Adaptive Tube Library for Safe Online Planning under Unknown Tracking Performance

Cherie Ho, Jay Patrikar, Rogerio Bonatti, Sebastian Scherer
The Robotics Institute, Carnegie Mellon University
Email: {cherieh, jpatrika, rbonatti, basti}@cs.cmu.edu

I. Introduction

Robots are increasingly deployed into applications where safety plays a critical role, ranging from delivery [3], to security [4] and even cinematography [2]. Safe navigation remains a challenge due to uncertain vehicle dynamics and imperfect trajectory controllers. Engineers often employ simplified dynamics models to handle the high-dimensionality of kinodynamic planning, and consequently leave the burden of trajectory-following for the controller. Such scheme generates large tracking errors, which can lead to crashes despite a nominal collision-free path.

A common strategy to achieve safety is to inflate obstacles, or to craft safety tubes around trajectories. Experts hand-tune static safety margins for particular missions based on previous experiences and costly trial-and-error [6]. However, this handcrafted approach works only under low variability of vehicle dynamics and environments. One can also use worst-case margins by assuming high dynamics variation [7, 5, 8, 1], but overly conservative approaches can lead to no feasible planning solutions. Our work develops a middle ground: we *adapt* margins *on-the-fly* to safely capture unknown varying dynamics without overly sacrificing performance.

Our key insight to achieve real-time adaptive behavior is to precompute safety tubes at discretized levels of dynamics variance. When planning, our system queries appropriate safety margins based on vehicle dynamics uncertainty estimated during flight. Our contributions are three-fold: (i) A mathematical formulation for designing probabilistic-safe safety tubes; (ii) A tube library framework that leverages precomputation to adapt margins online for safe real-time planning; (iii) Validation through field experiments in challenging scenarios showing improvements over baseline static margin methods. Supplementary Video: https://youtu.be/nrcfQx3rJnw

II. Problem Definition

We wish to control a vehicle with unknown tracking performance to stay probabilistically safe while navigating as close as possible to a reference trajectory ξ_{ref}, given a similarity metric J_{sim}. Let $\xi : [0, t_f] \rightarrow \mathbb{R}^3 \times SO(2)$, be a planned trajectory as a mapping from time to a position and heading, *i.e.*, $\xi(t) = (x(t), y(t), z(t), \psi(t))^T$. Let $\mathbf{x} \in \mathbf{X} \subseteq \mathbb{R}^n$ be the vehicle state, and $\pi : (\mathbf{x}, \xi, t) \mapsto \mathbf{u} \in \mathbf{U} \subseteq \mathbb{R}^m$ be a controller that maps the current state and trajectory to a control command. Actual acceleration dynamics function $\mathbf{f}_{true}$, is defined as the summation of a nominal model $\hat{\mathbf{f}}$ and a disturbance $\mathbf{d}$. We define safety as a confidence bound $1-\epsilon$ that the vehicle's actual state $\mathbf{x}(t)$ remains in free space $\mathbf{X}_{free}$ when following planned trajectory ξ. The overall problem is:

$$\begin{aligned} \xi^* = \ & \text{argmin}_{\xi} \quad J_{sim}(\xi, \xi_{\text{ref}}), \quad \text{where} \quad \xi(0) = \mathbf{x}_0 \\ s.t. \quad & P(\mathbf{x}(t) \in \mathbf{X}_{free}) \geq 1-\epsilon, \quad \ddot{\mathbf{x}}(t) = \mathbf{f}_{true}(\mathbf{x}(t), \mathbf{u}(t)), \\ & \mathbf{u}(t) = \pi(\mathbf{x}(t), \xi(t)) \end{aligned} \tag{1}$$

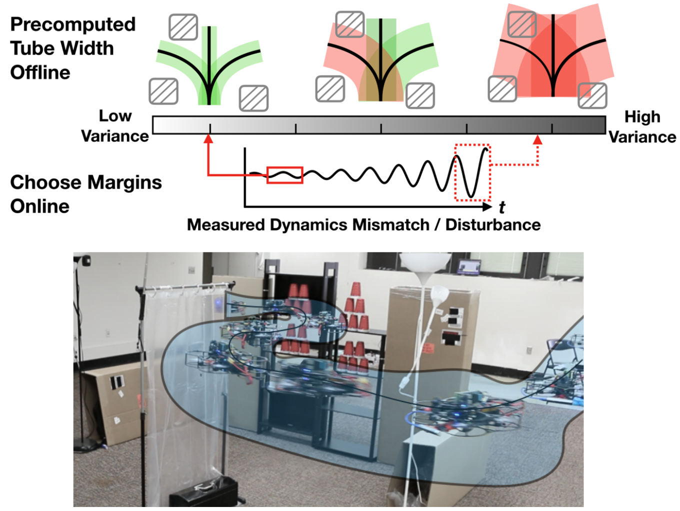

Fig. 1: **Adaptive Tube Library**: The proposed framework adapts safety margins online for a trajectory library with dynamics uncertainties estimated during flight to generate safe plans without overly sacrificing performance. To enable real-time adaptation, the library queries from a safety margin lookup table precomputed for each tube at discretized levels of dynamics variance.

