

Generalized Planning With Procedural Domain Control Knowledge

Javier Segovia and Sergio Jiménez and Anders Jonsson

Dept. Information and Communication Technologies

Universitat Pompeu Fabra

Roc Boronat 138, 08018 Barcelona, Spain

{javier.segovia, sergio.jimenez, anders.jonsson}@upf.edu

Abstract

Generalized planning is the task of generating a single solution that is valid for a set of planning problems. In this paper we show how to represent and compute generalized plans using procedural Domain Control Knowledge (DCK). We define a *divide and conquer* approach that first generates the procedural DCK solving a set of planning problems representative of certain subtasks and then compile it as callable procedures of the overall generalized planning problem. Our procedure calling mechanism allows nested and recursive procedure calls and is implemented in PDDL so that classical planners can compute and exploit procedural DCK. Experiments show that an off-the-shelf classical planner, using procedural DCK as callable procedures, can compute generalized plans in a wide range of domains including non-trivial ones, such as sorting variable-size lists or DFS traversal of binary trees with variable size.

Introduction

Domain Control Knowledge (DCK) refers to an overall strategy or suggestion of how planning problems from a certain domain should be solved. There are diverse approaches for representing DCK: *macros* (Fikes, Hart, and Nilsson 1972), *control rules* (Veloso et al. 1995), *temporal logic formulae* (Bacchus and Kabanza 1999), *HTNs* (Erol, Hendler, and Nau 1996), *reactive policies* (Yoon, Fern, and Givan 2008; De la Rosa et al. 2011), *procedural DCK* (Baier, Fritz, and McIlraith 2007) or *finite state automata* (Bonet, Palacios, and Geffner 2010).

Macro-actions (i.e. action subsequences) were among the first suggestions to speed up planning and there are several examples in the literature of computing macros (Botea et al. 2005; Coles and Smith 2007; Jonsson 2009). Incorporating macros into a planning problem can help solve it faster, but even when macros are parameterized (which is not always the case), a solution involving macros may not be applicable to other problems. Consider the navigation tasks in Figure 1, an action subsequence for reaching *G1* starting from *I* is no longer valid if we change the initial state.

A generalized plan is a single solution that is valid for a set of planning problems. Generalized plans are typically built with branching and repetition constructs which

G4				G3
			I	
G1				G2

Figure 1: Example navigation tasks in a 5×5 grid.

allows them to solve arbitrarily large problems or problems with partial observability and non-deterministic actions (Bonet, Palacios, and Geffner 2010; Pralet et al. 2010; Hu and Levesque 2011; Srivastava et al. 2011; Hu and De Giacomo 2013).

In this work we focus on *generalized plans* in the form of *planning programs* (Jiménez and Jonsson 2015). Planning programs can model conditional statements and loops which allows them to represent compact solutions to individual planning tasks, as well as generalized plans. Figure 2 shows a planning program for navigating to the $(0, 0)$ cell in a grid starting from any cell. Variables x and y represent the current position. Instructions $\text{dec}(x)$ and $\text{dec}(y)$ decrement the value of x and y . Conditional goto instructions $\text{goto}(0, !(x=0))$ and $\text{goto}(2, !(y=0))$ jump to line 0 when $x \neq 0$ and to line 2 when $y \neq 0$. Finally end is a marker that indicates program termination.

Since the number of possible planning programs is exponential in the number of programmable lines, our approach is to first generate procedural DCK by solving a set of planning problems representative of a certain subtask, and then compile the DCK as callable procedures in the overall generalized planning problem. We implement procedure calling using a stack modeled in PDDL so that an off-the-shelf planner can compute and exploit the DCK. Compared to previous work our contributions are:

1. To the best of our knowledge, ours is the first approach to compute procedural DCK over a wide range of domains and exploit it with an off-the-shelf classical planner. In previous work procedural DCK is hand-coded by human experts (Baier, Fritz, and McIlraith 2007) and tested in a reduced number of domains.
2. This is the first PDDL implementation of procedural DCK with nested and recursive procedure calls. This allows

```

0. dec(x)
1. goto(0,!(x=0))
2. dec(y)
3. goto(2,!(y=0))
4. end

```

Figure 2: Generalized plan in the form of a *planning program* for navigating to the (0, 0) position in a grid.

an off-the-shelf classical planner to compute generalized plans for non-trivial tasks, such as sorting lists with variable size or DFS traversal of binary trees with variable size. Fritz, Baier, and McIlraith (2008) already formalized callable procedures but never implemented them in PDDL. Jiménez and Jonsson (2015) implemented callable procedures in PDDL but only for 1-level procedure calls.

3. With respect to the original compilation (Jiménez and Jonsson 2015), we reduce the number of actions needed to execute *conditional goto* instructions.

Background

Here we introduce classical planning with conditional effects and define generalized planning as we consider it in the paper. We also review the planning programs that we use to represent generalized plans and the compilation to compute them with a classical planner.

Classical Planning

As is common in generalized planning, our work is based on the formulation of classical planning that includes conditional effects. This way a generalized plan can repeatedly refer to the same action, but the actual effect of the action depends on the state in which it is applied.

We describe states and conditional effects in terms of literals. Formally, given a set of fluents F , a literal l is a valuation of a fluent in F , i.e. $l = f$ or $l = \neg f$ for some $f \in F$. A set of literals L thus represents a partial assignment of values to fluents (WLOG we assume that L does not assign conflicting values to any fluent). Given L , let $\neg L = \{\neg l : l \in L\}$ be the complement of L . A *state* s is a set of literals such that $|s| = |F|$, i.e. a total assignment of values to fluents.

A classical planning problem is a tuple $P = \langle F, A, I, G \rangle$, where F is a set of fluents, A a set of actions, I an initial state and G a goal condition, i.e. a set of literals. Each action $a \in A$ has a set of literals $\text{pre}(a)$ called the *precondition* and a set of conditional effects $\text{cond}(a)$. Each conditional effect $C \triangleright E \in \text{cond}(a)$ is composed of sets of literals C (the condition) and E (the effect). Even though I is a total assignment of values to fluents, we often describe the initial state compactly in terms of the fluents that are true in I .

Action a is applicable in state s if and only if $\text{pre}(a) \subseteq s$, and the resulting set of *triggered effects* is

$$\text{eff}(s, a) = \bigcup_{C \triangleright E \in \text{cond}(a), C \subseteq s} E,$$

i.e. effects whose conditions hold in s . The result of applying a in s is a new state $\theta(s, a) = (s \setminus \neg \text{eff}(s, a)) \cup \text{eff}(s, a)$.

