

Master in Artificial Intelligence and Robotics (AIRO)

Electives in AI Reasoning Agents

Fabio Patrizi

Sapienza University of Rome, Italy
patrizi@diag.uniroma1.it

A.Y. 2021-2022

Planning for LTL_f and LDL_f Goals*

*slides based on G. De Giacomo's (www.diag.uniroma1.it/degiacomo)

- 1 LTL_f/LDL_f : LTL/LDL on finite traces
- 2 LTL_f/LDL_f and automata
- 3 Planning for LTL_f/LDL_f goals: deterministic domains
- 4 $FOND_{sp}$ for LTL_f/LDL_f goals: nondeterministic domains
- 5 Conclusion

- 1 LTL_f/LDL_f : LTL/LDL on finite traces
- 2 LTL_f/LDL_f and automata
- 3 Planning for LTL_f/LDL_f goals: deterministic domains
- 4 $FOND_{sp}$ for LTL_f/LDL_f goals: nondeterministic domains
- 5 Conclusion

LTL_f: the language

$$\varphi ::= A \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \bigcirc\varphi \mid \varphi_1 \mathcal{U} \varphi_2$$

- A : atomic propositions
- $\neg\varphi, \varphi_1 \wedge \varphi_2$: boolean connectives
- $\bigcirc\varphi$: “next step exists and at next step (of the trace) φ holds”
- $\varphi_1 \mathcal{U} \varphi_2$: “eventually φ_2 holds, and φ_1 holds until φ_2 does”
- $\bullet\varphi \doteq \neg\bigcirc\neg\varphi$: “if next step exists then at next step φ holds” (*weak next*)
- $\diamond\varphi \doteq \text{true} \mathcal{U} \varphi$: “ φ will eventually hold”
- $\square\varphi \doteq \neg\diamond\neg\varphi$: “from current till last instant φ will always hold”
- $\text{Last} \doteq \neg\bigcirc\text{true}$: denotes last instant of trace.

Main formal properties:

- **Expressibility:** FOL over finite sequences or Star-free RE
- **Reasoning:** satisfiability, validity, entailment PSPACE-complete
- **Model Checking:** linear on TS, PSPACE-complete on formula

Some interesting LTL_f formulas:

<i>name of template</i>	<i>LTL semantics</i>
<i>responded existence(A, B)</i>	$\Diamond A \Rightarrow \Diamond B$
<i>co-existence(A, B)</i>	$\Diamond A \Leftrightarrow \Diamond B$
<i>response(A, B)</i>	$\Box(A \Rightarrow \Diamond B)$
<i>precedence(A, B)</i>	$(\neg B \ U \ A) \vee \Box(\neg B)$
<i>succession(A, B)</i>	$\text{response}(A, B) \wedge \text{precedence}(A, B)$
<i>alternate response(A, B)</i>	$\Box(A \Rightarrow \bigcirc(\neg A \ U B))$
<i>alternate precedence(A, B)</i>	$\text{precedence}(A, B) \wedge \Box(B \Rightarrow \bigcirc(\text{precedence}(A, B)))$
<i>alternate succession(A, B)</i>	$\text{alternate response}(A, B) \wedge \text{alternate precedence}(A, B)$
<i>chain response(A, B)</i>	$\Box(A \Rightarrow \bigcirc B)$
<i>chain precedence(A, B)</i>	$\Box(\bigcirc B \Rightarrow A)$

<i>name of template</i>	<i>LTL semantics</i>
<i>not co-existence(A, B)</i>	$\neg(\Diamond A \wedge \Diamond B)$
<i>not succession(A, B)</i>	$\Box(A \Rightarrow \neg(\Diamond B))$
<i>not chain succession(A, B)</i>	$\Box(A \Rightarrow \bigcirc(\neg B))$

<i>name of template</i>	<i>LTL semantics</i>
<i>existence(1, A)</i>	$\Diamond A$
<i>existence(2, A)</i>	$\Diamond(A \wedge \bigcirc(\text{existence}(1, A)))$
...	...
<i>existence(n, A)</i>	$\Diamond(A \wedge \bigcirc(\text{existence}(n-1, A)))$
<i>absence(A)</i>	$\neg \text{existence}(1, A)$
<i>absence(2, A)</i>	$\neg \text{existence}(2, A)$
<i>absence(3, A)</i>	$\neg \text{existence}(3, A)$
...	...
<i>absence(n+1, A)</i>	$\neg \text{existence}(n+1, A)$
<i>init(A)</i>	A

