Automated Planning

Domain-Configurable Planning: Hierarchical Task Networks

Jonas Kvarnström Department of Computer and Information Science Linköping University

jonas.kvarnstrom@liu.se – 2019

Assumptions



- Recall the fundamental assumption that we <u>only</u> specify
 - Structure: Objects and state variables
 - Initial state and goal
 - Physical preconditions and physical effects of actions

<u>We</u> only specify what <u>can</u> be done The <u>planner</u> should decide what <u>should</u> be done

But even the most sophisticated heuristics and domain analysis methods lack our intuitions and background knowledge...

Domain-Configurable Planners



How can we make <u>a planner</u> take advantage of what <u>we</u> know?

- Planners taking advantage of additional knowledge can be called:
 - Knowledge-rich
 - Domain-configurable
 - (Sometimes incorrectly called "domain-dependent")

Comparisons (1)



4 John Maid

Comparisons (2)



Larger problem classes can be handled efficiently

Domain-configurable

Easier to improve expressivity and efficiency
 Often practically useful for a larger set of domains!

"Domain-independent"

Should be useful for a wide range of domains

Domain-specific

Only works in a single domain

HTNs: Ideas



Classical Planning vs. Hierarchical Task Networks:



Provides guidance but still requires planning

Total-Order Simple Task Networks

A simple form of Hierarchical Task Network, as defined in the book

Terminology 1: Primitive Task



A primitive task corresponds directly to an action

stack(A,B) stack(A,B) load(cranel, loc3, cont5, ...)

load(cranel, loc3, cont5, ...)

Primitive tasks: Dark green (in this presentation)

Corresponding actions: Black

- As in classical planning, **what is primitive** depends on:
 - The execution system
 - How detailed you want your plans to be
- Example:
 - For you, fly(here,there) may be a primitive task
 - For the pilot, it may be decomposed into many smaller steps
- Tasks can be ground or non-ground: stack(A,?x)
 - No separate terminology, as in operator/action

Terminology 2: Non-Primitive Task

A <u>non-primitive task</u>:

- Cannot be directly executed
- Must be <u>decomposed</u> into 0 or more <u>subtasks</u>



Should be decomposed to pickup, putdown, stack, unstack tasks / actions!

Terminology 3: Method



• A <u>method</u> specifies one way to decompose a non-primitive task



- The decomposition is a **graph** $\langle N, E \rangle$
 - Nodes in N correspond to subtasks to perform
 - Can be primitive or not!
 - Edges in E correspond to <u>ordering relations</u>

Totally Ordered STNs



In <u>Totally Ordered Simple Task Networks (STN)</u>, each method must specify a <u>sequence</u> of subtasks

• Can still be modeled as a graph $\langle N, E \rangle$

 $buy-ticket (airport(x), airport(y)) \rightarrow travel (x, airport(x)) \rightarrow fly(airport(x), airport(y)) \rightarrow travel (airport(y), y)$

- Alternatively: A sequence $\langle t_1, \dots, t_k \rangle$
 - <buy-ticket(airport(x), airport(y)),
 travel(x, airport(x)),
 fly(airport(x), airport(y)),
 travel(airport(y), y) >

Totally Ordered STNs (2)



We can <u>illustrate</u> the <u>entire decomposition</u> in this way (horizontal arrow → sequence)



Multiple Methods



- A non-primitive task can have <u>many methods</u>
 - So: You still need to <u>search</u>, to determine which method to use



...and to determine parameters (shown later)

Composition

14 IA

- An HTN plan:
 - <u>Hierarchical</u>
 - Consist of <u>tasks</u>
 - Based on graphs ≈ <u>networks</u>



Domains, Problems, Solutions

- An STN planning domain specifies:
 - A set of <u>tasks</u>
 - A set of <u>operators</u> used for primitive tasks
 - A set of <u>methods</u>

General HTNs:

Can have additional constraints to be enforced

- An STN **problem instance** specifies:
 - An STN planning domain
 - An initial state
 - An **initial task network**, which should be ground (no variables)
 - Total Order STN example:
 <travel(home,work); do-work(); travel(work,home)>



Domains, Problems, Solutions (2)

- Suppose you:
 - Start with the initial task network
 - Recursively apply <u>methods</u> to non-primitive tasks, expanding them
 - Continue until <u>all non-primitive tasks are expanded</u>



- Totally ordered → yields an action <u>sequence</u>
 - If this is executable: A <u>solution</u>
 - (No goals to check implicit in the method structure!)

