalent to HTN planning and only differ in their solution criteria. By making the specification of the problem class explicit, extensions to the language can easily add new classes. The initial task network is nested in this section. It has the same form as the methods' subtask networks. The other description languages for HTN planning also allow for a similar definition of the initial plan. Again, all of them use a slightly different syntax to describe them.

In the given example, the planning process is started with two unordered *deliver* tasks, one for each package.

In the original PDDL standard, the domain designer has to specify a state-based goal. This is often not specified in HTN planning, so we made it optional.

4   Full Syntax Definition

The following is defined as close as possible to the STRIPS part (i.e., language level 1) of PDDL 2.1 by Fox and Long (2003). Large parts are **identical** to their definition. Changes and extensions are discussed in the following.

The domain definition has been extended by definitions for compound tasks (line 6) and methods (line 7).

```
1 <domain> ::= (define (domain <name>)
2    [<require-def>]
3    [<types-def>]:typing
4    [<constants-def>]
5    [<predicates-def>]
6    <comp-task-def>*
7    <method-def>*
8    <action-def>*)
```

The definition of the basic elements is nearly unchanged.

```
9 <require-def> ::=
      (:requirements <require-key>+)
10 <require-key> ::= ...
11 <types-def> ::= (:types
      <typed list (name)>* <base-type>+)
12 <base-type> ::= <name>
13 <constants-def> ::=
      (:constants <typed list (name)>)
14 <predicates-def> ::= (:predicates
      <atomic-formula-skeleton>+)
```

```
15 <atomic-formula-skeleton> ::=
      (<predicate> <typed list (variable)>)
16 <predicate> ::= <name>
17 <variable> ::= ?<name>
18 <typed list (x)> ::= x+ - <type>
      [<typed list (x)>]
19 <primitive-type> ::= <name>
20 <type> ::= (either <primitive-type>+)
21 <type> ::= <primitive-type>
```

We changed the definition of `<types-def>` (given in line 11 and 12) in combination with the definition of `<typed list (name)>` (line 18). In the PDDL2.1 standard, this can be realized by a list of names, e.g. in an untyped way. Our intention was to enforce a typed model and therefore allow for untyped elements only in the type definition. There, it is necessary to define the base type(s). In every other definition that includes `<typed list (name)>` (e.g. parameter and constant definitions), we wanted to enforce a typed list.

Abstract tasks are defined similar to actions.

```
22 <comp-task-def> ::= (:task <task-def>)
23 <task-def> ::= <task-symbol>
      :parameters (<typed list (variable)>)
24 <task-symbol> ::= <name>
```

In a standard HTN setting, methods consist of a parameter list (line 26), the abstract task they decompose (line 27), and the resulting task network (line 29).

```
25 <method-def> ::= (:method <name>
26     :parameters (<typed list (variable)>)
27     :task (<task-symbol> <term>*)
28     [:precondition <gd>]
29     <tasknetwork-def>)
```

We use the same syntax definition for method subnetworks and the initial task network. Here, the keyword `subtasks` would seem odd. Therefore the syntax also allows for the keys `tasks` and `ordered-tasks` (line 31) that are supported to be used in the initial task network. The definition includes subtasks (line 31), ordering constraints (line 32), and variable constraints (line 33) between any method parameters.

```
30 <tasknetwork-def> ::=
31     [:[ordered-][sub]tasks
          <subtask-defs>]
32     [:order[ing] <ordering-defs>]
33     [:constraints <constraint-defs>]
```

When the key `:ordered-subtasks` is used, the network is regarded to be totally ordered. In the other cases, ordering relations may be defined explicitly. This is done by including ids into the task definition that can then be referenced in the ordering definition.