A plan for P is an action sequence $\pi = \langle a_1, \dots, a_n \rangle$ that induces a state sequence $\langle s_0, s_1, \dots, s_n \rangle$ such that $s_0 = I$ and, for each i such that $1 \leq i \leq n$, a_i is applicable in s_{i-1} and generates the successor state $s_i = \theta(s_{i-1}, a_i)$. The plan π *solves* P if and only if $G \subseteq s_n$, i.e. if the goal condition is satisfied following the application of π in I .

Generalized Planning

Our definition of generalized planning is based on that of Hu and De Giacomo (2011), who define a generalized planning problem $\mathcal{P} = \{P_1, \dots, P_T\}$ as a set of multiple individual planning problems that share the same observations and actions. Although actions are shared, an action may produce different results due to conditional effects. A solution to a generalized planning problem \mathcal{P} is a generalized plan that solves each individual problem P_t , $1 \leq t \leq T$.

We restrict the above definition in two ways: 1) states are fully observable, so observations are equivalent to states; and 2) each action has the same (conditional) effects in each individual problem. As a consequence, individual problems $P_1 = \langle F, A, I_1, G_1 \rangle, \dots, P_T = \langle F, A, I_T, G_T \rangle$ are classical planning problems that share fluents and actions, and thus only differ in the initial state and the goal.

Early work for generalized planning followed an inductive approach, solving individual planning problems in \mathcal{P} separately using a classical planner. An individual solution is then merged with the current generalized plan, and the process is repeated until the generalized plan solves the full set of problems (Winner and Veloso 2003; Srivastava et al. 2011). Another approach is compiling the generalized planning task into a conformant planning problem (Bonet, Palacios, and Geffner 2010), i.e. synthesizing a strong plan for a planning task with multiple possible initial states that encode the different individual planning problems under the assumption of null run-time observability.

Generalized plans can have diverse forms that range from *DS-planners* (Winner and Veloso 2003) or *generalized policies* (Martín and Geffner 2004), to finite state machines (FSMs) (Bonet, Palacios, and Geffner 2010). In this paper we represent and compute generalized plans using the recent formalism of *planning programs* (Jiménez and Jonsson 2015) that we next proceed to define.

Planning Programs

Given a classical planning problem $P = \langle F, A, I, G \rangle$, a planning program Π is a numbered list of *instructions* such that the instruction w on each line i of the program is either:

1. A *sequential instruction*, i.e. $w \in A$.
2. A *conditional goto instruction*, i.e. $w = \text{goto}(i', !f)$, where i' is the target program line and $f \in F$ the condition. Let \mathcal{I}_{go} be the set of conditional goto instructions.
3. A *termination instruction* marking the end of the program.

To execute a planning program Π we maintain a current state s , initialized to I , and a program counter pc , initialized to 0. Let w be the instruction on the line indicated by pc . If $w \in A$, we update s as a result of applying w and increment pc . If $w = \text{goto}(i', !f)$, we set pc to i' if f is false in the

current state s , and increment pc otherwise (as in the original paper, we jump whenever f is false). Eventually, if w is a termination instruction, execution ends successfully.

Since conditional goto instructions may cause infinite loops, execution fails whenever we reach a pair of state and program counter (s, pc) already visited. A planning program Π solves a classical planning problem P if the execution of Π ends successfully and the goal condition holds in the resulting state, i.e. $G \subseteq s$. A planning program Π solves a generalized planning problem $\mathcal{P} = \{P_1, \dots, P_T\}$ if it solves every classical planning problem in \mathcal{P} .

Computing Programs with Classical Planning

Jiménez and Jonsson (2015) introduced a compilation that simultaneously computes a planning program and verifies that it solves a set of planning problems. Given a classical planning task $P = \langle F, A, I, G \rangle$ the result of the compilation is a new classical planning task $P_n = \langle F_n, A_n, I_n, G_n \rangle$ where n bounds the number of lines of the program.

To specify P_n we have to introduce prior notation. Let $F_{pc} = \{pc_i : 0 \leq i \leq n\}$ be the fluents for coding the program counter and let $F_{ins} = \{ins_{i,w} : 0 \leq i \leq n, w \in A \cup \mathcal{I}_{go} \cup \{\text{nil}, \text{end}\}\}$ be the fluents coding the instruction w on line i . Here, nil denotes a line that has not yet been programmed, while end denotes the end of the program. Finally, let done be a fluent modeling we are done programming.

For each $a \in A$, let $a_i, 0 \leq i < n$, be a classical planning action with precondition $\text{pre}(a_i) = \text{pre}(a) \cup \{pc_i\}$ and conditional effects $\text{cond}(a_i) = \text{cond}(a) \cup \{\emptyset \triangleright \{\neg pc_i, pc_{i+1}\}\}$. Likewise, for each goto instruction $\text{goto}(i', !f) \in \mathcal{I}_{go}$, let $\text{go}_i^{i',f}, 0 \leq i < n$, be a classical action defined as

$$\begin{aligned} \text{pre}(\text{go}_i^{i',f}) &= \{pc_i\}, \\ \text{cond}(\text{go}_i^{i',f}) &= \{\emptyset \triangleright \{\neg pc_i\}, \\ &\quad \{\neg f\} \triangleright \{pc_{i'}\}, \{f\} \triangleright \{pc_{i+1}\}\}. \end{aligned}$$

Let $\text{end}_i, 0 < i \leq n$, be a classical action defined as $\text{pre}(\text{end}_i) = \{pc_i\}$ and $\text{cond}(\text{end}_i) = \{\emptyset \triangleright \{\text{done}\}\}$, corresponding to the termination instruction.