Domains, Problems, Solutions (3)



- HTN planning uses <u>only</u> the methods specified for a given task
 - Will <u>not</u> try arbitrary actions...
 - For this to be useful, you must <u>have</u> useful "recipes" for all tasks

DWR Example: Moving the Topmost Container

A simple "template expansion"

DWR

Let's switch to Dock Worker Robots...

- Example Tasks:
 - Primitive all DWR actions
 - Move the <u>topmost</u> container between piles
 - Move an <u>entire stack</u> from one pile to another
 - Move a stack, but keep it in the <u>same order</u>
 - Move <u>several stacks</u> in the same order





Methods



- To move the topmost container from one pile to another:
 - <u>task</u>: move-topmost-container(pile1, pile2)
 - <u>method</u>: take-and-put(cont, crane, loc, pile1, pile2, c1, c2)

precond: attached(pile1, loc), attached(pile2, loc), belong(crane, loc), top(cont, pile1), on(cont, c1), top(c2, pile2)

The *task* has parameters **given from above**

A method can have additional parameters, whose values are **chosen by the planner** – just as in classical planning!

The precond adds constraints: crane must be some crane in the same loc as the piles, cont must be the topmost container of pile1, ...

Interpretation:

If you are asked to **move-topmost-container**(pile1, pile2), check all possible values for **cont, crane, loc, c1, c2** where the preconds are satisfied

Methods (2)



• To **move the topmost container** from one pile to another:

- <u>task</u>: move-topmost-container(pile1, pile2)
- <u>method</u>:
 take-and-put(cont, crane, loc, pile1, pile2, c1, c2)
- **precond**: attached(pile1, loc), attached(pile2, loc), belong(crane, loc), top(cont, pile1), on(cont, c1), top(c2, pile2)
- <u>subtasks</u>: <take(crane, loc, cont, c1, pile1), put(crane, loc, cont, c2, pile2)>





DWR Example: Moving a Stack of Containers

Iteration with no predetermined bound

Moving a Stack of Containers

- How can we implement the task move-stack(pile1, pile2)?
 - Should move <u>all</u> containers in a stack
 - There is no <u>limit</u> on how many there might be...



Recursion (1)



- We need a <u>loop</u> with a <u>termination condition</u>
 - HTN planning allows <u>recursion</u>
 - Move the <u>topmost</u> container (we know how to do that!)
 - Then move the <u>rest</u>
 - First attempt:
 - **task:** move-stack(pile1, pile2)
 - method: recursive-move(pile1, pile2)
 - **precond**: true
 - subtasks: <move-topmost-container(pile1, pile2), move-stack(pile1, pile2)>



Recursion (2)



But consider the BW and DWR "pile models"...



The bottom block is not "on" anything

The bottom block is "on" the pallet, a "special container"

What if the pallet is "topmost"? We don't want to move it!

Recursion (3)



• To fix this:

Add two method params – "non-natural", as in "ordinary" planning; does not give the planner a real choice

- <u>**Task</u>**: move-stack(pile1, pile2)
 </u>
 - method: recursive-move(pile1, pile2, cont, x)
 - precond: top(cont, pile1), on(cont, x)
 - subtasks: <move-topmost-container(pile1, pile2), move-stack(pile1, pile2)>

cont is on top of something (x), so cont can't be the pallet

Recursion (4)



The planner can now create a structure like this:



But when will the recursion end?

Recursion (5)

take-and-put(...)

take(...)



