Open-loop POMDP Simplification and Safe Skipping of Replanning with Formal Performance Guarantees

Da Kong Vadim Indelman

Autonomous Navigation and Perception Lab (ANPL)
Technion – Israel Institute of Technology

WAFR 2026

Outline

Talk structure

Motivation: POMDPs Are Powerful but Expensive

- POMDPs are a principled model for sequential decision-making under uncertainty.
- Online planning maintains a belief tree over action–observation histories.
- The observation branches worsen the curse of history.
- In online planning, the agent has to conduct replanning after execution.

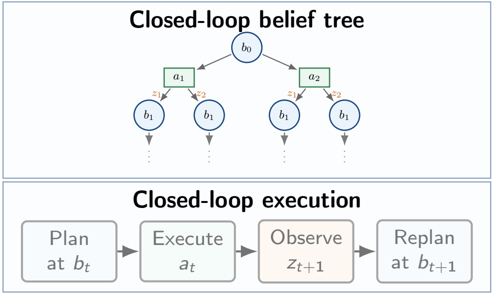

Core question

Can we simplify online POMDP planning via *open-loop* without losing the performance guarantee relative to the original POMDP?

Open-Loop Helps

Planning level

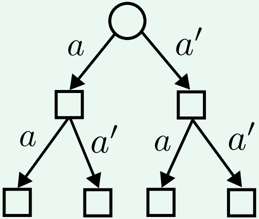

Open-loop planning simplifies the belief tree by removing observation branching.

Execution level

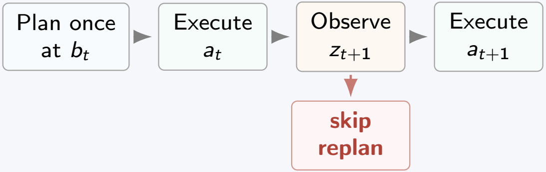

Open-loop execution can skip replanning.

Contributions

Open-loop simplification with performance guarantees

- **1. Adaptive simplification structure.** We introduce an adaptive structure that can interleave open-loop and closed-loop belief-tree nodes.
- **2. Practical bounds for performance guarantees.** We derive computationally tractable lower/upper bounds that support performance guarantees for root action optimality.
- **3. Practical solvers.** We develop practical solvers with adaptive open-loop simplification.
- **4. Open-loop at the execution level.** We extend open-loop planning to safe skipping of replanning.

Related Work

Open-loop / macro-actions

- He et al. (JAIR 2011): macro-actions in POMDPs.
- Amato et al. (JAIR 2019): Dec-POMDP macro-actions.
- Flaspohler et al. (NeurIPS 2020): VoI-based macro-action bounds.
- Lee et al. (RSS 2021): learned macro-actions.

Simplification with guarantees

- Elimelech & Indelman (IJRR 2022): belief sparsification.
- Lim et al. (JAIR 2023): particle-belief approximation.
- Lev-Yehudi et al. (AAAI 2024): observation-model simplification.
- Shienman & Indelman (ISRR 2022): data association aware simplification.
- Zhitnikov & Indelman (IJRR 2024): adaptive multi-level simplification.

Replanning decisions

Honda et al. (ICRA 2024): learning-based adaptive replanning, without formal POMDP performance guarantees.

Outline

Talk structure

Adaptive Open-Loop Topology

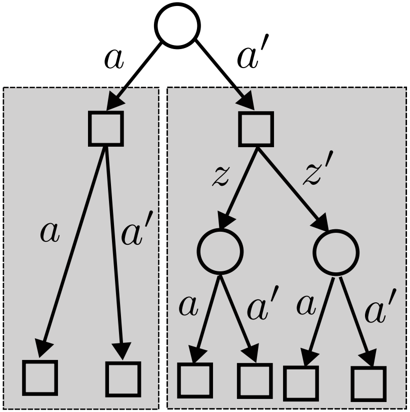

Topology τ is defined via the indicator $\beta^{\tau}(h)$:

$$\beta^{\tau}(h) = \begin{cases} 1, & \text{open-loop: drop next observation,} \\ 0, & \text{closed-loop: branch on observation.} \end{cases}$$

- Fully closed-loop recovers the original POMDP.
- Fully open-loop recovers a macro-action-style tree.

