When to Commit to an Action in Online Planning

Tianyi Gu, Wheeler Ruml, Shahaf Shperberg,
Eyal Shlomo Shimony, Erez Karpas

University of New Hampshire
Ben-Gurion University of the Negev
TECHNION Israel Institute of Technology

Motivation

- Online planning using real-time search: agent has bounded time to search next action for execution (deterministic, single agent)
- The meta-level problem: commit or not commit an action?

Previous Approaches

- always commit one [1] is too conservative

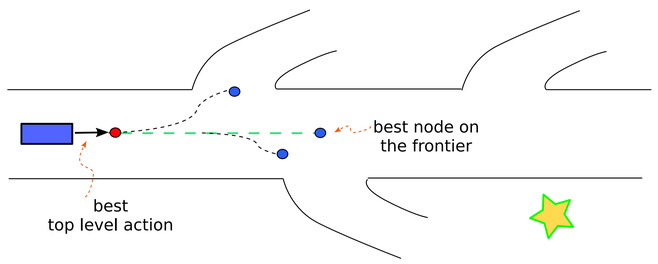

- always commit all [2,3] is too risky

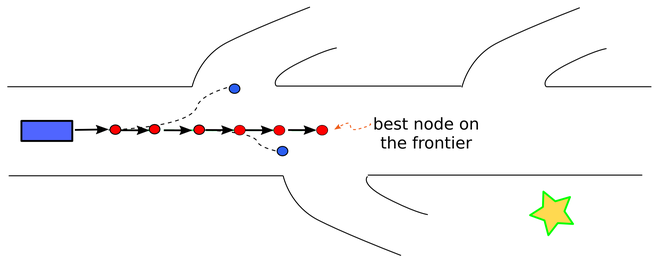

[1] Korf 1990.
[2] Koenig and Sun 2008.
[3] Burns et al 2013.

Flexible Action Commitment Search

- Our idea:
  - commit if an action in prefix is certainly the best
  - to gain more planning time for next iteration

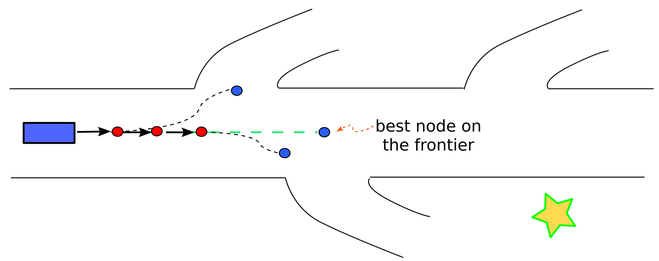

- Assumptions:
  - system can't be uncontrolled, so force to commit if action queue is empty
  - search tree structure (order of decisions is fixed)
  - no replanning required
  - deterministic system
  - only propose commitment strategy
- We propose a principled way to make meta-level decision:

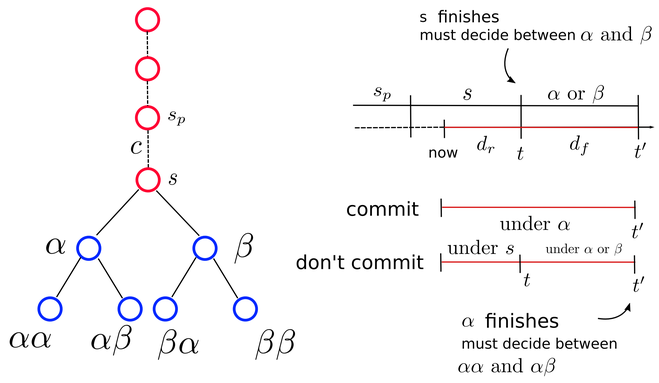

The Effect of More Search

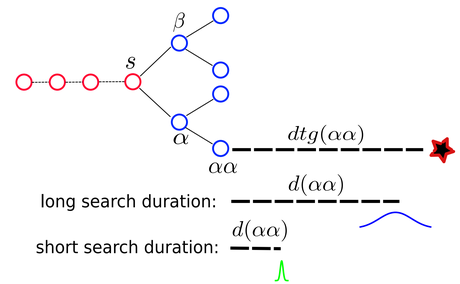

$$X^d_{\alpha\alpha} \sim \mathcal{N}(\hat{f}(\alpha\alpha), (\bar{\epsilon}_\alpha \cdot dtg(\alpha\alpha))^2 \cdot \min(1, \frac{\frac{d}{ed}}{dtg(\alpha\alpha)}))$$

Compute Utility:

- $U_{\text{commit}} = \mathbb{E}\left[\min(X^d_{\alpha\alpha}, X^d_{\alpha\beta})\right]$
- $U_{\text{don't commit}} = P_{\text{choose}\alpha} \cdot U_\alpha + (1 - P_{\text{choose }\alpha}) \cdot U_\beta$

commit when $U^{t'}_{\text{commit}} > U^{t'}_{\text{don't commit}}$

Experiments

Synthetic Grid Pathfinding:

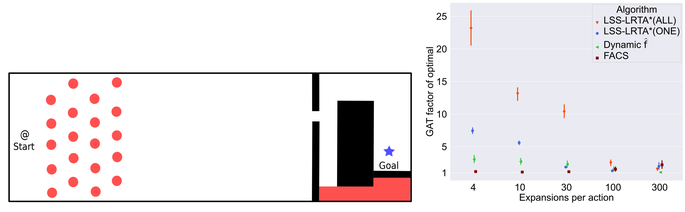

- Left: tar pit area → high cost for reckless committing
- Right: corridor area → need long lookahead to observe the local minima
- Middle: empty area → gain lookahead, no harm to commit

FACS consistently performs better than fixed strategies!