LDL_f: the language
$$\varphi ::= \text{tt} \mid \text{ff} \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \langle \rho \rangle \varphi \mid [\rho] \varphi \quad \rho ::= \phi \mid \varphi? \mid \rho_1 + \rho_2 \mid \rho_1; \rho_2 \mid \rho^*$$

- **tt** and **ff** stand for **true** and **false**
- ϕ : **propositional formula** on current state/instant
- $\neg\varphi, \varphi_1 \wedge \varphi_2$: **boolean connectives**
- ρ is a **regular expression** on propositional formulas
- $\langle \rho \rangle \varphi$: **exists** an “execution” of RE ρ that ends with φ holding
- $[\rho] \varphi$: **all** “executions” of RE ρ (*along the trace!*) end with φ holding

In the infinite trace setting, such enhancement strongly advocated by industrial model checking (ForSpec, PSL).

Main formal properties:

- **Expressibility**: **MSO** over finite sequences: adds the power of recursion (as RE)
- **Reasoning**: **satisfiability**, **validity**, **entailment** PSPACE-complete
- **Model Checking**: **linear** on TS, PSPACE-complete on formula

Example

- All coffee requests from person p will eventually be served:

$$[\text{true}^*](\text{request}_p \supset \langle \text{true}^* \rangle \text{coffee}_p)$$

- Every time the robot opens door d it closes it immediately after:

$$[\text{true}^*](\text{openDoor}_d \supset \text{closeDoor}_d)$$

- Before entering restricted area a the robot must have permission for a :

$$\langle (\neg \text{inArea}_a^* ; \text{getPermission}_a ; \neg \text{inArea}_a^* ; \text{inArea}_a)^* ; \neg \text{inArea}_a^* \rangle \text{end}$$

Note that the first two properties (not the third one) can be expressed also in LTL_f:

$$\Box(\text{request}_p \supset \Diamond \text{coffee}_p)$$

$$\Box(\text{openDoor}_d \supset \bigcirc \text{closeDoor}_d)$$

LDL_f: Linear Dynamic Logic on finite traces

LDL_f, not LTL_f, is able to easily express procedural constraints [BaierFritzMcIlraith07].

Let's introduce a sort of propositional variant of GOLOG

$$\delta ::= A \mid \varphi? \mid \delta_1 + \delta_2 \mid \delta_1; \delta_2 \mid \delta^* \mid \text{if } \phi \text{ then } \delta_1 \text{ else } \delta_2 \mid \text{while } \phi \text{ do } \delta$$

where **if** and **while** can be seen as abbreviations for LDL_f path expression, namely:

$$\text{if } \phi \text{ then } \delta_1 \text{ else } \delta_2 \doteq (\phi?; \delta_1) + (\neg\phi?; \delta_2) \quad \text{while } \phi \text{ do } \delta \doteq (\phi?; \delta)^*; \neg\phi?$$

Example (LDL_f procedural constraints)

- “At every point, if it is hot then, if the air-conditioning system is off, turn it on, else don't turn it off”:

$$[\text{true}^*](\text{if } (\text{hot}) \text{ then} \\ \quad \text{if } (\neg\text{airOn}) \text{ then } \text{turnOnAir} \\ \quad \text{else } \neg\text{turnOffAir}) \text{true}$$

- “alternate till the end the following two instructions: (1) while is hot if the air-conditioning system is off turn it on, else don't turn it off; (2) do something for one step”

$$\langle (\text{while } (\text{hot}) \text{ do} \\ \quad \text{if } (\neg\text{airOn}) \text{ then } \text{turnOnAir} \\ \quad \text{else } \neg\text{turnOffAir}; \\ \text{true})^* \rangle \text{end}$$

Example (LDL_f captures finite domain variant of GOLOG in SitCalc)

GOLOG – finite domain variant

$$\delta ::= A \mid \varphi? \mid \delta_1 + \delta_2 \mid \delta_1; \delta_2 \mid \delta^* \mid \pi x. \delta(x) \mid \text{if } \phi \text{ then } \delta_1 \text{ else } \delta_2 \mid \text{while } \phi \text{ do } \delta$$