Planning Performance Guarantee: Bounds

Theorem 1: for a given topology τ

$$\underbrace{Q^{\pi^{\mathrm{AOL},\tau *}}(b_0, a_0)}_{\mathrm{lb}(\tau, b_0, a_0)} \leq Q^{\pi^*}(b_0, a_0)$$

Lower bound: Adaptive Open-Loop

- The proposed adaptive open-loop simplification works as the lower bound.
- Extreme case becomes fully open-loop.

Planning Performance Guarantee: Bounds

Theorem 1: for a given topology τ

$$\underbrace{Q^{\pi^{\mathrm{AOL},\tau *}}(b_0, a_0)}_{\mathrm{lb}(\tau, b_0, a_0)} \leq Q^{\pi^*}(b_0, a_0) \leq \underbrace{Q^{\pi^{\mathrm{AFO},\tau *}}(b_0, a_0)}_{\mathrm{ub}(\tau, b_0, a_0)}.$$

Lower bound: Adaptive Open-Loop

- The proposed adaptive open-loop simplification works as the lower bound.
- Extreme case becomes fully open-loop.

Upper bound: Adaptive Fully Observable

- Use the same topology to introduce adaptive fully observable nodes: $P_z(z|x) = \delta(z - x)$.
- Works as the upper bound with more information than the original POMDP.
- Extreme case becomes QMDP.

Planning Performance Guarantee: Bounds

Theorem 1: for a given topology τ

$$\underbrace{Q^{\pi^{\mathrm{AOL},\tau*}}(b_0, a_0)}_{\mathrm{lb}(\tau, b_0, a_0)} \leq Q^{\pi^*}(b_0, a_0) \leq \underbrace{Q^{\pi^{\mathrm{AFO},\tau*}}(b_0, a_0)}_{\mathrm{ub}(\tau, b_0, a_0)}.$$

Lower bound: Adaptive Open-Loop

- The proposed adaptive open-loop simplification works as the lower bound.
- Extreme case becomes fully open-loop.

Upper bound: Adaptive Fully Observable

- Use the same topology to introduce adaptive fully observable nodes: $P_z(z|x) = \delta(z - x)$.
- Works as the upper bound with more information than the original POMDP.
- Extreme case becomes QMDP.

Both lower and upper bounds can be computed more easily than solving the original POMDP.

Root-Action Performance Guarantee

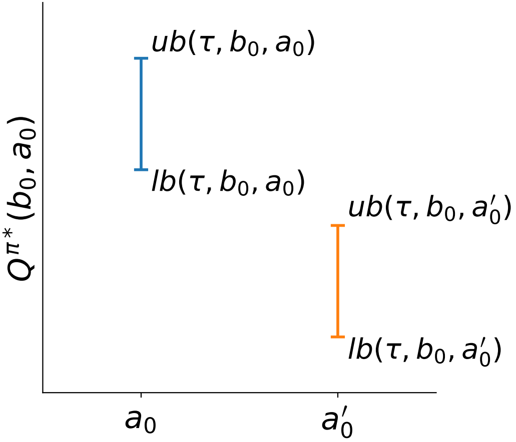

- Compute a lower/upper interval for each root action.
- If the best lower bound exceeds all competing upper bounds, the root action has a performance guarantee.
- If bounds overlap, refine τ by switching selected nodes back to closed-loop.

Performance guarantee condition

$$\mathrm{lb}(\tau, b_0, a) \geq \max_{a' \neq a} \mathrm{ub}(\tau, b_0, a')$$

$$\Rightarrow a = \arg\max_{a \in \mathcal{A}} Q^{\pi^*}(b_0, a).$$

Online Solvers with Adaptive Open-Loop Simplification

AT-SparsePFT

- Sampling-based estimator for the AOL/AFO bounds.
- Finite-time probabilistic guarantees.
- Refines topology until action bounds no longer overlap.
- Reuses cached belief nodes across refinements.

AT-POMCP

- Anytime MCTS-style solver.
- Progressive topology adaptation during simulations.
- Starts from a simplified topology, then refines toward the full POMDP.
- Converges in probability to the optimal value.

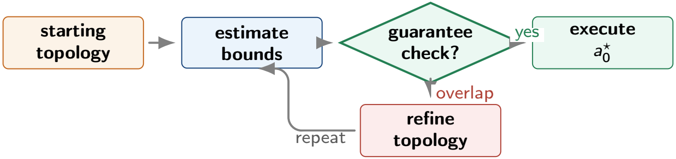

Outline

Talk structure

Safe Skip Replanning

Standard online execution

Plan $\rightarrow$ Act $\rightarrow$ Observe $\rightarrow$ Replan

- Replanning is usually triggered after every action.
- During the execution time, the planner does nothing.