- At some point, only the pallet will be left in the stack
 - Then recursive-move will not be applicable

put(...)

But we <u>must</u> execute <u>some</u> form of move-stack!





move-topmost-container(pile1, pile2)



move-stack(pile1, pile2)

pile1 is empty! No applicable methods... Planner would backtrack!

Recursion (6)





DWR Example: Moving a stack, in the same order

Ordering (1)



Using move-stack inverts a stack:



Ordering (2)



• To avoid this: Use an intermediate pile



Ordering (3)



• Example:

Task:

- move-stack-same-order(pile1, pile2)
- method: move-each-twice(pile1, pileX, pile2, loc)
 - precond: top(pallet, pileX),
 pile1 != pileX, pile2 != pileX, pile1 != pile2,
 attached(...), // All in the same location

. . .

Unlike classical planning, someone specifies the task + pile1 and pile2

The planner must choose a matching method ("implementation") to use

The planner must choose added method params *pileX* and *loc* to satisfy the precond

subtasks: ; move twice:
 <move-stack(pile1, pileX), move-stack(pile2)>

Why does **pileX** have to be empty initially?

Because the second *move-stack* moves *all* containers from the intermediate pile...

DWR Example: Moving Three Stacks

Overall Objective



- Our overall **objective** is:
 - Moving three entire stacks of containers, preserving order



Overall Objective: Defining a Task

Define a <u>task</u> for this objective

- <u>Task:</u> move-three-stacks()
 - method: move-each-twice()
 - **precond**: ; no preconditions apart from the subtasks'
 - subtasks: ; move each stack twice:
 <move-stack-same-order(p1a,p1c), move-stack-same-order(p2a,p2c), move-stack-same-order(p3a,p3c) >

 Use this task as the initial task network







DWR Example: Moving *n* stacks

Goal Predicates in HTNs



- Here the <u>entire</u> objective was encoded in the initial network
 - move-three-stacks
 - <move-stack-same-order(p1a,p1c), move-stack-same-order(p2a,p2c), move-stack-same-order(p3a,p3c) >

• To avoid this:

- New predicate should-move-same-order(pile, pile) encoding the goal
- <u>Task:</u> move-as-necessary()
 - method: move-and-repeat(pile1, pile2)
 - precond: should-move-same-order(pile1, pile2)
 - subtasks: <move-stack-same-order(pile1, pile2), ;; makes should-move... false! move-as-necessary>
- **<u>Task:</u>** move-as-necessary()
 - method: all-done
 - **precond**: not exists pile1, pile2 [should-move-same-order(pile1, pile2)]
 - subtasks: <>

Uninformed Planning in HTNs

Can even do <u>uninformed unguided planning</u>

- Doing something, anything:
 - Task <u>do-something</u>
 - Task <u>do-something</u>
 - Task <u>do-something</u>
 - Task <u>do-something</u>
- Repeating:
 - Task <u>achieve-goals</u>

- → operator <u>pickup(x)</u>
- → operator **<u>putdown(x)</u>**
- \rightarrow operator <u>stack(x,y)</u>
- → operator <u>unstack(x,y)</u>

Planner chooses all parameters

→ <do-something, achieve-goals>

- Ending:
 - Task <u>achieve-goals</u>

→ <>, with precond: entire goal is satisfied

Or combine <u>aspects</u> of this model with <u>other aspects</u> of "standard" HTN models!

Useful Modeling Strategies:

Delivery Example – Delivering a package

Modeling "conditional" actions

Delivery 1: First Variation

Delivery:

- A single truck
- Pick up a package, drive to its destination, unload
- **Task:** deliver(package, dest)
 - method: move-by-truck(package, packageloc, dest)
 - precond: at(package, packageloc)
 - subtasks: <driveto(packageloc), load(package), driveto(dest), unload(package)>

What if the truck is already *at* the package location? First driveto is unnecessary!