- $\pi x. \delta(x)$ stands for $\sum_{o \in Obj} \delta(o)$
- $\text{if } \phi \text{ then } \delta_1 \text{ else } \delta_2$ stands for $(\phi?; \delta_1) + (\neg\phi?; \delta_2)$
- $\text{while } \phi \text{ do } \delta$ stands for $(\phi?; \delta)^* \neg\phi?$

- $\langle \delta \rangle \phi$ in LDL_f captures SitCalc formula $\exists s'. Do(\delta, s, s') \wedge s \leq s' \leq last \wedge \phi(s')$.
- $[\delta] \phi$ in LDL_f captures SitCalc formula $\forall s'. Do(\delta, s, s') \wedge s \leq s' \leq last \supset \phi(s')$.

($\phi(s)$ "uniform" in s .)

- 1 LTL_f/LDL_f : LTL/LDL on finite traces
- 2 LTL_f/LDL_f and automata
- 3 Planning for LTL_f/LDL_f goals: deterministic domains
- 4 $FOND_{sp}$ for LTL_f/LDL_f goals: nondeterministic domains
- 5 Conclusion

Key point

LTL_f/LDL_f formulas can be translated into **nondeterministic finite state automata (NFA)**.

$$t \models \varphi \text{ iff } t \in \mathcal{L}(\mathcal{A}_\varphi)$$

where \mathcal{A}_φ is the NFA φ is translated into.

We can compile reasoning into automata based procedures!

Both LTL_f and LDL_f formulas can be translated in **exponential time** to **nondeterministic automata on finite words (NFA)**.

NFA \mathcal{A}_φ associated with an LTL_f formula φ (in NNF)

Auxiliary rules

$$\begin{aligned}
 \delta(A, \Pi) &= \text{true if } A \in \Pi \\
 \delta(A, \Pi) &= \text{false if } A \notin \Pi \\
 \delta(\neg A, \Pi) &= \text{false if } A \in \Pi \\
 \delta(\neg A, \Pi) &= \text{true if } A \notin \Pi \\
 \delta(\varphi_1 \wedge \varphi_2, \Pi) &= \delta(\varphi_1, \Pi) \wedge \delta(\varphi_2, \Pi) \\
 \delta(\varphi_1 \vee \varphi_2, \Pi) &= \delta(\varphi_1, \Pi) \vee \delta(\varphi_2, \Pi) \\
 \delta(\bigcirc\varphi, \Pi) &= \begin{cases} \varphi & \text{if } \text{Last} \notin \Pi \\ \text{false} & \text{if } \text{Last} \in \Pi \end{cases} \\
 \delta(\diamond\varphi, \Pi) &= \delta(\varphi, \Pi) \vee \delta(\bigcirc\varphi, \Pi) \\
 \delta(\varphi_1 \mathcal{U} \varphi_2, \Pi) &= \delta(\varphi_2, \Pi) \vee (\delta(\varphi_1, \Pi) \wedge \delta(\bigcirc(\varphi_1 \mathcal{U} \varphi_2), \Pi)) \\
 \delta(\bullet\varphi, \Pi) &= \begin{cases} \varphi & \text{if } \text{Last} \notin \Pi \\ \text{true} & \text{if } \text{Last} \in \Pi \end{cases} \\
 \delta(\square\varphi, \Pi) &= \delta(\varphi, \Pi) \wedge \delta(\bullet\square\varphi, \Pi) \\
 \delta(\varphi_1 \mathcal{R} \varphi_2, \Pi) &= \delta(\varphi_2, \Pi) \wedge (\delta(\varphi_1, \Pi) \vee \delta(\bullet(\varphi_1 \mathcal{R} \varphi_2), \Pi))
 \end{aligned}$$

Observe these are the rules defining the transition function of the AFW!

