

Width and Complexity of Belief Tracking in Non-Deterministic Conformant and Contingent Planning

Blai Bonet¹ and Hector Geffner²

¹Universidad Simón Bolívar

²ICREA & Universitat Pompeu Fabra

AAAI, Toronto, Canada, July 2012



Motivation

Planning in the **non-deterministic** and **partially observable** setting

Setting is similar to qualitative POMDPs, where uncertainty is encoded by **sets of states** rather than probability distributions

Need to solve two **fundamental tasks**, both intractable for problems in compact form:

1. representation and tracking of belief states
2. planning (searching) for goals in belief space

Main Contributions

We focus on **belief tracking**:

1. Palacios and Geffner (2009) showed that belief tracking for **deterministic** conformant planning is exponential in a **width** parameter that is often bounded and small
2. Results extended to **deterministic** contingent planning by Albore, Palacios and Geffner (2009)
3. This paper generalizes these results to **non-deterministic** conformant and contingent planning for which **new and effective belief tracking algorithms** are developed
4. Purely semantic approach (no translations involved)

Model for Non-Deterministic Contingent Planning

Contingent model $\mathcal{S} = \langle S, S_0, S_G, A, F, O \rangle$ given by

- finite **state space** S
- non-empty subset of **initial states** $S_0 \subseteq S$
- non-empty subset of **goal states** $S_G \subseteq S$
- **actions** A where $A(s) \subseteq A$ are the actions applicable at state s
- **non-deterministic** transition function $F(s, a) \subseteq S$ for $s \in S, a \in A(s)$
- **non-deterministic** sensor model $O(s', a) \subseteq O$ for $s' \in S, a \in A$

Language: Factored Representation of the Model

Model expressed in **compact form** as tuple $P = \langle V, A, I, G, V', W \rangle$ where

- V is set of multi-valued variables, each X has finite domain D_X
- A is set of actions; each action $a \in A$ has precondition $Pre(a)$ and conditional **non-deterministic** effects $C \rightarrow E^1 | \dots | E^n$
- Sets of V -literals I and G defining the initial and goal states
- V' is set of observable variables (not necessarily disjoint from V). Observations o are **valuations** over V'
- **Sensing model** is formula $W_a(\ell)$ for each $a \in A$ and observable literal ℓ that tells the states that may be obtained after applying a

Note: a literal is an atom of the form ' $X = x$ ' or ' $X \neq x$ '

From Language to Model

- states S are valuations over state variables V
- initial states S_0 are states that satisfy the clauses in I
- goal states S_G are states that satisfy the literals in G
- action $A(s)$ applicable at s are those whose precondition hold at s
- transition function $F(s, a)$ defined as in (non-det) planning
- observations o are valuations over observable variables V'
- observation $o \in O(s, a)$ iff $s \models W_a(\ell)$ for each literal ℓ with $o \models \ell$

Basic Algorithm: Flat Belief Tracking

Explicit representation of beliefs states as sets of states

Definition (Flat Tracking)

Given belief b at time t , and action a (applied) and observation o (obtained), the belief at time $t + 1$ is the belief b_a^o given by

$$b_a = \{s' : s' \in F(s, a) \text{ and } s \in b\}$$

$$b_a^o = \{s' : s' \in b_a \text{ and } s' \models W_a(\ell) \text{ for each } \ell \text{ s.t. } o \models \ell\}$$

- Flat belief tracking is sound and complete for **every formula**
- Time complexity is **exponential in $|V \cap V_U|$** where $V_U = V \setminus V_K$ and V_K are the variables that are **always known**
- In planning, however, only need to check **preconditions and goals**

Belief Tracking in Planning (BTP)

Definition

Given execution $\tau = \langle a_0, o_0, a_1, o_1, \dots, a_n, o_n \rangle$ and precondition or goal literal ℓ , **determine** whether

- execution τ is possible, and
- if τ is possible, whether b_τ , the belief that results of executing τ , makes literal ℓ true