Delivery 2: Second Variation

- **<u>Alternative</u>**: Two alternative methods for *deliver*
 - **Task:** deliver(*package, dest*)
 - method: move-by-truck-1(*package*, *packageloc*, *truckloc*, *dest*)
 - precond: at(truck, truckloc), at(package, packageloc), packageloc = truckloc
 - subtasks: <load(package), driveto(dest), unload(package)>
 - Task: deliver(package, dest)
 - method: move-by-truck-2(*package*, *packageloc*, *truckloc*, *dest*)
 - precond: at(*truck*, *truckloc*), at(*package*, *packageloc*),
 packageloc != *truckloc*
 - subtasks: <driveto(packageloc), load(package), driveto(dest), unload(package)>

Do we really have to repeat the entire task? Many "conditional" subtasks -> combinatorial explosion

Delivery 3: Third variation

- Make the choice in the subtask instead!
 - **Task**: deliver(*package, dest*)
 - method: move-by-truck-3(package, packageloc, truckloc, dest)
 - **precond**: at(truck, truckloc), at(*package*, *packageloc*)
 - subtasks: <be-at(packageloc), load(package), be-at(dest), unload(package)>
 - Task: be-at(loc)
 - method: drive(loc)
 - **precond**: !at(truck, *loc*)
 - subtasks: <driveto(loc)>
 - Task: be-at(loc)
 - **method**: already-there
 - **precond**: at(truck, *loc*)
 - subtasks: <>

One possible search space for Total Order Simple Task Networks:

Total Order Forward Decomposition – TFD

Search Spaces



Need a <u>search space</u>

I) A <u>node structure</u> defining what information is in a node

2) A way of creating an <u>initial node</u> from a problem instance

3) A <u>successor function</u> / branching rule returning all successors

4) A <u>solution criterion</u>, detecting if a node corresponds to a solution

5) A <u>plan extractor</u>, telling us which plan a solution node corresponds to

Different alternatives exist!

Total Order?



- Basic <u>assumption</u>: Total Order Simple Task Networks
 - Any initial task is totally ordered
 - Any <u>decomposition method</u> is totally ordered



Forward Decomposition?

Different <u>decomposition orders</u> are still possible:



3. Multiple alternatives; which to decompose next?

Choose what to decompose, which method to use, how to parameterize it → Need search!

Forward Decomposition!

- **Forward** decomposition: One of many possibilities
 - Go "depth first, left to right"
 - Like forward state space search:
 - Generates actions in the same order in which they'll be executed
 - → When we decompose a task, we know the "current" state of the world!



Total Order Forward Decomposition





Total Order Forward Decomposition: The Search Space

TFD Node Structure



- [1] A node structure defining what information is in a node
 - Plan so far
 - Current state possible due to forward decomposition
 - Remaining tasks to expand



TFD Successors



[3] Successors:

- We know which task to decompose —
- Find all applicable methods and apply them

take(...)put(...)S2mtc(pile1, pile2)move-stack(pile1, pile2)

- [4] Solution test
 - No more tasks → done
- [5] Solution extraction
 - The resulting search node *contains* a sequential plan

Total Order Forward Decomposition: The Search Space

Solving Total-Order STN Problems

- TFD takes a search node
 - π a sequence of actions
 - s the current state
 - <t1,...,tk> a list of tasks to be achieved in the specified order
- We also assume:
 - O the available operators (with params, preconds, effects)
 - M the available methods (with params, preconds, subtasks)

Returns:

- A sequential plan
 - Loses the hierarchical structure of the final plan
 - Simplifies the presentation but the structure *could* also be kept!