The subtask definition may contain several subtasks. A single task consists of a task symbol and a list of parameters. In case of a method's subnetwork, these parameters have to be included in the method's parameters, in case of the initial task network, they have to be defined as constants in s_0 or in a dedicated parameter list (see definition of the initial task network, line 81).

```
34 <subtask-defs> ::= () | <subtask-def>
      | (and <subtask-def>+)
35 <subtask-def> ::= (<task-symbol> <term>*)
      | (<subtask-id> (<task-symbol>
      <term>*))
36 <subtask-id> ::= <name>
```

The ordering constraints are defined via the task ids. They have to induce a partial order.

```
37 <ordering-defs> ::= () | <ordering-def>
      | (and <ordering-def>+)
38 <ordering-def> ::=
      ("<" <subtask-id> <subtask-id>+)
39 <constraint-defs> ::= () | <constraint-def
      > | (and <constraint-def>+)
40 <constraint-def> ::= ()
      | (not (= <term> <term>))
      | (= <term> <term>)
```

The original action definition of PDDL has been split to reuse its body in the task definition.

```
41 <action-def> ::= (:action <task-def>
42     [:precondition <gd>]
43     [:effects <effect>])
```

We restricted the definition of preconditions and effects to level 1, i.e., the STRIPS part of the overall language.

```
44 <gd> ::= ()
45 <gd> ::= <atomic formula (term)>
46 <gd> ::=:negative-preconditions <literal (term)>
47 <gd> ::= (and <gd>*)
48 <gd> ::=:disjunctive-preconditions (or <gd>*)
49 <gd> ::=:disjunctive-preconditions (not <gd>)
50 <gd> ::=:disjunctive-preconditions (imply <gd> <gd>)
51 <gd> ::=:existential-preconditions
      (exists (<typed list (variable)>*)
              <gd>)
52 <gd> ::=:universal-preconditions
      (forall (<typed list (variable)>*)
              <gd>)
53 <gd> ::= (= <term> <term>)
54 <literal (t)> ::= <atomic formula(t)>
55 <literal (t)> ::=
      (not <atomic formula(t)>)
56 <atomic formula(t)> ::= (<predicate> t*)
57 <term> ::= <name>
58 <term> ::= <variable>
59 <effect> ::= ()
60 <effect> ::= (and <c-effect>*)
61 <effect> ::= <c-effect>
62 <c-effect> ::=:conditional-effects
      (forall (<variable>*) <effect>)
63 <c-effect> ::=:conditional-effects
      (when <gd> <cond-effect>)
64 <c-effect> ::= <p-effect>
65 <p-effect> ::=
      (not <atomic formula(term)>)
66 <p-effect> ::= <atomic formula(term)>
67 <cond-effect> ::= (and <p-effect>*)
68 <cond-effect> ::= <p-effect>
```

The problem definition includes the initial task network (line 73). Since a state-based goal definition is often not in-

cluded in HTN planning, we made the goal definition optional (line 75).

```
69 <problem> ::= (define (problem <name>)
70     (:domain <name>)
71     [<require-def>]
72     [<p-object-declaration>]
73     [<p-htn>]
74     <p-init>
75     [<p-goal>])
76 <p-object-declaration> ::=
        (:objects <typed list (name)>)
77 <p-init> ::= (:init <init-el>*)
78 <init-el> ::= <literal (name)>
79 <p-goal> ::= (:goal <gd>)
```

The initial task network contains the definition of the problem class (line 80). In this first definition we only included standard HTN planning, but we integrated this definition to allow for other classes, e.g. HTN planning with task insertion.

```
80 <p-htn> ::= (<p-class>
81     [:parameters (<typed list (variable)>)]
82     <tasknetwork-def>)
83 <p-class> ::= :htn
```

5   Discussion

We consider the language proposed in this paper as a first step towards a standardized language for hierarchical planning problems and hope that it helps to find a minimal set of features supported by the diverse systems.

First of all, we think it is important to remain as close as possible to PDDL and to reuse its features to allow domain designers to create both hierarchical and non-hierarchical problems with minimal learning effort. Then, we must decide which features have to be at the core of the language, and which ones are secondary and possibly could be ignored. This is especially important to establish a competition to compare the performance of different systems (as recently proposed by Behnke et al., 2019).