Let $w \in A \cup \mathcal{I}_{go} \cup \{\text{end}\}$ be an instruction and let w_i be the corresponding classical planning action that executes instruction w on line $0 \leq i \leq n$. Since w may be executed multiple times, we define two versions: $P(w_i)$, that is only applicable on an empty line i and programs w on that line, and $R(w_i)$, that is only applicable when instruction w already appears on line i and repeats the execution of w :

$$\begin{aligned} \text{pre}(P(w_i)) &= \text{pre}(w_i) \cup \{ins_{i,\text{nil}}\}, \\ \text{cond}(P(w_i)) &= \{\emptyset \triangleright \{\neg ins_{i,\text{nil}}, ins_{i,w}\}\}, \\ \text{pre}(R(w_i)) &= \text{pre}(w_i) \cup \{ins_{i,w}\}, \\ \text{cond}(R(w_i)) &= \text{cond}(w_i). \end{aligned}$$

Now we are ready to define $P_n = \langle F_n, A_n, I_n, G_n \rangle$:

- $F_n = F \cup F_{pc} \cup F_{ins} \cup \{\text{done}\}$,
- $A_n = \{P(a_i), R(a_i) : a \in A, 0 \leq i < n\}$
 $\cup \{P(\text{go}_i^{i',f}), R(\text{go}_i^{i',f}) : \text{goto}(i', !f) \in \mathcal{I}_{go}, 0 \leq i < n\}$
 $\cup \{P(\text{end}_i), R(\text{end}_i) : 0 < i \leq n\}$,

```

main:  0. p1      | p2:  0. inc(x)
       1. p2      |       1. goto(0,!(x=n))
       2. p3      |       2. end
       3. p4      | p3:  0. inc(y)
       4. end     |       1. goto(0,!(y=n))
                    |       2. end
                    | p4:  0. dec(x)
                    |       1. goto(0,!(x=0))
                    |       2. end

```

Figure 3: Planning program for visiting the 4 corners of a $n \times n$ grid starting from any initial position and using 4 auxiliary procedures (p1 is defined by the program in Figure 2).

- $I_n = I \cup \{ins_{i,\text{nil}} : 0 \leq i \leq n\} \cup \{pc_0\}$,
- $G_n = G \cup \{\text{done}\}$.

The compilation can be extended to a generalized planning problem $\mathcal{P} = \{P_1, \dots, P_T\}$. After computing a planning program Π and verifying that it solves P_1 , simulating the end instruction resets the program counter to 0 and the state to I_2 , the initial state of P_2 . To solve P_n , the classical plan has to simulate the execution of all planning problems in \mathcal{P} , thus verifying that the planning program Π solves them all.

Planning Programs with 1-Level Procedure Calls

Jiménez and Jonsson (2015) extended planning programs with 1-level callable procedures. Figure 3 shows a planning program with 1-level callable procedures for visiting the corners of a square grid starting from any initial position.

Planning programs with 1-level procedure calls have a new set of instructions \mathcal{I}_{call} , for calling auxiliary procedures (in Figure 3, $\mathcal{I}_{call} = \{p1, p2, p3, p4\}$). Program execution proceeds as explained before except in the particular case where the instruction to execute is a calling procedure instruction. In that case the program counter is set to the first line of the called procedure and execution continues from there. Program execution always starts on the first line of the main procedure, auxiliary procedures can only be called from the main procedure, and control always returns to the main procedure when the auxiliary procedure ends.

Jiménez and Jonsson (2015) also defined a compilation to compute planning programs with 1-level procedure calls. The result of the compilation is a classical planning task $P_{n,b} = \langle F_{n,b}, A_{n,b}, I_{n,b}, G_{n,b} \rangle$ where b is a new parameter to bound the number of procedures of the planning program. $F_{n,b}$ contains the fluents of F_n with the following modifications: Fluents $pc_{i,j}$ and $ins_{i,j,w}$ are parameterized with the associated procedure $0 \leq j \leq b$ and a new fluent main is included to reflect that control is currently with the main procedure. Actions in $A_{n,b}$ are modified as follows:

- Actions $a_{i,j}$ and $\text{go}_{i,j}^{i',f}$ include the associated procedure $0 \leq j \leq b$, and have extra precondition main for $j = 0$.
- New actions $\text{call}_{i,0}^j$ implement the call to procedure 1 $\leq j \leq b$ on line $0 \leq i < n$ of the main program, and are defined as $\text{pre}(\text{call}_{i,0}^j) = \{pc_{i,0}, \text{main}\}$ and $\text{cond}(\text{call}_{i,0}^j) =$

$\{\emptyset \triangleright \{\neg pc_{i,0}, \neg main, pc_{i+1,0}, pc_{0,j}\}\}$.

- New actions $end_{i,j}$ are defined differently for the main procedure, $j = 0$, and for the auxiliary procedures $j > 0$:

$$\begin{aligned} pre(end_{i,0}) &= \{pc_{i,0}, main\}, \\ cond(end_{i,0}) &= \{\emptyset \triangleright \{done\}\}, \\ pre(end_{i,j}) &= \{pc_{i,j}\}, j > 0, \\ cond(end_{i,j}) &= \{\emptyset \triangleright \{\neg pc_{i,j}, main\}\}, j > 0. \end{aligned}$$

Nested Procedure Calls

In this section we extend the 1-level procedure calls of Jiménez and Jonsson (2015) to nested and recursive procedure calls. In particular, we represent a *call stack* and extend the semantics of *call* and *termination* instructions:

- *Call* instructions now (1) increment the current program counter; and (2) push information onto the stack about the new procedure and its program counter.
- *Termination* instructions now pop information about the current procedure and program counter from the stack.

The stack model

We model a stack of finite size inspired by the compilation of *fault tolerant planning* into classical planning (Domshlak 2013). We define the stack as a tuple $\langle K, m \rangle$ where $K \subseteq F$ is a given set of *stackable fluents*, i.e., fluents that can be allocated in the stack, and m is the maximum size of the stack. Implicitly our stack model defines:

- A set of fluents $K^m = \{f^k : f \in K, 1 \leq k \leq m\}$ that contains replicas of the fluents in K parameterized with the stack level k . These fluents represent the m partial states that can be stored in the stack.
- A set of fluents $F_{top}^m = \{top^k\}_{1 \leq k \leq m}$ representing the top level of the stack at the current time.
- Actions $push_Q$ and pop are the canonical stack operations, with $push_Q$ pushing a subset of stackable fluents $Q \subseteq K$ to the top level of the stack and pop popping any fluent in K from the top level of the stack.