Algorithm

```

algorithm LTLf2NFA
input LTLf formula  $\varphi$ 
output NFA  $\mathcal{A}_\varphi = (2^{\mathcal{P}}, S, \{s_0\}, \varrho, \{s_f\})$ 
 $s_0 \leftarrow \{\varphi\}$  ▷ single initial state
 $s_f \leftarrow \emptyset$  ▷ single final state
 $S \leftarrow \{s_0, s_f\}, \varrho \leftarrow \emptyset$ 
while ( $S$  or  $\varrho$  change) do

    if ( $q \in S$  and  $q' \models \bigwedge_{(\psi \in q)} \delta(\psi, \Pi)$ )
         $S \leftarrow S \cup \{q'\}$  ▷ update set of states
         $\varrho \leftarrow \varrho \cup \{(q, \Pi, q')\}$  ▷ update transition relation
    
```

NFA \mathcal{A}_φ associated with an LDL_f formula φ (in NNF)

Auxiliary rules

$\delta(\text{tt}, \Pi)$	=	true
$\delta(\text{ff}, \Pi)$	=	false
$\delta(\varphi_1 \wedge \varphi_2, \Pi)$	=	$\delta(\varphi_1, \Pi) \wedge \delta(\varphi_2, \Pi)$
$\delta(\varphi_1 \vee \varphi_2, \Pi)$	=	$\delta(\varphi_1, \Pi) \vee \delta(\varphi_2, \Pi)$
$\delta(\langle \phi \rangle \varphi, \Pi)$	=	$\begin{cases} \text{false} & \text{if } \Pi \not\models \phi \text{ or } \Pi = \epsilon \text{ (trace ended)} \\ \mathbf{e}(\varphi) & \text{o/w } (\phi \text{ propositional}) \end{cases}$
$\delta(\langle \psi? \rangle \varphi, \Pi)$	=	$\delta(\langle \psi \rangle \varphi, \Pi) \wedge \delta(\varphi, \Pi)$
$\delta(\langle \rho_1 + \rho_2 \rangle \varphi, \Pi)$	=	$\delta(\langle \rho_1 \rangle \varphi, \Pi) \vee \delta(\langle \rho_2 \rangle \varphi, \Pi)$
$\delta(\langle \rho_1; \rho_2 \rangle \varphi, \Pi)$	=	$\delta(\langle \rho_1 \rangle \langle \rho_2 \rangle \varphi, \Pi)$
$\delta(\langle \rho^* \rangle \varphi, \Pi)$	=	$\delta(\varphi, \Pi) \vee \delta(\langle \rho \rangle \mathbf{f}_{\langle \rho^* \rangle \varphi}, \Pi)$
$\delta([\phi] \varphi, \Pi)$	=	$\begin{cases} \text{true} & \text{if } \Pi \not\models \phi \text{ or } \Pi = \epsilon \text{ (trace ended)} \\ \mathbf{e}(\varphi) & \text{o/w } (\phi \text{ propositional}) \end{cases}$
$\delta([\psi?] \varphi, \Pi)$	=	$\delta(\text{nnf}(\neg \psi), \Pi) \vee \delta(\varphi, \Pi)$
$\delta([\rho_1 + \rho_2] \varphi, \Pi)$	=	$\delta([\rho_1] \varphi, \Pi) \wedge \delta([\rho_2] \varphi, \Pi)$
$\delta([\rho_1; \rho_2] \varphi, \Pi)$	=	$\delta([\rho_1][\rho_2] \varphi, \Pi)$
$\delta([\rho^*] \varphi, \Pi)$	=	$\delta(\varphi, \Pi) \wedge \delta([\rho] \mathbf{t}_{[\rho^*] \varphi}, \Pi)$
$\delta(\mathbf{f}_{\psi}, \Pi)$	=	false
$\delta(\mathbf{t}_{\psi}, \Pi)$	=	true

$\mathbf{e}(\varphi)$ replaces in φ all occurrences of \mathbf{t}_{ψ} and \mathbf{f}_{ψ} by $\mathbf{e}(\psi)$

Algorithm

```

algorithm LDLf2NFA
input LDLf formula  $\varphi$ 
output NFA  $\mathcal{A}_\varphi = (2^{\mathcal{P}}, \mathcal{S}, \{s_0\}, \varrho, \{s_f\})$ 
 $s_0 \leftarrow \{\varphi\}$  ▷ single initial state
 $s_f \leftarrow \emptyset$  ▷ single final state
 $\mathcal{S} \leftarrow \{s_0, s_f\}, \varrho \leftarrow \emptyset$ 
while ( $\mathcal{S}$  or  $\varrho$  change) do

    if ( $q \in \mathcal{S}$  and  $q' \models \bigwedge_{(\psi \in q)} \delta(\psi, \Pi)$ )
         $\mathcal{S} \leftarrow \mathcal{S} \cup \{q'\}$  ▷ update set of states
         $\varrho \leftarrow \varrho \cup \{(q, \Pi, q')\}$  ▷ update transition relation
    
```