A feature that was present in the early HTN formalisms (e.g. the one by Erol, Hendler, and Nau, 1994) is the possibility to define more elaborated constraints in task networks. Recent work in hierarchical planning was not based on such a rich definition language, but on rather minimalistic formalisms like the one introduced by Geier and Bercher, (2011). In this first definition we only included very basic constraints: ordering constraints, variable constraints, and method preconditions. However, we think that a constraint set as given in PDDL3 might be a nice extension beneficial for domain designers. To foster the application in real world domains, it may be necessary to integrate support for numbers and time. Since our definition builds on the PDDL2.1 definition, at least the extension of the language in that direction could easily be done. Another possible extension is the support for preconditions and effects in the definition of abstract tasks (see Bercher et al. (2016) for an overview).

Besides new features, it might be interesting to include new problem classes like *HTN planning with task insertion*, *decompositional planning*, or *HGN planning*, which comes

with the ability to decompose not tasks, but also goals (Shivashankar et al. 2012) and that even has been combined with task decomposition (Alford et al. 2016b).

6   Conclusion

We propose a common description language for hierarchical planning problems. We argue that the core feature set underlying many planners from the last years is HTN planning and introduced its elements as an extension of PDDL. We defined the language in a way that can easily be extended by further features as has been done in PDDL. We introduced our novel language elements "by example" and discussed our design choices, the syntax used in related work, and the proposed meaning.

Acknowledgements

References

Alford, R.; Bercher, P.; and Aha, D. W. 2015a. Tight bounds for HTN planning. In *Proceedings of the 25th International Conference on Automated Planning and Scheduling (ICAPS)*, 7–15. AAAI Press.

Alford, R.; Bercher, P.; and Aha, D. W. 2015b. Tight bounds for HTN planning with task insertion. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, 1502–1508. AAAI Press.

Alford, R.; Behnke, G.; Höller, D.; Bercher, P.; Biundo, S.; and Aha, D. W. 2016a. Bound to plan: Exploiting classical heuristics via automatic translations of tail-recursive HTN problems. In *Proceedings of the 26th International Conference on Automated Planning and Scheduling (ICAPS)*, 20–28. AAAI Press.

Alford, R.; Shivashankar, V.; Roberts, M.; Frank, J.; and Aha, D. W. 2016b. Hierarchical planning: Relating task and goal decomposition with task sharing. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, 3022–3029. IJCAI/AAAI Press.

Alford, R.; Kuter, U.; and Nau, D. S. 2009. Translating HTNs to PDDL: A small amount of domain knowledge can go a long way. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, 1629–1634.

Bäckström, C., and Nebel, B. 1995. Complexity results for SAS+ planning. *Computational Intelligence* 11(4):625–656.

Behnke, G.; Höller, D.; Bercher, P.; Biundo, S.; Fiorino, H.; Pellier, D.; and Alford, R. 2019. Hierarchical planning in the IPC. In *Proceedings of the Workshop on the IPC (WIPC)*.

Behnke, G.; Höller, D.; Schmid, A.; Bercher, P.; and Biundo, S. 2020. On succinct groundings of HTN planning problems. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*. AAAI Press.

Behnke, G.; Höller, D.; and Biundo, S. 2018a. totSAT – Totally-ordered hierarchical planning through SAT. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, 6110–6118. AAAI Press.

Behnke, G.; Höller, D.; and Biundo, S. 2018b. Tracking branches in trees – A propositional encoding for solving partially-ordered HTN planning problems. In *Proceedings of the 30th International Conference on Tools with Artificial Intelligence (ICTAI)*, 73–80. IEEE Computer Society.

Behnke, G.; Höller, D.; and Biundo, S. 2019a. Bringing order to chaos – A compact representation of partial order in SAT-based HTN planning. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, 7520–7529. AAAI Press.

Behnke, G.; Höller, D.; and Biundo, S. 2019b. Finding optimal solutions in HTN planning – A SAT-based approach. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 5500–5508. IJCAI.

Bercher, P.; Alford, R.; and Höller, D. 2019. A survey on hierarchical planning – One abstract idea, many concrete realizations. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 6267–6275. ijcai.org.