Note: contingent setting has the conformant setting as a special case

Factored Belief Tracking: Roadmap

- 1 Show that Belief Tracking in Planning for problem P can be **decomposed** into belief tracking for **subproblems** P_X for each variable X that is a precondition or goal variable
- 2 Moreover, a **width** parameter $width(P)$ can be defined so that the size (# of vars) of all subproblems P_X is bounded by $width(P)$
- 3 **Fundamental property:** a literal ' $X = x$ ' is true in P after a possible execution τ iff it is true in subproblem P_X after τ
- 4 Thus, flat belief tracking over each subproblem P_X yields an **algorithm for belief tracking in planning** for problem P that is exponential in $width(P)$

Next: define subproblems P_X and $width(P)$ from structure of P

Causal Relevance

Definition (Direct Cause)

Variable X is **direct cause** of Y if $X \neq Y$, and either:

- there is an effect $C \rightarrow E^1 | \dots | E^n$ such that X occurs in C and Y occurs in some E^i , or
- X occurs in some formula $W_a(Y = y)$ for obs var $Y \in V'$

Definition

Variable X is **causally relevant** to Y if $X = Y$, X is a direct cause of Y , or X is causally relevant to Z that is causally relevant to Y

I.e., causally relevant is the smaller transitive and reflexive relation that includes the direct cause relation

Relevance and Contexts

The relevance relation captures causal and evidential relations due to **observations**

Definition

Variable X is **relevant** to Y if either:

- a) X is causally relevant to Y ,
- b) both X and Y are causally relevant to an observable variable Z , or
- c) X is relevant to Z that is relevant to Y

Definition (Contexts)

The **context** of variable X , $Ctx(X)$, is the set of **state variables** that are relevant to X

Width


Definition (Width of Variable)

The **width** of variable X is the number of variables in its context that are not known: $width(X) = |Ctx(X) \cap V_U|$ where $V_U = V \setminus V_K$

Definition (Width)


The **width** of a problem is $width(P) = \max_X width(X)$ where X ranges over the **goal or precondition variables**

Example: NON-DET-Ring-Key



- windows W_1, \dots, W_n that can be open, closed, or locked
- agent doesn't know its position, windows' status, or key position
- goal is to have **all windows locked**
- when unlocked, **windows open/close non-det.** when agent moves
- to lock window: must close and then lock it with **key**
- key's position is unknown and must be **grabbed** to lock windows
- **possible plan:** repeat n times $\langle \text{Grab, Fwd} \rangle$ followed by repeat n times $\langle \text{Close, Lock, Fwd} \rangle$

Example: NON-DET-Ring-Key



- Variables:
 - ▶ windows' status: $W_i \in \{open, closed, locked\}$
 - ▶ position of agent (Loc) and key (KLoc)
- Actions:
 - ▶ Close: $W_i = open, Loc = i \rightarrow W_i = closed$
 - ▶ Lock: $W_i = closed, Loc = i, KLoc = hand \rightarrow W_i = locked$
 - ▶ Grab: $Loc = i, KLoc = i \rightarrow KLoc = hand$
 - ▶ Fwd: $Loc = i \rightarrow Loc = i + 1 \pmod n$
 $W_i \neq locked \rightarrow W_i = open \mid W_i = closed$
- Contexts: $Ctx(W_i) = \{W_i, Loc, KLoc\}$, $width(W_i) = 3$, $width(P) = 3$

Subproblems P_X

Subproblem P_X is problem P **projected** on the vars in $Ctx(X)$

Basically, P_X has:

- variables $Ctx(X)$ but same observable variables V'
- only precondition and effects relevant to $Ctx(X)$ are kept
- sensing formulas $W_a(Y = y)$ are **logically projected** on $Ctx(X)$

Theorem (Flat Belief Tracking on P_X)

Flat belief tracking on P_X is exponential in $width(X)$ which is less than or equal to $width(P)$ for precondition or goal variable X

Factored Belief Tracking: Properties

Theorem

- 1) *an execution $\tau = \langle a_0, o_0, \dots \rangle$ is possible in P iff it is possible over all subproblems P_X for goal or precondition variables X*
- 2) *a literal $X = x$ or $X \neq x$ is known in belief state b that results from possible execution τ on P iff it is known to be true in the belief b_X that results from the same execution on P_X*

Theorem (Soundness and Completeness)