LTL_f/LDL_f satisfiability (φ SAT)

- 1: Given LTL_f/LDL_f formula φ
- 2: Compute NFA for φ (*exponential*)
- 3: Check NFA for nonemptiness (*NLOGSPACE*)
- 4: Return result of check

LTL_f/LDL_f validity (φ VAL)

- 1: Given LTL_f/LDL_f formula φ
- 2: Compute NFA for $\neg\varphi$ (*exponential*)
- 3: Check NFA for nonemptiness (*NLOGSPACE*)
- 4: Return complemented result of check

LTL_f/LDL_f logical implication ($\Gamma \models \varphi$)

- 1: Given LTL_f/LDL_f formulas Γ and φ
- 2: Compute NFA for $\Gamma \wedge \neg\varphi$ (*exponential*)
- 3: Check NFA for nonemptiness (*NLOGSPACE*)
- 4: Return complemented result of check

Thm:[IJCAI13] All above reasoning tasks are PSPACE-complete. (*As for infinite traces.*)
 (*Construction of NFA can be done while checking nonemptiness.*)


Relationship to Classical Planning

Let Ψ_{domain} describe action domain (LTL_f formula), ϕ_{init} initial state (prop. formula), and G goal (prop. formula). **Classical planning** amounts to LTL_f satisfiability of:

$$\phi_{init} \wedge \Psi_{domain} \wedge \diamond G$$

Complexity: PSPACE-complete.

Example (Automata for some LTL_f/LDL_f formulas)



(online software for LTL_f2DFA : <http://ltlf2dfa.diag.uniroma1.it>)

(online software for LDL_f2DFA : <https://ffloat.herokuapp.com>)

- 1 LTL_f/LDL_f : LTL/LDL on finite traces
- 2 LTL_f/LDL_f and automata
- 3 Planning for LTL_f/LDL_f goals: deterministic domains
- 4 $FOND_{sp}$ for LTL_f/LDL_f goals: nondeterministic domains
- 5 Conclusion

Deterministic domain (including initial state)

$\mathcal{D} = (2^{\mathcal{F}}, \mathcal{A}, s_0, \delta, \alpha)$ where:

- \mathcal{F} **fluents** (atomic propositions)
- \mathcal{A} **actions** (atomic symbols)
- $2^{\mathcal{F}}$ set of states
- s_0 initial state (initial assignment to fluents)
- $\alpha(s) \subseteq \mathcal{A}$ represents **action preconditions**
- $\delta(s, a) = s'$ with $a \in \alpha(s)$ represents **action effects (including frame)**.

Traces

A **trace** for \mathcal{D} is a finite sequence:

$$s_0, a_1, s_1, \dots, a_n, s_n$$

where s_0 is the initial state, and $a_i \in \alpha(s_i)$ and $s_{i+1} = \delta(s_i, a_{i+1})$ for each i .

Goals, planning, and plans

Goal = propositional formula G on fluents

Planning = find a trace $s_0, a_1, s_1, \dots, a_n, s_n$ such that $s_n \models G$. (PSPACE-complete)

Plan = project traces on actions, i.e., return a_1, \dots, a_n .

Deterministic planning domains as automata

Let's transform the planning domain $\mathcal{D} = (2^{\mathcal{F}}, \mathcal{A}, s_0, \delta, \alpha)$ into a DFA recognizing all its traces.

DFA $A_{\mathcal{D}}$ for \mathcal{D}

$A_{\mathcal{D}} = (2^{\mathcal{F} \cup \mathcal{A}}, (2^{\mathcal{F}} \cup \{s_{init}\}), s_{init}, \rho, F)$ where:

- $2^{\mathcal{F} \cup \mathcal{A}}$ alphabet (actions \mathcal{A} include dummy *start* action)
- $2^{\mathcal{F}} \cup \{s_{init}\}$ set of states
- s_{init} dummy initial state
- $F = 2^{\mathcal{F}}$ (all states of the domain are final)
- $\rho(s, [a, s']) = s'$ **with** $a \in \alpha(s)$, **and** $\delta(s, a) = s'$
 $\rho(s_{init}, [start, s_0]) = s_0$

(notation: $[a, s']$ stands for $\{a\} \cup s'$)

Traces


Each trace $s_0, a_1, s_1, \dots, a_n, s_n$ of the domain \mathcal{D} becomes a finite sequence:

$$[start, s_0], [a_1, s_1], \dots, [a_n, s_n]$$

recognized by the DFA $A_{\mathcal{D}}$.

