Stochastic Fairness and Language-Theoretic Fairness in Planning in Nondeterministic Domains

Benjamin Aminof, Giuseppe De Giacomo, Sasha Rubin

ICAPS2020

Context: Fairness in nondeterministic domains

- Nondeterministic environments capture uncertainty that an agent has at planning time about the effects of its actions.
- Fairness conditions restrict the possible nondeterminism.
- Stochastic fairness: the goal should be achieved with probability 1.
- Language-theoretic fairness: the goal should be achieved on fair traces. E.g., state-action fairness: repeating an action in a state infinitely often results in all possible effects.

Purpose of the paper

Study the relationship between Stochastic Fairness and State-action Fairness

- For reachability goals, the two notions coincide (see discussion in D'Ippolito et al JAIR 2018)
- For temporally extended goals (LTL/LTLf), the two notions differ ...
- Neglecting to observe this difference has resulted in incorrect algorithms, e.g., IJCAI 13,18,19.

Counterexample

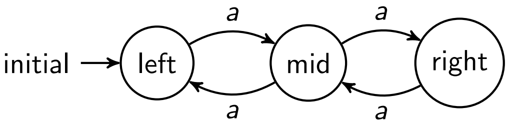

Goal: eventually reach left, and two steps afterwards left again.

- There is only one policy (always do action *a*)
- Under stochastic fairness, the environment essentially flips a coin (of fixed, but unknown bias)
  - so the goal is achieved with probability 1.
- Under state-action fairness, the environment infinitely often goes left and infinitely often goes right
  - so the goal is not achieved on the fair trace that alternates in the middle.

Main computational problem

How to solve planning for LTL/LTLf goals under fairness?

- There are algorithms for solving stochastic fairness.
- What about state-action fairness? this paper.

Automata-theoretic Approach

Reduce planning on (D, ψ) to planning on product domain which synchronously simulates the original domain D and a deterministic automaton A_ψ for the goal formula, and new goal induced by the acceptance condition of A_ψ.

How to make this work for fair planning problems?

- Works for strong FOND (cf. Vardi-Wolper LICS 1986).
- Works for stochastic fairness (Vardi FOCS 1985).
  - Intuitively, stochastic-fairness is preserved under products
- But state-action fairness is not preserved under products!
  - Intuitively, the reason is that there may be fair traces in D that do not induce any fair trace in the product.
- Instead, reduce it to a non-fair problem $(D, \varphi_{D,fair} \to \psi)$.
  - Explicitly express fairness as part of the goal.
  - Take care to get optimal algorithm!
  - Use Rabin acceptance condition.

Technical Contribution

Theorem

The complexity of solving planning with LTL/LTLf goals assuming state-action fairness is 2EXPTIME-complete; moreover,

- *Goal complexity is 2EXPTIME-complete.*
- *Domain complexity is in NEXPTIME (it is known to be EXPTIME-hard).*

Lower-bounds

Inspired by Courcoubetis-Yannakakis (JACM 1995).

- Also works for no-fairness, and stochastic-fairness.

Summary

Stochastic fairness and state-action fairness differ.

- Virtually all work in nondeterministic Planning assumes stochastic fairness:
  - practical algorithms for reachability goals, FOND strong-cyclic planners (NDP, FIP, myND, Gamer, PRP, GRENADE, FOND-SAT),
  - simple algorithm for LTLf goals (cf. IJCAI 2018),
  - more sophisticated algorithm for LTL goals (cf. Courcoubetis-Yannakakis JACM 1995).
  - Principle: nondeterminism is resolved by rolling a dice.
- State-action fairness:
  - No practical algorithms yet for LTL/LTLf goals.
  - Used because it is a language-theoretic alternative to stochastic fairness for reachability goals
  - Is there a principle behind it? Or is it just an ad-hoc fairness condition?