Factored belief tracking over subproblems P_X , for precondition or goal variable X , is a sound and complete tracking algorithm for planning

Theorem (Complexity)

Complexity of factored belief tracking is exponential in $width(P)$

Experiments: Conformant Ring

| n | steps | exp. | time |
|-----|-------|-------|-------|
| 10 | 68 | 355 | < 0.1 |
| 20 | 138 | 705 | 0.1 |
| 30 | 208 | 1,055 | 0.9 |
| 40 | 277 | 1,400 | 3.1 |
| 50 | 345 | 1,740 | 8.3 |
| 60 | 415 | 2,090 | 18.6 |
| 70 | 476 | 2,395 | 34.5 |
| 80 | 545 | 2,740 | 62.8 |
| 90 | 610 | 3,065 | 106.4 |
| 100 | 679 | 3,410 | 171.0 |

DET-Ring-Key

| n | steps | exp. | time |
|-----|-------|-------|-------|
| 10 | 118 | 770 | < 0.1 |
| 20 | 198 | 1,220 | 0.8 |
| 30 | 278 | 1,670 | 4.2 |
| 40 | 488 | 3,210 | 15.2 |
| 50 | 438 | 2,570 | 34.4 |
| 60 | 468 | 2,660 | 52.2 |
| 70 | 543 | 3,080 | 100.6 |
| 80 | 616 | 3,480 | 172.9 |
| 90 | 682 | 3,880 | 285.6 |
| 100 | 1,111 | 7,220 | 783.1 |

NON-DET-Ring-Key

- Solved with a greedy A* algorithm with eval function $f(n) = h(n)$
- Heuristic is $h(b) = \sum_{i=1}^n h(b_i)$ where $h(b_i)$ is fraction of states in projection over $Ctx(W_i)$ where $W_i \neq locked$
- Planner KACMBP by Cimatti et al. (2004) solves up to 20 windows, planner T0 cannot be used because problem is non-det

Experiments: Variation of Wumpus

| dimension | #objects | avg. steps | avg. time |
|-----------|----------|-------------|-----------------|
| 10 × 10 | 0 | 57.4 ± 46 | 43.6 ± 37 |
| 10 × 10 | 1 | 137.6 ± 204 | 113.7 ± 167 |
| 10 × 10 | 2 | 145.8 ± 200 | 195.7 ± 259 |
| 10 × 10 | 3 | 191.2 ± 177 | 538.0 ± 438 |
| 10 × 10 | 4 | 114.0 ± 57 | 953.6 ± 506 |
| 10 × 10 | 5 | 48.0 ± 34 | 1,552.6 ± 1,001 |
| 10 × 10 | 6 | 129.6 ± 105 | 8,714.7 ± 4,716 |

- Agent navigates grid, searching for gold while avoiding pits and wumpus
- Agent gets signal when next to hazard or at same cell of gold
- Each hazard (either wumpus or pit) has **unique** feedback signal
- Solved with action selection mechanism based on a lookahead tree of fixed depth, explored with Anytime AO* (Bonet & Geffner, AAAI-12)

Related Work: Other Accounts of Width

Accounts of Palacios and Geffner, and Albore et al.:

- they deal with with deterministic problems
- our notion of width is similar but not equivalent on deterministic problems:
 - ▶ newer notion is simpler
 - ▶ it is defined over multi-valued variables
 - ▶ but, it is slightly less tight when initial uncertainty does not encode multi-valued variables (see paper)

Related Work: Bayesian Networks

Notion of width is not the same as in BNs:

- exploit knowledge that some variables are **not observable**
- exploit knowledge that some variables are **always known**
- **difference** between preconditions and conditions of effects

Notion of relevance is also different:

- not necessarily symmetric as in BNs
- influenced by which variables are observable or not

Summary and Future Work

- First account of **width** (as far as we know) in **non-deterministic** conformant and contingent planning
- Factored belief tracking that is **sound and complete** for planning (i.e., wrt preconditions and goal literals; **not every formula**)
- Time complexity of factored belief tracking is exponential in the width and (low) polynomial on the rest of the parameters
- **Currently working** on approximate tracking for problems with unbounded width (e.g., Battleship, Minesweeper, Wumpus, etc.)