TFD 1: Base case



total-order-forward-decomposition(problem) {

 $\begin{array}{l} \textit{initial-node} \leftarrow \langle [], \textit{problem.initialstate, problem.initialtask} // [2] \\ \textit{open} \leftarrow \{\textit{initial-node} \} \\ \underline{\textit{while}} (\textit{open} \neq \emptyset) \\ \{ \\ \langle \pi, s, < t_1, \dots, t_k > \rangle \leftarrow \textit{search-strategy-remove-from}(\textit{open}) \\ // [6] \textit{TFD uses depth first} \end{array}$

// [4] If we have no tasks left to decompose... <u>if</u> k = 0 then <u>return</u> π // [5] plan extraction

TFD 2: Ground Primitive Tasks

56

total-order-forward-decomposition(problem) {

// initial-node, open **while** (open $\neq \emptyset$) { $\langle \pi, s, < t_1, ..., t_k \rangle \rangle \leftarrow ...$ **<u>if</u>** k = 0 **then** <u>return</u> π For simplicity: The case where all **tasks to achieve** are **ground**

if $(t_1 \text{ is primitive})$ then// A primitive task is decomposed into a single action!// May be many to choose from (e.g. method has more params than task).actions \leftarrow ground instances of operators in Ocandidates \leftarrow { a |a \in actions andname(a) = t_1 anda is applicable in s }



TFD 3: Successors



total-order-forward-decomposition(problem) {

// initial-node, open **while** (open $\neq \emptyset$) { $\langle \pi, s, < t_1, ..., t_k \rangle \rangle \leftarrow ...$ **<u>if</u>** k = 0 **then** <u>return</u> π

For simplicity: The case where all **tasks to achieve** are **ground**

if $(t_1 ext{ is primitive})$ then// A primitive task is decomposed into a single action!// May be many to choose from (e.g. method has more params than task).actions \leftarrow ground instances of operators in Ocandidates \leftarrow { a | a \in actions andname(a) = t_1 anda is applicable in s }

for all $a \in candidates$:

π'	$\leftarrow \pi + a$	// Add action at the end
s'	$\leftarrow \gamma(s, a)$	// Apply the action, find the new state
rest	← <t2,,tk></t2,,tk>	
open	$\leftarrow open \cup \{ \langle \pi', s', rest \rangle \}$	



Total Order Forward Decomposition: Non-Ground Primitive Tasks

TFD 4: Lifted Primitive Tasks



total-order-forward-decomposition(problem) {

The case where tasks to achieve are **non-ground**: move(container1, *X*)

The **plan** will still be **ground**!



TFD 5: Lifted Primitive Successors

total-order-forward-decomposition(problem) {

// initial-node, open

while (open $\neq \emptyset$) {

 $\langle \pi, s, \langle t_1, \dots, t_k \rangle \rangle \leftarrow \dots$

The case where tasks to achieve are **non-ground**: move(container1, *X*)

if k = 0 then return π **<u>if</u>** (t_1 is primitive) **<u>then</u>** // A primitive task is decomposed into a single action! // May be many to choose from (e.g. method has more params than task). \leftarrow ground instances of operators in O actions candidates \leftarrow { (a, σ) | a \in actions and σ is a substitution s.t. name(a) = $\sigma(t_1)$ and a is applicable in s } **for all** $(a, \sigma) \in candidates:$ π' $\leftarrow \pi + a$ // Add action at the end $\leftarrow \gamma(s, a)$ // Apply the action, find the new state s'*rest* $\leftarrow \sigma(\langle t2,...,tk \rangle) //$ Must have the same variable bindings! open \leftarrow open $\cup \{ \langle \pi', s', rest \rangle \}$ $\sigma(t2) =$ put(<u>crane1</u>, ...) (*italics* = variables) t1 = take(*crane*, loc1, cont2, *cont*, pile8) t2=put(*crane*, ...) S **chosen:** a = take(**crane1**, loc1, cont2, **cont5**, pile8) $\sigma = \{ crane \mapsto crane1, cont \mapsto cont5 \}$ take(crane2, loc1, cont2, cont5, pile8) $\{ crane \mapsto crane2, cont \mapsto cont5 \}$

Total Order Forward Decomposition: (Non-Ground) Non-Primitive Tasks

TFD 6: Non-Primitive Tasks

62 billowide

total-order-forward-decomposition(problem) {

// initial-node, open **while** (open $\neq \emptyset$) { $\langle \pi, s, < t_1, ..., t_k > \rangle \leftarrow ...$ **if** k = 0 **then** <u>return</u> π **if** $(t_1$ is primitive) <u>**then** ...</u>