Example (Simplified Yale shooting domain)

- Domain \mathcal{D} :



- DFA $A_{\mathcal{D}}$:



Deterministic planning domains as automata

Planning in deterministic domains

Planning = find a trace of DFA $A_{\mathcal{D}}$ for deterministic domain \mathcal{D} such that is also a trace for the DFA for $\diamond G$ where G is the goal. That is:

CHECK for nonemptiness $A_{\mathcal{D}} \cap A_{\diamond G}$: extract plan from witness.

(Computable on-the-fly, PSPACE in \mathcal{D} , constant in G . i.e., optimal)

Example (Simplified Yale shooting domain)

$A_{\mathcal{D}}$



$A_{\diamond \neg \text{alive}}$



$A_{\mathcal{D}} \cap A_{\diamond G}$:



Generalization: planning for LTL_f/LDL_f goals in deterministic domains

Planning in deterministic domains for LTL_f/LDL_f goals

Planning = find a trace of DFA $A_{\mathcal{D}}$ for deterministic domain \mathcal{D} such that is also accepted by NFA A_{φ} for the LTL_f/LDL_f formula φ . That is:

CHECK for nonemptiness $A_{\mathcal{D}} \cap A_{\varphi}$: extract plan from witness.

(Computable on-the-fly, PSPACE in \mathcal{D} , PSPACE also in φ i.e., optimal)
(We can use NFA directly since we are checking for **existence** of a trace satisfying φ)

Example (Simplified Yale shooting domain)


$A_{\mathcal{D}}$



$A_{\diamond \square \neg \text{alive}}$



$A_{\mathcal{D}} \cap A_{\diamond \square \neg \text{alive}}$:



Planning for LTL_f/LDL_f goals

Algorithm: Planning for LDL_f/LTL_f goals

- 1: Given LTL_f/LDL_f domain \mathcal{D} and goal φ
- 2: Compute corresponding NFA (exponential)
- 3: Compute intersection with DFA of \mathcal{D} (polynomial)
- 5: Check nonemptiness of resulting NFA (NLOGSPACE)
- 6: Return plan

Theorem

Planning for LTL_f/LDL_f goals is:

- *PSPACE-complete in the domain;*
- *PSPACE-complete in the goal.*

- 1 LTL_f/LDL_f : LTL/LDL on finite traces
- 2 LTL_f/LDL_f and automata
- 3 Planning for LTL_f/LDL_f goals: deterministic domains
- 4 **FOND_{sp} for LTL_f/LDL_f goals: nondeterministic domains**
- 5 Conclusion

Nondeterministic domain (including initial state)

$\mathcal{D} = (2^{\mathcal{F}}, \mathcal{A}, s_0, \delta, \alpha)$ where:

- \mathcal{F} **fluents** (atomic propositions)
- \mathcal{A} **actions** (atomic symbols)
- $2^{\mathcal{F}}$ set of states
- s_0 initial state (initial assignment to fluents)
- $\alpha(s) \subseteq \mathcal{A}$ represents **action preconditions**
- $\delta(s, a, s')$ with $a \in \alpha(s)$ represents **action effects (including frame)**.

Who controls what?

Fluents controlled by **environment**

Actions controlled by **agent**

Observe: $\delta(s, a, s')$

Goals, planning, and plans

Goal = propositional formula G on fluents

Planning = **game** between two players:

agent tries to force eventually reaching G no matter how other **environment** behave.

Plan = **strategy** to **win** the game.

(FOND_{sp} is EXPTIME-complete)

Let's transform the nondeterministic domain $\mathcal{D} = (2^{\mathcal{F}}, \mathcal{A}, s_0, \delta, \alpha)$ into an automaton recognizing all its traces.