Bercher, P.; Höller, D.; Behnke, G.; and Biundo, S. 2016. More than a name? On implications of preconditions and effects of compound HTN planning tasks. In *Proceedings of the 22nd European Conference on Artificial Intelligence (ECAI)*, 225–233. IOS Press.

Bercher, P.; Behnke, G.; Höller, D.; and Biundo, S. 2017. An admissible HTN planning heuristic. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 480–488. AAAI Press.

Bit-Monnot, A.; Smith, D. E.; and Do, M. 2016. Delete-free reachability analysis for temporal and hierarchical planning. In *Proceedings of the 22nd European Conference on Artificial Intelligence (ECAI)*, 1698–1699. IOS Press.

de Silva, L.; Lallement, R.; and Alami, R. 2015. The HATP hierarchical planner: Formalisation and an initial study of its usability and practicality. In *Proceedings of the International Conference on Intelligent Robots and Systems (IROS)*, 6465–6472. IEEE.

Erol, K.; Hendler, J. A.; and Nau, D. S. 1994. UMCP: A sound and complete procedure for hierarchical task-network planning. In *Proceedings of the 2nd International Conference on Artificial Intelligence Planning Systems (AIPS)*, 249–254. AAAI Press.

Fox, M., and Long, D. 2003. PDDL2.1: An extension to PDDL for expressing temporal planning domains. *Journal of Artificial Intelligence Research* 20:61–124.

Geier, T., and Bercher, P. 2011. On the decidability of HTN planning with task insertion. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, 1955–1961. AAAI Press.

Ghallab, M.; Nau, D. S.; and Traverso, P. 2004. *Automated Planning – Theory and Practice*. Elsevier.

González-Ferrer, A.; Fernández-Olivares, J.; and Castillo, L. 2009. JABBAH: A java application framework for the translation between business process models and HTN. In *Proceedings of the International Competition on Knowledge Engineering for Planning and Scheduling (ICKEPS)*, 28–37.

Höller, D.; Bercher, P.; Behnke, G.; and Biundo, S. 2018. A generic method to guide HTN progression search with classical heuristics. In *Proceedings of the 28th International Conference on Automated Planning and Scheduling (ICAPS)*, 114–122. AAAI Press.

Höller, D.; Bercher, P.; Behnke, G.; and Biundo, S. 2019. On guiding search in HTN planning with classical planning heuristics. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 6171–6175. ijcai.org.

McDermott, D.; Ghallab, M.; Howe, A.; Knoblock, C.; Ram, A.; Veloso, M.; Weld, D.; and Wilkins, D. 1998. PDDL – the planning domain definition language. Technical Report CVC TR-98-003/DCS TR-1165, Yale Center for Computational Vision and Control.

Nau, D. S.; Au, T.; Ilghami, O.; Kuter, U.; Murdock, J. W.; Wu, D.; and Yaman, F. 2003. SHOP2: An HTN planning system. *Journal of Artificial Intelligence Research* 20:379–404.

Ramoul, A.; Pellier, D.; Fiorino, H.; and Pesty, S. 2017. Grounding of HTN planning domain. *International Journal on Artificial Intelligence Tools* 26(5):1–24.

Schreiber, D.; Pellier, D.; Fiorino, H.; and Balyo, T. 2019. Tree-REX: SAT-based tree exploration for efficient and high-quality HTN planning. In *Proceedings of the 29th International Conference on Automated Planning and Scheduling (ICAPS)*, 382–390. AAAI Press.

Shivashankar, V.; Alford, R.; and Aha, D. W. 2017. Incorporating domain-independent planning heuristics in hierarchical planning. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, 3658–3664. AAAI Press.

Shivashankar, V.; Kuter, U.; Nau, D. S.; and Alford, R. 2012. A hierarchical goal-based formalism and algorithm for single-agent planning. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 981–988. IFAAMAS.

Smith, D.; Frank, J.; and Cushing, W. 2008. The ANML language. In *Proceedings of the Workshop on Knowledge Engineering for Planning and Scheduling (KEPS)*.