New question

Can we make use of the execution time and safe skip replanning in some future steps?

Safe Skip Replanning

Standard online execution

Plan $\rightarrow$ Act $\rightarrow$ Observe $\rightarrow$ Replan

- Replanning is usually triggered after every action.
- During the execution time, the planner does nothing.

New question

Can we make use of the execution time and safe skip replanning in some future steps?

Future-action bounds: Theorem 4

For a topology with first k open-loop steps:

$$\mathrm{lb}^k(\tau, b_0, a_{0:k}) \leq Q^{\pi^*}(b_k, a_k) \leq \mathrm{ub}^k(\tau, b_0, a_{0:k}).$$

These bounds are computed during the current planning session $t = 0$.

- Future belief b_k is unknown at $t = 0$.
- The performance guarantee must cover all possible future beliefs.

Allowed observation set. All-observation bounds can be loose; prove guarantees over $\bar{\mathcal{Z}}_{1:k}$, with $\bar{\mathcal{Z}}_i \subseteq \mathcal{Z}$.

Leads to tighter bounds:

$$\mathrm{lb}^k(\tau, b_0, a_{0:k}, \bar{\mathcal{Z}}_{1:k}) \;\leq\; Q^{\pi^*}(b_k, a_k) \;\leq\; \mathrm{ub}^k(\tau, b_0, a_{0:k}, \bar{\mathcal{Z}}_{1:k})$$

for all observations $z_{1:k} \in \bar{\mathcal{Z}}_{1:k}$

Skipping Replanning Guarantee: Allowed Observation Sets

Allowed observation set. All-observation bounds can be loose; prove guarantees over $\bar{\mathcal{Z}}_{1:k}$, with $\bar{\mathcal{Z}}_i \subseteq \mathcal{Z}$.

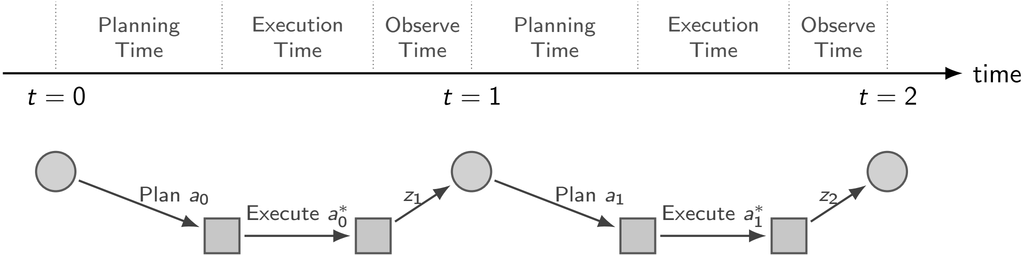

Skipping Replanning Guarantee: Allowed Observation Sets

Allowed observation set. All-observation bounds can be loose; prove guarantees over $\bar{\mathcal{Z}}_{1:k}$, with $\bar{\mathcal{Z}}_i \subseteq \mathcal{Z}$.

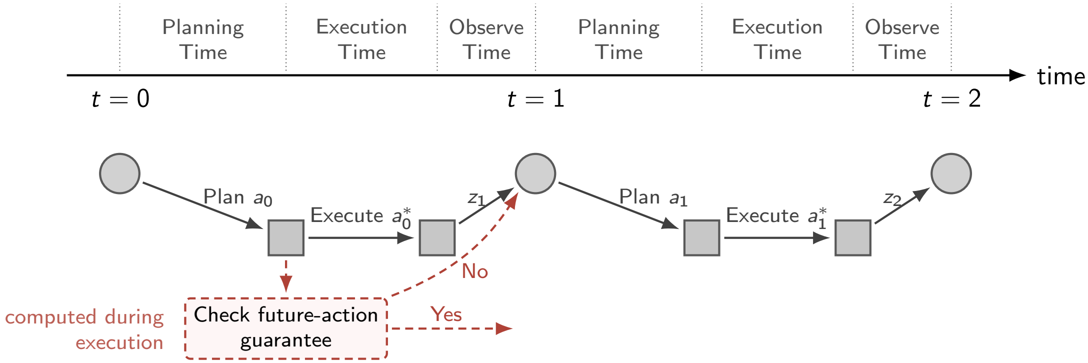

Skipping Replanning Guarantee: Allowed Observation Sets

Allowed observation set. All-observation bounds can be loose; prove guarantees over $\bar{\mathcal{Z}}_{1:k}$, with $\bar{\mathcal{Z}}_i \subseteq \mathcal{Z}$.

