Effects of online disturbance tube adaptation on conservativeness and the safety-performance tradeoff

Adaptive tube-based MPC for autonomous search-and-rescue robots

Nicolas Laplagne
Dr. Anahita Jamshidnejad
Sven Schoonebeek
TU Delft

1. Background

- Autonomous robots are increasingly in demand in search-and-rescue (SaR) – they can handle dangerous tasks to let rescue workers focus on victims and coordinating efforts.
- Model predictive control (MPC) is commonly used for SaR robots – selects the best control action by predicting each action's effect over a short time horizon.
- Tube MPC (TMPC) makes MPC robust to disturbances by applying *all* possible disturbances to the outcome of each control action, but can therefore be overly conservative.
- In [1], Surma and Jamshidnejad propose state-dependent dynamic TMPC (SDD-TMPC), a novel extension to TMPC which dynamically adjusts tube size by estimating per-state maximum disturbance magnitude.
- Has been shown to outperform MPC and TMPC in specific scenarios while maintaining safety guarantees, but has yet to be tested in non-toy scenarios.
- We develop a a grid-based SaR simulation to evaluate SDD-TMPC in a dynamic, partially observable environment.

2. Research Questions

Main Question:
How does online adaptation of the disturbance tube influence the conservativeness of safety margins and the trade-off between empirical safety and performance in a dynamic, partially observable search-and-rescue environment?

Sub-questions:

1. What is the quantifiable impact of an adaptive tube on mission feasibility and completion times compared to standard MPC and fixed-tube TMPC?
2. How does adaptive tube sizing influence the robot's trajectory and action choices?

[1] F. Surma and A. Jamshidnejad, 'State-dependent dynamic tube mpc: A novel tube mpc method with a fuzzy model of disturbances', International Journal of Robust and Nonlinear Control, 2023. eprint: arXiv:2310.19997.

3. Methodology

Simulation: Grid-based SaR simulation. Procedurally generated tilemaps for terrain, victims, area vulnerability, and fire spread. Agents can move one cell per step and scan in a radius around themselves each step. Cell confidence increases each time that cell is scanned and decays over time. Agents can *slip* one cell in any direction when moving into vulnerable and high-risk cells, slip probability decreases as confidence increases.

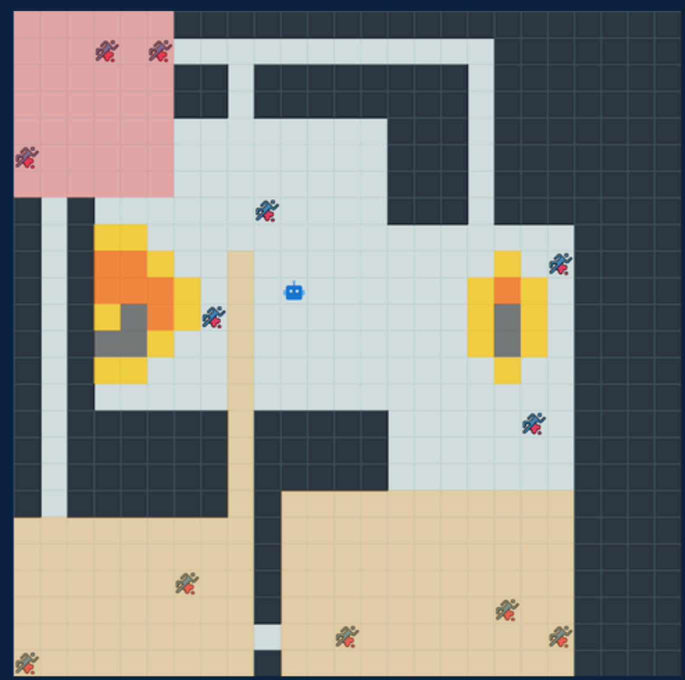

Visualisation of SaR environment:

- Robot: agent
- Body: victim
- Black: wall
- White: safe
- Pale yellow: vulnerable
- Pale red: high-risk
- Yellow: flammable
- Orange: burning
- Grey: burnt

Controllers:

- Shared objective function – rewards exploration via newly explored tiles and distance to closest unexplored tile, along with global confidence.
- Shared constraints – agent must stay within bounds of grid, must not occupy a cell which is burning, obstructed, or contains a victim.
- MPC is unaware of disturbances – can enter any cell which does not violate constraints.
- TMPC is only aware of maximum disturbance magnitude – cannot enter any cells adjacent to those which violate constraints.
- SDD-TMPC is aware of disturbance mechanics – can only enter cells adjacent to those which violate constraints when slip probability is 0.

Experimental Setup: Each controller attempted the same 500 seeds once, with the goal of scanning all victims. Missions were considered infeasible and marked as a *timeout* if longer than 300 steps (limit determined from previous testing). For each mission, we recorded the outcome, duration, average safety margin, and the number of times the agent collided with terrain, a victim, or a burning tile.