Planning Programs with Nested Procedure Calls

Here we formalize our extension to planning programs with nested and recursive procedure calls. We add a third compilation parameter m that models the stack size and hence bounds the depth of allowed nested procedure calls. The result of our extended compilation is then a classical planning task $P_{n,b}^m = \langle F_{n,b}^m, A_{n,b}^m, I_{n,b}^m, G_{n,b}^m \rangle$ defined as:

- $F_{n,b}^m = (F \setminus K) \cup K^m \cup F_{pc}^m \cup F_{ins} \cup F_{top}^m \cup \{done\}$, i.e. stackable fluents in $K \subseteq F$ are replicated and F_{pc}^m contains stackable fluents indicating the line and procedure currently being executed: $F_{pc}^m = \{pc_i^k : 0 \leq i \leq n, 1 \leq k \leq m\} \cup \{proc_j^k : 0 \leq j \leq b, 1 \leq k \leq m\}$.
- $A_{n,b}^m$ parameterizes the actions in $A_{n,b}$ with the stack level k and adds a precondition top^k . Also, the actions for calling and ending procedures are redefined. For each line i , $0 \leq i < n$, pair of procedures $j, j', 0 \leq j, j' \leq b$, and

stack level k , $1 \leq k < m$, action $call_{i,j}^{j',k}$ simulates a call to j' from line i of procedure j on stack level k , and action $end_{i,j}^{k+1}$ simulates the termination on line i of procedure j on stack level $k+1$:

$$pre(call_{i,j}^{j',k}) = \{top^k, pc_i^k, proc_j^k\}, \quad (1)$$

$$\begin{aligned} cond(call_{i,j}^{j',k}) &= \{\emptyset \triangleright \{\neg pc_i^k, pc_{i+1}^k, \neg top^k, top^{k+1}\}\} \\ &\cup \{\emptyset \triangleright \{pc_0^{k+1}, proc_{j'}^{k+1}\}\}, \quad (2) \end{aligned}$$

$$pre(end_{i,j}^{k+1}) = \{top^{k+1}, pc_i^{k+1}, proc_j^{k+1}\},$$

$$\begin{aligned} cond(end_{i,j}^{k+1}) &= \{\emptyset \triangleright \{\neg top^{k+1}, \neg pc_i^{k+1}, \neg proc_j^{k+1}\}\} \\ &\cup \{\emptyset \triangleright \{top^k\}\} \cup \{\emptyset \triangleright \{\neg f^{k+1} : f \in K\}\}. \end{aligned}$$

Action $end_{i,0}^1$, $0 < i \leq n$, is defined as $pre(end_{i,0}^1) = \{top^1, pc_i^1, proc_0^1\}$ and $cond(end_{i,0}^1) = \{\emptyset \triangleright \{done\}\}$.

- The initial state copies stackable fluents onto stack level 1, initializes empty program lines and sets the procedure on stack level 1 to 0: $I_{n,b}^m = (I \setminus K) \cup \{f^1 : f \in I \cap K\} \cup \{ins_{i,j,nil} : 0 \leq i \leq n, 0 \leq j \leq b\} \cup \{top^1, pc_0^1, proc_0^1\}$.
- The goal is the same as before, i.e. $G_{n,b}^m = G \cup \{done\}$.

Action $call_{i,j}^{j',k}$ is defined for any procedure pair j, j' ; specifically the case where $j = j'$ corresponds to a recursive call. The definition of $call_{i,j}^{j',k}$ included here does not copy stackable fluents and thus assumes $K = \emptyset$; in a later section we redefine call actions such that they copy stackable fluents.

Improving the Computation of Planning Programs

Apart from implementing nested and recursive procedure calls, we introduce another improvement to the compilation that reduces the number of actions needed to execute *conditional goto* instructions. In the original compilation, the number of $goto_{i,i'}^{i',f}$ actions is $|F| \cdot n^2$, since any fluent $f \in F$ can be a condition and there are n^2 combinations of line pairs (i, i') . Here we reduce this number to $(|F| + n) \cdot n$.

The idea is to split $goto_{i,i'}^{i',f}$ actions into two actions: $eval_i^f$, that evaluates condition f on line i , and $jmp_i^{i'}$, that performs the conditional jump according to the evaluation outcome. This is inspired by assembly languages that separate comparison instructions that modify flags registers, e.g., *CMP* and *TEST* in the *x86 assembly* language, from jump instructions that update the program counter according to the flag registers, e.g., *JZ* and *JNZ* in *x86 assembly*.

To implement the split we introduce two new fluents acc and $eval$, initially false. Fluent acc records the outcome of the evaluation, while $eval$ indicates that the evaluation has been performed. Actions $eval_i^f$ and $jmp_i^{i'}$ are defined as

$$pre(eval_i^f) = \{pc_i, \neg eval\},$$

$$cond(eval_i^f) = \{\{f\} \triangleright \{acc\}\} \cup \{\emptyset \triangleright \{eval\}\},$$

$$pre(jmp_i^{i'}) = \{pc_i, eval\},$$

$$cond(jmp_i^{i'}) = \{\emptyset \triangleright \{\neg pc_i, \neg eval\}\}$$

$$\cup \{\{\neg acc\} \triangleright \{pc_{i'}\}\}$$

$$\cup \{\{acc\} \triangleright \{pc_{i+1}, \neg acc\}\}.$$

<pre>main: 0. p1(x) 1. p1(y) 2. p2(x) 3. p2(y) 4. p1(x) 5. end</pre>	<pre>p1(aux): 0. dec(aux) 1. goto(0,!(aux=0)) 2. end p2(aux): 0. inc(aux) 1. goto(0,!(aux=n)) 2. end</pre>
--	---

Figure 4: Planning program with parameterized procedures for visiting the four corners of a $n \times n$ grid. Procedure $p1(aux)$ decrements variable aux until reaching value 0 while $p2(aux)$ increments aux until reaching value n .

Actions for *programming* a conditional goto remain the same, but we introduce new actions $R(\text{eval}_i^f)$ and $R(\text{jmp}_i^{f'})$ for repeating the execution of a goto instruction.

Parameterized Procedures

In many programs the number of required auxiliary procedures can be reduced parameterizing the procedures, which also decreases the total number of program lines. For example, Figure 4 shows a planning program with two parameterized procedures for visiting the corners of a grid. Compared to the program in Figure 3, the new program is significantly more compact.

In this section we extend nested procedure calls to parameterized procedures. We do so extending the semantics of call instructions. Now, apart from updating the program counter and pushing the new procedure onto the stack, *Call instructions* also push the parameters of the procedure onto the stack.

To explain parameter passing, we expand on the notation using variables to represent values. For a given finite domain of values D , let $V(D)$ be a set of variables for storing a value in D . Let $F_{V(D)} = \{\text{assign}_{v,x} : x \in D, v \in V(D)\}$ be a set of fluents encoding assignments of type $v = x$. The key characteristic of this notation is that variables are represented as *objects*.