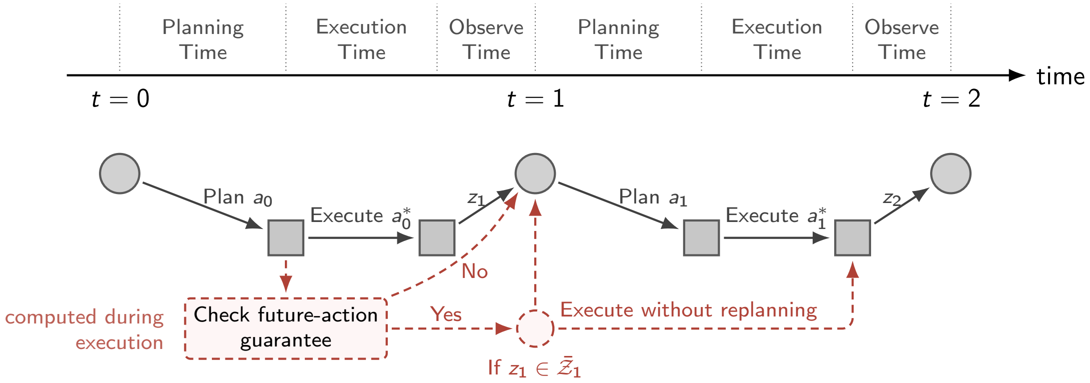

Outline

Talk structure

Experiments: Planning Speedup

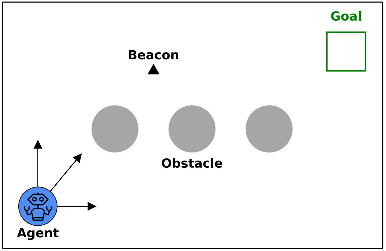

Beacon Navigation

Method	Return	Runtime	Speedup
SparsePFT	12.64	345.9s	1.0×
AT-SparsePFT	12.50	20.66s	**16.7**×
POMCP	7.456	2.5s	1.0×
AT-POMCP	7.533	0.05s	**50**×

- SparsePFT: 10-step simulations, averaged over 100 runs.
- POMCP: similar-return to compare runtime.
- Adaptive topology substantially reduces online computational cost.

Experiments: Safe Skipping of Replanning

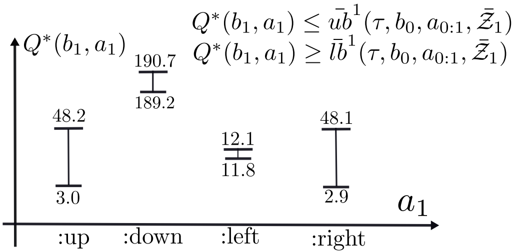

Method	Return	Skip ratio
No skipping	235.7	0%
With skipping	231.3	**24%**

24%
of replanning calls safely skipped

- Skip only when the future-action performance guarantee holds.
- Otherwise, fall back to ordinary replanning; return remains close to the no-skipping baseline.

Takeaways

Three takeaways

1. **Adaptive topology** simplifies only the parts of the belief tree where observation branching is not needed.
2. **AOL/AFO bounds** relate simpler surrogate problems back to the original POMDP Q-function, yielding performance guarantees for action selection.
3. **Future-action bounds** extend the same guarantee mechanism to execution, allowing replanning to be skipped when the guarantee holds.

Open-loop simplification works on the planning and execution levels with performance guarantees.

Dirty Laundry

Reviewer concerns and current limitations

Empirical scope	Experiments are simulated navigation domains; broader benchmarks and continuous state/action settings remain to be tested.
Refinement strategy	Currently choosing which nodes to refine is still heuristic or random.
Skip replanning	Current bound dependent on positive-Q and likelihood-ratio terms can be conservative.

Next steps

- Broader empirical validation, including real-world robotics tasks.
- Intelligent decision on topology.
- Improved bounds for safe skipping of replanning.

Thank you

Q&A

Open-loop POMDP Simplification and Safe Skipping of Replanning with Formal Performance Guarantees

Da Kong | Autonomous Navigation and Perception Lab | WAFR 2026