4. Results

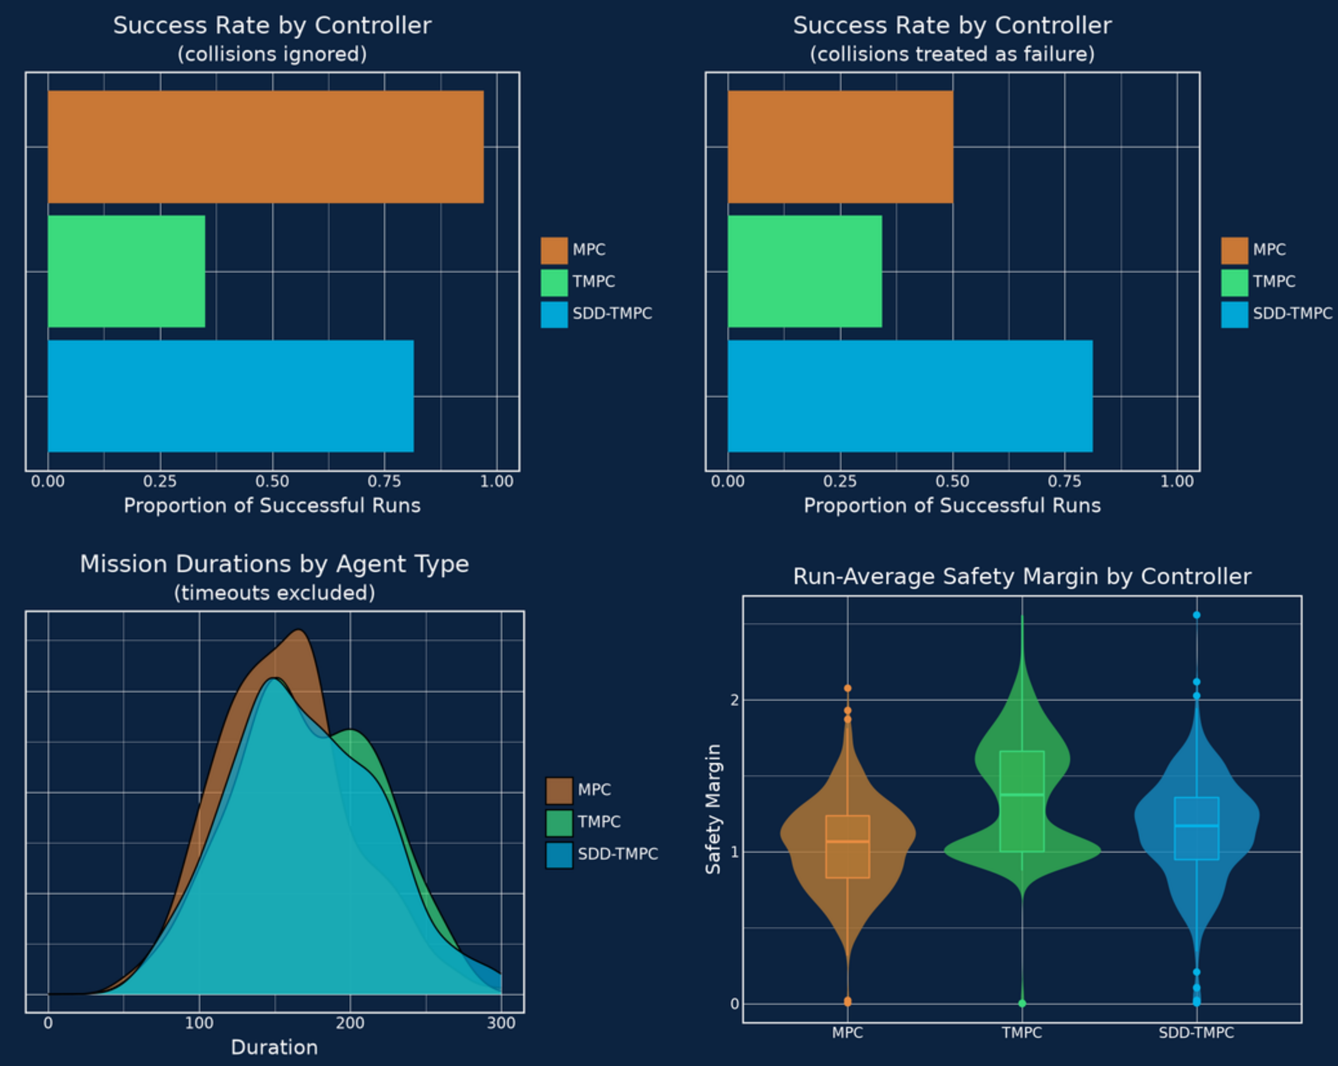

- TMPC and SDD-TMPC never collided with terrain or a victim, but *were* cornered by fire in a small number of runs. MPC, on the other hand, collided in 48.0% of runs, and experienced up to 17 terrain collisions in a single run.
- SDD-TMPC's timeouts were found to be a strong predictor of MPC collisions (OR=3.54, CI=[20.9,5.97]), while TPMC's were not, suggesting that SDD-TMPC's conservativeness was usually justified, while TMPC's was less often so.

5. Conclusions and Future Work

- SDD-TMPC performed much better than TMPC, increasing success rate from 35.0% to 81.6% and decreasing average duration by 13 steps.
- MPC beat SDD-TMPC when facing no consequence for collisions, but fell behind when collisions were treated as fatal. SDD-TMPC was very rarely more conservative than MPC when it was not required to maintain safety.
- Future work should focus on non-discrete environments to allow testing of a continuously varying adaptive tube.