Automaton $A_{\mathcal{D}}$ for \mathcal{D} is a **DFA!!!**


$A_{\mathcal{D}} = (2^{\mathcal{F} \cup \mathcal{A}}, (2^{\mathcal{F}} \cup \{s_{init}\}), s_{init}, \rho, F)$ where:

- $2^{\mathcal{F} \cup \mathcal{A}}$ alphabet (actions \mathcal{A} include dummy *start* action)
- $2^{\mathcal{F}} \cup \{s_{init}\}$ set of states
- s_{init} dummy initial state
- $F = 2^{\mathcal{F}}$ (all states of the domain are final)
- $\rho(s, [a, s']) = s'$ **with** $a \in \alpha(s)$, **and** $\delta(s, a, s')$ $\rho(s_{init}, [start, s_0]) = s_0$


(notation: $[a, s']$ stands for $\{a\} \cup s'$)

Example (Simplified Yale shooting domain variant)

- Domain \mathcal{D} :



- DFA $A_{\mathcal{D}}$:

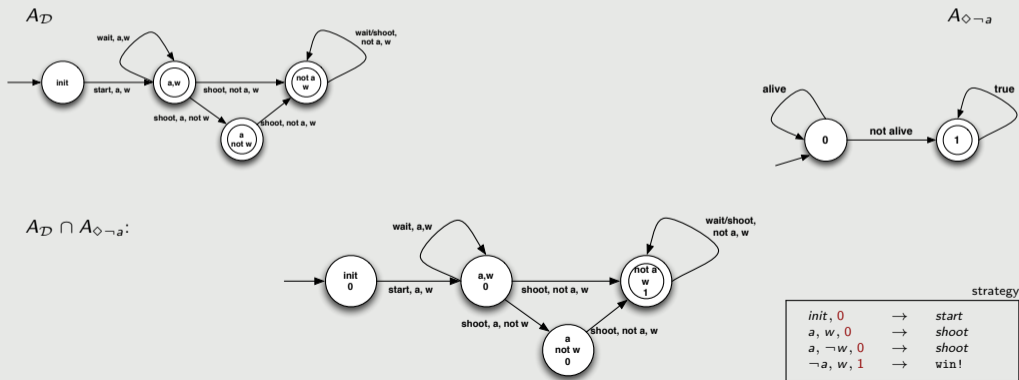


Nondeterministic domains as automata


FOND_{sp}: strong planning in nondeterministic domains

- Set the **arena** formed by all traces that satisfy both the DFA $A_{\mathcal{D}}$ for \mathcal{D} and the DFA for $\diamond G$ where G is the goal.
- Compute a **winning strategy**. *(EXPTIME-complete in \mathcal{D} , constant in G)*


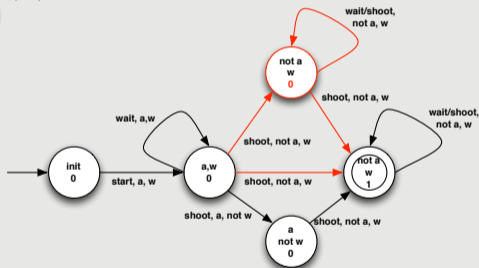
Example (Simplified Yale shooting domain)



Example (Simplified Yale shooting domain)



$A_{\mathcal{D}} \cap A_{\diamond \square \neg a}$:



Can we use directly NFA's?

No, because of a **basic mismatch**


- NFA have perfect **foresight**, or **clairvoyance**
- Strategies must be runnable: **depend only on past**, not future

(*angelic nondeterminism*)

(*devilish nondeterminism*)

We need first to determinize the NFA for LTL_f/LDL_f formula

NFA for $\diamond\Box\neg a$




corresponding DFA




(DFA can be exponential in NFA in general)


Example (Simplified Yale shooting domain)



$A_{\diamond\Box\neg a}$



$A_D \cap A_{\diamond\Box\neg a}$:



strategy

$init, 0$	\rightarrow	start
$a, w, 0$	\rightarrow	shoot
$a, \neg w, 0$	\rightarrow	shoot
$\neg a, w, 1$	\rightarrow	win!

DFA games

A **DFA game** $\mathcal{G} = (2^{\mathcal{F} \cup \mathcal{A}}, S, s_{init}, \varrho, F)$, is such that:

- \mathcal{F} controlled by **environment**; \mathcal{A} controlled by **agent**;
- $2^{\mathcal{F} \cup \mathcal{A}}$, alphabet of game;
- S , states of game;
- s_{init} , initial state of game;
- $\varrho : S \times 2^{\mathcal{F} \cup \mathcal{A}} \rightarrow S$, transition function of the game: given current state s and a choice of action a and resulting fluents values E the resulting state of game is $\varrho(s, [a, E]) = s'$;
- F , final states of game, where game can be considered terminated.