We now associate each procedure j' with a parameter list $(D_1, u_1), \dots, (D_r, u_r)$ where $D_q, 1 \leq q \leq r$, is a finite domain and $u_q \in V(D_q)$ is a corresponding (fixed) variable. The set of stackable fluents of j' is $F_{V(D_1)} \cup \dots \cup F_{V(D_r)}$. We use the action $\text{call}_{i,j}^{j',k}$ defined in Equations (1) and (2) to define the actions for calling procedure j' . For each variable combination $v_1, \dots, v_r \in V(D_1) \times \dots \times V(D_r)$, we introduce a new action $\text{call}_{i,j}^{j',k}(v_1, \dots, v_r)$ defined as

$$\begin{aligned} \text{pre}(\text{call}_{i,j}^{j',k}(v_1, \dots, v_r)) &= \text{pre}(\text{call}_{i,j}^{j',k}), \\ \text{cond}(\text{call}_{i,j}^{j',k}(v_1, \dots, v_r)) &= \text{cond}(\text{call}_{i,j}^{j',k}) \\ \cup \{ \{ \text{assign}_{v_q,x}^k \} \triangleright \{ \text{assign}_{u_q,x}^{k+1} \} : 1 \leq q \leq r, x \in D_q \}. \end{aligned}$$

In other words, $\text{call}_{i,j}^{j',k}(v_1, \dots, v_r)$ has the effect of *copying* the value of each variable $v_q, 1 \leq q \leq r$, on level k of the stack to the (fixed) variable u_q on level $k + 1$ of the stack. Since each stack level has a separate fluent set, we can reuse variables for different procedures and procedure calls (cf. Figure 4 where $p1$ and $p2$ share variable aux).

Planning Programs As Control Knowledge

The compilation by Jiménez and Jonsson (2015) for generating planning programs with a classical planner does not show clear benefit from including auxiliary procedures. Indeed the configurations of the compilation that drop auxiliary procedures and increase the number of program lines of the main program were reported to achieve better, or at least equivalent, experimental results. Considering this and that the number of programs is exponential in the number of program lines, its applicability is limited to tasks solvable with planning programs of small size (they do not report planning programs of more than 6 lines including main and auxiliary procedures).

Here we extend the applicability of planning programs to address more challenging generalized planning tasks. This is done by incorporating existing planning programs as given DCK in the form of auxiliary procedures. We now support nested procedure calls so this given DCK can represent hierarchies of planning programs.

Exploiting Procedural Domain Control Knowledge

In the above compilations all program lines are initially empty. However, nothing prevents us from programming some of the instructions as part of the initial state. Specifically, we consider partial planning programs such that the program lines of the main procedure are empty, but the lines of the auxiliary procedures are already programmed. For example, when generating the planning program shown in Figure 3, the five program lines of the main procedure are initially empty but the lines of the four auxiliary procedures $p1, p2, p3$ and $p4$ appear as programmed in the initial state of the classical planning problem.

The benefits of this approach compared to the original compilation are:

- The number of empty program lines is reduced, decreasing the number of possible planning programs and hence the complexity of the planning problem resulting from the compilation.
- Each procedure solves a clearly defined subtask so the benefit of decomposing a program into procedures becomes more apparent (depending, of course, on how the procedure is obtained).

Evidently, the benefit of including a given procedure is contingent on how much the procedure contributes to solving the overall problem. The question, then, is how to generate useful procedures. One option is for a domain expert to hand-craft auxiliary procedures (Baier, Fritz, and McIlraith 2007), and this might be the best choice if such control knowledge is readily available.

Automatic Generation of Procedural DCK

Here we show how to take another approach to generating useful procedural DCK, which is to automatically compute the auxiliary procedures. Since each procedure is in effect a program of its own, we can use the same compilation as in Jiménez and Jonsson (2015) to compute this program from

examples (albeit without auxiliary procedures). In the navigation example from Figure 1, to compute the auxiliary procedure $p1$, i.e., a program for navigating to position $(0, 0)$, we simply define a series of planning problems with different initial states whose goal condition is to be at position $(0, 0)$. Similarly, we can define suitable planning problems to compute the remaining auxiliary procedures in Figure 3.

We assume the existence of a specific decomposition of the overall problem into a set of subtasks, and appropriately extend the definition of a generalized planning problem. Formally, a generalized planning problem with subtasks is a tuple $\mathcal{P}_{sub} = \langle \mathcal{P}, \mathcal{P}^1, \dots, \mathcal{P}^b \rangle$, where $\mathcal{P} = \{P_1, \dots, P_T\}$ is the generalized planning problem that we ultimately want to solve. For each j , $1 \leq j \leq b$, $\mathcal{P}^j = \{P_1^j, \dots, P_T^j\}$ is also a generalized planning problem whose purpose is to compute the j -th auxiliary procedure that corresponds to the j -th subtask.

Given a generalized planning problem with subtasks \mathcal{P}_{sub} , for each j , $1 \leq j \leq b$, we separately compute a planning program Π^j that solves the generalized planning problem \mathcal{P}^j . We then compile a partial planning program with b auxiliary procedures corresponding to the planning programs Π^1, \dots, Π^b , but whose main procedure has empty program lines. Finally, we compute a solution to the compiled planning problem, corresponding to a generalized plan in the form of a planning program with procedures that solves the generalized planning problem \mathcal{P} . This process can be iterated generating a hierarchy of planning programs.

Evaluation

In all experiments, we run the classical planner Fast Downward (Helmert 2006) with the LAMA-2011 setting (Richter and Westphal 2010) on a processor Intel Core i5 3.10GHz x 4 with a 4GB memory bound and time limit of 3600s.

We evaluate our approach in Blocks, Gripper, Hall-A and Visual-Marker from Bonet, Palacios, and Geffner (2010). In **Blocks** the problem is to unstack blocks from a tower until a green block is found. In **Gripper** a robot has to move balls from one room to another. In **Hall-A** the problem is to visit the four corners of a grid as in Figure 1. In **Visual-Marker** the problem is moving a marker from the bottom-left corner of a grid to a cell with a green block. We also add the **Grid** domain, where an agent in an arbitrary cell of a grid must visit another arbitrary cell and **Visit-All**, where an agent in the bottom-left corner of a grid must visit all the grid cells.