TFD 6: Non-Primitive Tasks

total-order-forward-decomposition(problem) {

// initial-node, open while (open $\neq \emptyset$) { $\langle \pi, s, < t_1, ..., t_k \rangle \leftarrow ...$ if k = 0 then return π if $(t_1 \text{ is primitive})$ then ...

else // t1 is travel(LiU, Resecentrum), for example // A non-primitive task is decomposed into a new task list. // May have many methods to choose from: taxi-travel, bus-travel, walk, ... ground \leftarrow ground instances of methods in M candidates \leftarrow { (m, σ) | m \in ground and σ is a substitution s.t. task(m) = $\sigma(t_1)$ and m is applicable in s } // Methods have preconds! <u>for all</u> (m, σ) \in candidates: $\pi' \leftarrow \pi$ // No action added!

s' \leftarrow s // No state change! rest \leftarrow subtasks(m) + $\sigma(\langle t2,...,tk \rangle)$ // Prepend new list!

open
$$\leftarrow$$
 open $\cup \{ \langle \pi', s', rest \rangle \}$

In TFD

the "origin" of a task is discarded: No longer needed, only the subtasks are relevant





Limitations of Total-Order HTN Planning

Limitation of Ordered-Task Planning

- TFD requires totally ordered methods
 - Can't interleave subtasks of different tasks
- Suppose we want to <u>fetch one object</u> somewhere, then return to where we are now
 - Task: <u>fetch</u>(obj)
 - method: <u>get</u>(obj, mypos, objpos)
 - precond: <u>robotat</u>(mypos) & at(obj, objpos)
 - subtasks: <<u>travel(mypos, objpos)</u>, <u>pickup(obj)</u>, <u>travel(objpos, mypos)</u>>
 - Task: <u>travel(x, y)</u>
 - method: <u>walk</u>(x, y)
 - method: <u>stayat(x)</u>



Limitation of Ordered-Task Planning

Suppose we want to fetch <u>two</u> objects somewhere, and return

66

- (Simplified example consider "fetching all the objects we need")
- One idea: Just "fetch" each object in sequence
 - Task: <u>fetch-both</u>(obj1, obj2)
 - method: <u>get-both</u>(obj1, obj2, mypos, objpos1, objpos2)
 - precond:
 - subtasks: <<u>fetch</u>(obj1, mypos, objpos1), <u>fetch</u>(obj2, mypos, objpos2)>



Alternative Methods



- To generate more efficient plans using total-order STNs:
 - Use a different domain model!

method:

- Task: <u>fetch-both</u>(obj1, obj2)
 - **get-both**(obj1, obj2, mypos, objpos1, objpos2)
 - precond: objpos1 != objpos2 & at(obj1, objpos1) & at(obj2, objpos2)
 - subtasks: <<u>travel</u>(mypos, objpos1), <u>pickup</u>(obj1), travel(objpos1, objpos2), <u>pickup</u>(obj2), <u>travel</u>(objpos2, mypos)>
- Task: <u>fetch-both</u>(obj1, obj2)
 - method: get-both-in-same-place(obj1, obj2, mypos, objpos)
 - precond: <u>robotat</u>(mypos) & at(obj1, objpos) & at(obj2, objpos)
 - subtasks: <<u>trav</u>
- <<u>travel</u>(mypos, objpos), <u>pickup</u>(obj1), <u>pickup</u>(obj2), <u>travel</u>(objpos, mypos)>

Or: load-all; drive-truck; unload-all

HTN Planning with Partially Ordered Methods

Partially Ordered Methods

Partially ordered method:

• The subtasks are a **partially ordered** set $\{t_1, ..., t_k\}$ – a network



Partially Ordered Methods

With partially ordered methods, subtasks can be interleaved

onkv@ida



- Requires a more complicated planning algorithm: PFD
- SHOP2: implementation of PFD-like algorithm + generalizations