Winning Strategy:

- A play is **winning** for the agent if such a play leads from the initial to a final state.
- A **strategy** for the agent is a function $f : (2^{\mathcal{F}})^* \rightarrow \mathcal{A}$ that, given a **history of choices from the environment**, decides which action \mathcal{A} to do next.
- A **winning strategy** is a strategy $f : (2^{\mathcal{F}})^* \rightarrow \mathcal{A}$ such that for all traces π with $a_i = f(\pi_{\mathcal{F}}|_i)$ we have that π leads to a final state of \mathcal{G} .

Winning condition for DFA games

Let

$$PreC(S) = \{s \in S \mid \exists a \in \mathcal{A}. \forall E \in 2^{\mathcal{F}}. \rho(s, [a, E]) \in S\}$$

Compute the set Win of winning states of a DFA game \mathcal{G} , i.e., states from which the agent can win the game \mathcal{G} , by **least-fixpoint**:

- $Win_0 = F$ (the final states of \mathcal{G})
- $Win_{i+1} = Win_i \cup PreC(Win_i)$
- $Win = \bigcup_i Win_i$

(Computing Win is linear in the number of states in \mathcal{G})

Computing the winning strategy

Let's define $\omega : S \rightarrow 2^{\mathcal{A}}$ as:

$$\omega(s) = \{a \mid \text{if } s \in Win_{i+1} - Win_i \text{ then } \forall E. \rho(s, [a, E]) \in Win_i\}$$

- **Every way** of restricting $\omega(s)$ to return only one action (chosen arbitrarily) gives a **winning strategy** for \mathcal{G} .
- Note **s is a state of the game!** not of the domain only!
To phrase ω wrt the domain only, we need to return a **stateful transducer** with transitions from the game.

FOND_{sp} for $\text{LTL}_f/\text{LDL}_f$ goals

Algorithm: FOND_{sp} for $\text{LDL}_f/\text{LTL}_f$ goals

- 1: Given $\text{LTL}_f/\text{LDL}_f$ domain \mathcal{D} and goal φ
- 2: Compute NFA for φ (exponential)
- 3: Determinize NFA to DFA (exponential)
- 4: Compute intersection with DFA of \mathcal{D} (polynomial)
- 5: Synthesize winning strategy for DFA game (linear)
- 6: Return strategy

Theorem

FOND_{sp} for $\text{LTL}_f/\text{LDL}_f$ goals is:

- *EXPTIME-complete in the domain;*
- *2-EXPTIME-complete in the goal.*

- 1 LTL_f/LDL_f : LTL/LDL on finite traces
- 2 LTL_f/LDL_f and automata
- 3 Planning for LTL_f/LDL_f goals: deterministic domains
- 4 $FOND_{sp}$ for LTL_f/LDL_f goals: nondeterministic domains
- 5 **Conclusion**

Summary

In Planning we separate the **domain** from the **goal** (*this is not the case in synthesis*), for good reasons!

- **Domain:** it is a **representation of the world** in which the agent acts, hence **typically large**
 - Cost for $FOND_{sp}$, $FOND_{sc}$ is **EXPTIME-complete**, ...
 - ... independently from the goal being classical reachability, LTL_f or LDL_f
- **Goal:** it is an **objective the agent wants to obtain**, hence **typically small**
 - **Costs depends on the size of the DFA corresponding the LTL_f/LDL_f expressing the goal**
 - **Polynomial** for reachability, i.e., $\diamond G$, (G propositional), as well as for many LTL_f/LDL_f formulas that admit a small (bounded) DFA.
 - **Exponential** for those LTL_f/LDL_f that do not require to determinization
 - **2EXPTIME-complete**, in general

Two basic solvers

Two basic solvers *on which the planning community has the best know-how*:

- for **DFA games** ("eventually good"), i.e., a **FOND strong planner**
- for **fair DFA games** ("eventually good (under fairness)"), i.e., a **FOND strong cyclic planner**

See work in progress at: <http://fond4ltlfpctl.diag.uniroma1.it>

See papers by Alberto Camacho, Christian Muise, Jorge A. Baier, Sheila A. McIlraith at IJCAI18 and ICAPS18