In these domains we model integer variables v with fluents of type $v = a$, and define actions $inc(v)$ and $dec(v)$ with conditional effects to increment or decrement v . For instance, in Grid and Visit-All variables x and y model the position of the agent, and w and h model the width and height of the grid (allowing then variable grid sizes). In Grid we also model the goal cell using variables x_G and y_G . To obtain a general program, Grid models derived fluents $x = x_G$ and $y = y_G$ that reflect when the agent is at a goal position, and Visit-All includes $x = w$ and $y = h$ that reflect when the agent is at the last row or column. Both domains include a `visit` action that marks the current cell as visited.

A benefit of computing generalize plans in the form of

programs is that we can naturally model programming tasks:

- **Summatory** models computing the summatory of a natural number n , $y = 1 + \dots + n$, using integer variables n and y . The initial state of each test is $(n = m, y = 0)$ for some m and the goal is $y = \sum_{x=1}^m x$. **Fibonacci** models computing the n^{th} term of the *Fibonacci* sequence, $F_n = F_{n-1} + F_{n-2}$, using integer variables n , F_n , F_{n-1} and F_{n-2} . The initial state is $(n = m, F_n = 1, F_{n-1} = 0, F_{n-2} = 0)$ for some m and the goal is $F_n = F_m$.
- **Reverse** models the task of reversing a list of numbers with variable list size n using integer variables i and j (representing iterators) as well as n and v_r , $1 \leq r \leq n$. The domain includes action `swap-i-j` that swaps the values of v_i and v_j . The initial state is given by $(i = 1, j = n)$ and arbitrary n and v_r , $1 \leq r \leq n$, and the goal is to reverse the list represented by v_r , $1 \leq r \leq n$. **Sorting** models the task of ordering a list of numbers with variable list size n using integer variables i , j and k (representing iterators) as well as n and v_r , $1 \leq r \leq n$. The domain includes action `swap-i-k`, actions `inc(i)` and `inc(j)`, and an action that assigns the resulting value of j to k . The initial state is given by $i = j = k = 1$ and arbitrary n and v_r , $1 \leq r \leq n$, and the goal is to sort the list represented by v_r , $1 \leq r \leq n$.
- The **List** domain iterates over the nodes of a linked list. We define a finite domain D that includes all list nodes, and variables that point to nodes. Fluent `is-tail` indicates the tail node, action `next` accesses the next node and action `visit` marks nodes as visited. The **Tree** domain DFS visits all nodes of a binary tree. The domain D includes tree nodes, and fluents `is-internal` indicate internal nodes. Two actions `left` and `right` access the left and right child of a node, respectively, and we use variables `current` and `child`. The `visit` action marks nodes as visited.

For programming tasks, given two integer variables v and w we add action `assign(v,w)` to assign the value of w to v , and `add(v,w)`, adding the value of w to the current value of v . We make minimal assumptions about conditions in goto instructions. Given a generalized planning problem $\mathcal{P} = \{P_1 = \langle F, A, I_1, G_1 \rangle, \dots, P_T = \langle F, A, I_T, G_T \rangle\}$, any fluent in F can be a condition. Moreover, for selected pairs of integer variables v and w , we add derived fluents $(v = w)$, as well as $(v < w)$ for Reverse and Sorting.

Figure 5 shows the flow diagrams of the planning programs obtained for several evaluation domains and Table 1 summarizes the obtained experimental results. For each domain we report: number of **procedures** used to solve the generalized planning problem (including the main procedure) and kind of **solution** (or error) obtained. For each procedure (starting with the main procedure): **program lines**, number of **instances** used, **planning time** (the sum of all the individual times too) and **plan length**. When the solution kind is **One Procedure** (OP) only a main procedure is needed to obtain a generalized plan. This corresponds to not using DCK and is the case of Blocks, Gripper, Summatory, Fibonacci and Reverse, where a general program able to solve any instance of the domain was generated without

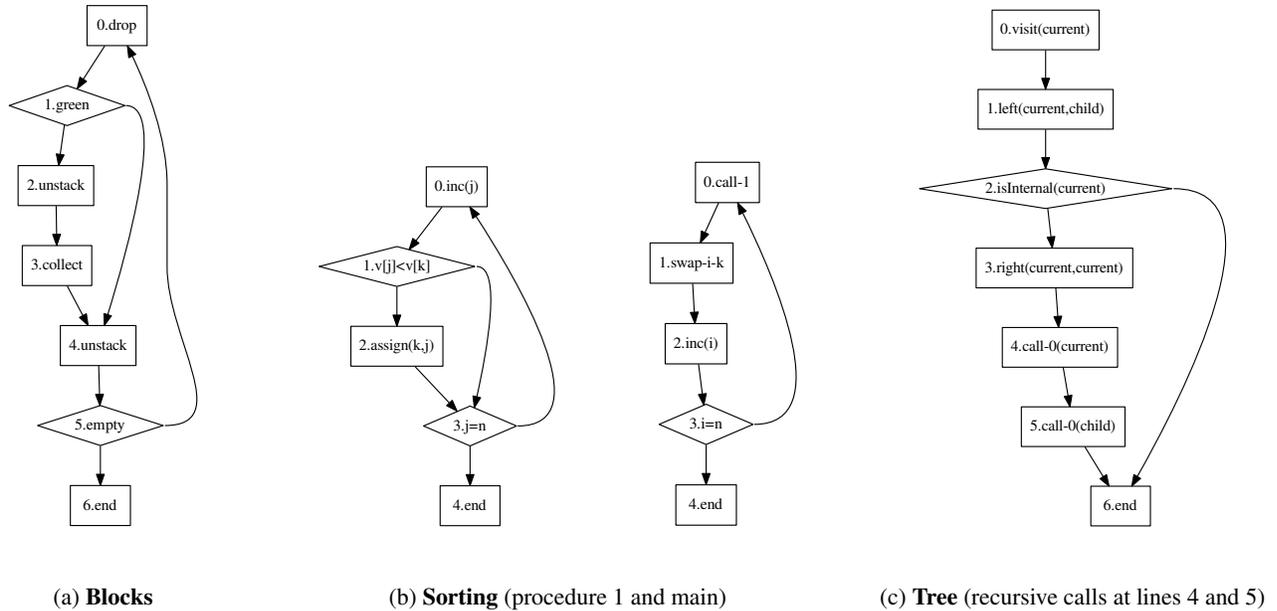


Figure 5: Flow diagrams of example planning programs obtained for selected domains.

DCK. In List and Tree only 1 procedure was used but with **Recursivity (R)** and **Recursivity with Parameters (RP)** respectively. **Nested Procedures (NP)** indicates that the number of procedures is greater than 1 so auxiliary procedures are used as DCK. In some cases, where procedural DCK pays off, the generalized planning problem could not be solved without DCK in the given memory bound, e.g., Hall-A, Grid, Sorting, Visit-All and Visual-Marker. In these domains procedural DCK was automatically generated as generalized plans for subtasks and then inserted as procedures in the overall generalized planning problem. Other domains such as Blocks, Fibonacci and Gripper can be solved without DCK, even though using DCK is better, in terms of time for the first two, and plan length for the latter.

We briefly describe the obtained auxiliary procedures. In **Hall-A**, each procedure solves the task of visiting one of the 4 corners. In **Grid**, one procedure solves the task of reaching the target column (x), while the other reaches the target row (y). In **Visit-All**, one procedure visits all cells in a grid with a single row from left to right, while the other returns the agent to the starting (i.e. leftmost) column. In **Visual Marker**, similar to **Visit-All**, one procedure moves the marker right until a green block is found (or the end of the grid is reached) and the other returns the marker back to the first column. In **Fibonacci** the procedure performs one update of the Fibonacci sequence. In **Sorting**, the procedure solves the generalized planning problem of selecting the smallest element in a partial list; hence the generalized plan corresponds to selection sort. The **Blocks** domain has a procedure to find the green block, and **Gripper** has two procedures, the first one collects two balls and moves to the other room and the second one drops the carried balls and

moves to the other room.

The compilation of Jiménez and Jonsson (2015) cannot generate recursive solutions, i.e. of types R and RP. Besides it fails to generate solutions that require a significant number of program lines (Grid, Hall-A, Sorting, Visit-all and Visual-Marker). The state-of-the-art in generalized planning is the automatic derivation of FSMs. This approach is able to solve the Hall-A, Blocks, Gripper and Visual-Marker (Bonet, Palacios, and Geffner 2010; Hu and De Giacomo 2013) domains as well as Visit-All and Sorting (Srivastava et al. 2011). The previous solution for Sorting assumed fixed list size and an existing action for inserting an element in a sorted list, while we allow variable list size and automatically compute the procedure for selecting the smallest element from a sublist. Remarkably, Hu and De Giacomo (2013) compute solutions to the FSM domains in a fraction of a second. FSMs can represent generalized plans more compactly than planning programs, reducing the search space, and the authors also hand-pick the conditions for branching while we allow any fluent to appear as a condition. However, there is currently no mechanism that enables procedure calls for FSMs, preventing FSMs from representing solutions to the **Tree** domain.

Related Work

Procedural DCK in the form of Golog-like programs (Baier, Fritz, and McIlraith 2007; Fritz, Baier, and McIlraith 2008) include conditionals and loops as well as non-deterministic action choices that constrain the search for a solution plan. Nevertheless they are not proper generalized plans since it is still necessary to apply a planner to solve each individual planning problem. In addition they are hand-crafted and do

Domain	Procedures	Solution	Lines	Instances	Time(s)	Total time (s)	Plan length
Blocks	2	NP	4,3	5,5	2,2	4	46,61
Blocks	1	OP	6	5	85	85	73
Fibonacci	2	NP	3,3	2,5	2,177	179	12,129
Fibonacci	1	OP	5	2	3570	3570	56
Grid	3	NP	5,5,2	2,2,4	611,631,2	1244	43,43,213
Grid	1	ME	10	5	-	-	-
Gripper	3	NP	3,3,3	2,2,2	1,1,2	4	12,12,54
Gripper	1	OP	4	2	1	1	77
Hall-A	5	NP	5,5,5,3,5	2,2,2,2,2	3,7,3,1,4	18	44,40,40,73,155
Hall-A	1	ME	14	2	-	-	-
List	1	R	5	6	5	5	120
Reverse	1	OP	4	2	22	22	38
Sorting	2	NP	4,4	4,3	14,16	30	73,188
Sorting	1	ME	7	4	-	-	-
Summatory	1	OP	3	2	1	1	26
Tree/DFS	1	RP	6	1	165	165	51
Visitall	3	NP	4,2,4	2,2,2	1,1,5	7	47,18,238
Visitall	1	ME	7	2	-	-	-
Visual-Marker	3	NP	4,2,4	4,2,5	2,1,10	13	82,18,205
Visual-Marker	1	ME	8	2	-	-	-

Table 1: Procedures (1: only main; >1: DCK as auxiliary procedures were used), solution kind (NP=Nested Procedures, OP=One Procedure, R=Recursivity, RP=Recursivity with Parameters, ME=Memory error) and per each procedure: program lines, instances used to generate each procedure, planning time and plan length.

not implement in PDDL mechanisms for procedure calling.

Planning programs (Jiménez and Jonsson 2015) can model generalized plans and be automatically computed from test cases but allow only one level of procedure calls which makes recursion unfeasible. Moreover the original compilation does not show benefits from including auxiliary procedures so it could only address generalized planning tasks that can be solved with small planning programs.

Previous work on computing FSMs (Bonet, Palacios, and Geffner 2010) also uses a compilation that interleaves programming a FSM with verifying that it satisfies a set of test cases. The generation of a FSM is however different since it starts from a partially observable planning model and uses a conformant to classical planning compilation (Palacios and Geffner 2009). FSMs can be understood as a way of representing and computing procedural DCK that (1) does not implement procedure calls and recursion; (2) does not reuse FSMs for similar tasks. Another difference is that our *conditional goto* instructions can branch on any fluent without the need to predefine a subset of observations, although we would also benefit from restricting the number of branch conditions.

Conclusion

In this paper we introduce a novel approach to computing and exploiting procedural DCK to solve generalized planning tasks. Our approach compiles procedural DCK into callable procedures that allow an off-the-shelf planner to solve challenging generalized planning tasks such as sorting lists with variable list size or DFS traversal of a binary tree with variable size. As far as we know this is the first approach to compute procedural DCK over a wide range of planning tasks using an off-the-shelf planner. Ours is

also the first PDDL implementation of procedural DCK that handles nested procedure calls and allows representing and computing recursive solutions with an off-the-shelf planner.

In experiments, the resulting generalized plan solves all instances of a given domain. In general, however, there are no such guarantees; an extension would be to define test instances and verify that a generalized plan solves all of them. We have not formally characterized the type of plans that we can represent but intuitively, single programs (or FSMs) correspond to regular expressions, and our additions make it possible to represent plans more compactly. Some planning domains require extra high-level state features, e.g., "well-placed", "above" or "good-tower" concepts in Blocksworld (Martín and Geffner 2004), for defining compact generalized plans. It remains as future work how to automatically generate such high-level state features.

The strongest assumption of our approach is the existence of a subtask decomposition of the overall generalized planning problems, necessary to compute the DCK in the form of procedures. An interesting research direction would be to automatically discover these decompositions. This would make it possible to automatically generate and solve a generalized planning problem with subtasks. Work on the automatic generation of planning hierarchies (Hogg, Munoz-Avila, and Kuter 2008) may be a good starting point. Another opportunity for future research is the identification of test cases that produce effective generalized plans for a given domain. This is also an open problem for FSMs (Bonet and Geffner 2015) and, in general, for planning and learning systems (Fern, Khardon, and Tadepalli 2011; Jiménez et al. 2012).

Acknowledgments Sergio Jiménez is partially supported by the *Juan de la Cierva* program funded by the Spanish government.

References

- Bacchus, F., and Kabanza, F. 1999. Using temporal logics to express search control knowledge for planning. *Artificial Intelligence* 116:2000.
- Baier, J. A.; Fritz, C.; and McIlraith, S. A. 2007. Exploiting procedural domain control knowledge in state-of-the-art planners. In *ICAPS*, 26–33.
- Bonet, B., and Geffner, H. 2015. Policies that generalize: Solving many planning problems with the same policy. *Proc IJCAI 2015*.
- Bonet, B.; Palacios, H.; and Geffner, H. 2010. Automatic derivation of finite-state machines for behavior control. In *AAAI Conference on Artificial Intelligence*.
- Botea, A.; Enzenberger, M.; Miller, M.; and Schaeffer, J. 2005. Macro-ff: Improving ai planning with automatically learned macro-operators. *Journal of Artificial Intelligence Research* 24:581–621.
- Coles, A., and Smith, A. 2007. Marvin: A heuristic search planner with online macro-action learning. *Journal of Artificial Intelligence Research* 28:119–156.
- De la Rosa, T.; Jiménez, S.; Fuentetaja, R.; and Borrajo, D. 2011. Scaling up heuristic planning with relational decision trees. *Journal of Artificial Intelligence Research* 767–813.
- Domshlak, C. 2013. Fault tolerant planning: Complexity and compilation. In *Proceedings of the Twenty-Third International Conference on Automated Planning and Scheduling, ICAPS*.
- Erol, K.; Hendler, J.; and Nau, D. 1996. Complexity Results for HTN Planning. *Annals of Mathematics and Artificial Intelligence* 18:69–93.
- Fern, A.; Khardon, R.; and Tadepalli, P. 2011. The first learning track of the international planning competition. *Machine Learning* 84(1-2):81–107.
- Fikes, R.; Hart, P.; and Nilsson, N. 1972. Learning and Executing Generalized Robot Plans. *Artificial Intelligence* 3(4):251–288.
- Fritz, C.; Baier, J. A.; and McIlraith, S. A. 2008. Congolog, sin trans: Compiling congolog into basic action theories for planning and beyond. In *Principles of Knowledge Representation and Reasoning*, 600–610.
- Helmert, M. 2006. The Fast Downward Planning System. *Journal of Artificial Intelligence Research* 26:191–246.
- Hogg, C.; Munoz-Avila, H.; and Kuter, U. 2008. Htn-maker: Learning htms with minimal additional knowledge engineering required. In *AAAI*, 950–956.
- Hu, Y., and De Giacomo, G. 2011. Generalized planning: Synthesizing plans that work for multiple environments. In *International Joint Conference on Artificial Intelligence*, 918–923.
- Hu, Y., and De Giacomo, G. 2013. A generic technique for synthesizing bounded finite-state controllers. In *International Conference on Automated Planning and Scheduling*.
- Hu, Y., and Levesque, H. J. 2011. A correctness result for reasoning about one-dimensional planning problems. In *International Joint Conference on Artificial Intelligence*, 2638–2643.
- Jiménez, S., and Jonsson, A. 2015. Computing Plans with Control Flow and Procedures Using a Classical Planner. In *Proceedings of the Eighth Annual Symposium on Combinatorial Search, SOCS-15*, 62–69.
- Jiménez, S.; De La Rosa, T.; Fernández, S.; Fernández, F.; and Borrajo, D. 2012. A review of machine learning for automated planning. *The Knowledge Engineering Review* 27(04):433–467.
- Jonsson, A. 2009. The role of macros in tractable planning. *Journal of Artificial Intelligence Research* 471–511.
- Martín, M., and Geffner, H. 2004. Learning generalized policies from planning examples using concept languages. *Applied Intelligence* 20:9–19.
- Palacios, H., and Geffner, H. 2009. Compiling uncertainty away in conformant planning problems with bounded width. *Journal of Artificial Intelligence Research* 35:623–675.
- Pralet, C.; Verfaillie, G.; Lemaître, M.; and Infantes, G. 2010. Constraint-based controller synthesis in non-deterministic and partially observable domains. In *European conference on artificial intelligence*, 681–686.
- Richter, S., and Westphal, M. 2010. The LAMA Planner: Guiding Cost-Based Anytime Planning with Landmarks. *Journal of Artificial Intelligence Research* 39:127–177.
- Srivastava, S.; Immerman, N.; Zilberstein, S.; and Zhang, T. 2011. Directed search for generalized plans using classical planners. In *International Conference on Automated Planning and Scheduling*, 226–233.
- Veloso, M.; Carbonell, J.; Perez, A.; Borrajo, D.; Fink, E.; and Blythe, J. 1995. Integrating planning and learning: The prodigy architecture. *Journal of Experimental & Theoretical Artificial Intelligence* 7(1):81–120.
- Winner, E., and Veloso, M. 2003. Distill: Learning domain-specific planners by example. In *International Conference on Machine Learning*, 800–807.
- Yoon, S.; Fern, A.; and Givan, R. 2008. Learning control knowledge for forward search planning. *The Journal of Machine Learning Research* 9:683–718.