III. Approach

We optimize Eq. 1 with a tube library that adapts with uncertainty. Our key idea is to define the probabilistic safety margin of Eq. 1 in terms of trajectory tracking error. We define a safety tube Ω_ξ, relative to a planned trajectory ξ. Given a stochastic disturbance function $f_{\mathbf{d}} : t \rightarrow \mathbb{R}^n$ and a controller π, we define Ω_ξ as the minimum volume in workspace that encloses the vehicle states when tracking ξ, with a confidence bound $1-\epsilon$. We can describe the shape of the tube by using a center line $\boldsymbol{\rho} : [0, t_f] \rightarrow \mathbb{R}^3$ and a 2D cross-section S_t centered at $\boldsymbol{\rho}_t$ and perpendicular to $\dot{\boldsymbol{\rho}}_t$, with form parametrized by $\boldsymbol{\theta} : [0, t_f] \rightarrow \mathbb{R}^p$, where $\boldsymbol{\theta}$ and $\boldsymbol{\rho}$ are Lipschitz-continuous. For simplicity, disturbance mean is set to zero, tubes are set to be centered along trajectory $\boldsymbol{\rho} = \xi$, and S_t is set as a circle with radius $\boldsymbol{\theta}_t$ constant over entire trajectory.

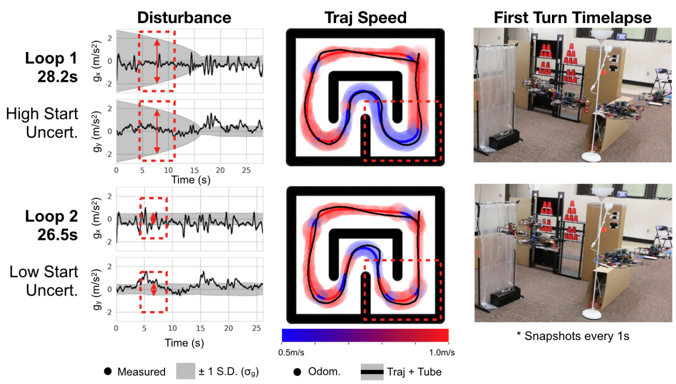

Fig. 2: **Exp. A - Higher performance with experience**: Safety margins adaptively decrease with lower uncertainty. At bottom-right corner, system picked the faster trajectory in the second loop when turning. Actual odometry (black) stays within tubes in both loops. Across two loops, average trajectory cost decreased from 0.249 to 0.218, and completion time from 28.2s to 26.5s.

TABLE I: **Exp. B - Baseline Comparisons:** Over 10 trials, our method computes more optimal plans while remaining safe. Square brackets in first column show margin width for each speed used in static margin experiments. * indicates over successful trials

Margin Type [0.5m/s, 1m/s]	**Success % 10 Trials**	**Comp. Time (s) Nominal: 18.4s**	**Avg. Cost**
Ours - Adaptive Tubes	**100%**	***26.5 ± 0.6***	***0.222 ± 0.006***
Conservative [40cm, 2m]	**100%**	39.9 ± 0.3	0.391 ± 0.002
Handtuned [15cm, 15cm]	30%	**24.5 ± 0.3***	**0.193 ± 0.005***
Handtuned [20cm, 20cm]	90%	24.9 ± 0.4*	0.195 ±0.002*

Offline Tube Shape Lookup Table (LUT) Generation:
Simulating Possible Trajectories: Given a set of disturbance standard deviations $\boldsymbol{\sigma} = \{\sigma_g\}_{g=0}^{G}$ and a trajectory library $\boldsymbol{\mathcal{L}} = \{\xi_k\}_{k=0}^{K}$, we run $N_{mc} \times K \times G$ Monte Carlo simulations to collect a set of possible vehicle trajectories $\mathcal{D}_{\mathcal{L},\boldsymbol{\sigma}}$. The simulated vehicle follows trajectory ξ_k with controller π whilst under nominal model dynamics $\hat{\mathbf{f}}$ with a disturbance at every timestep sampled from $\mathcal{N}(0, \sigma_g)$ and an initial state from an user-defined distribution N_{mc} times to generate a set $\mathcal{D}_{\xi_k,\sigma_g}$.
Tube Fitting: We find optimal $\boldsymbol{\rho}_{k,g}$ and $\boldsymbol{\theta}_{k,g}$ that define a minimum-volume tube that encloses predicted states in $\mathcal{D}_{\xi_k,\sigma_g}$. We discretize trajectory to segments by time and fit a zero-mean normal distribution to cross-track error of collected points in each segment. For each segment, we fit an interval (e.g., 2σ) that corresponds to desired confidence bound $1 - \epsilon$. The maximum fitted interval over all segments is then added to LUT, which can be queried by $lookup(\xi_k, \sigma_g) \mapsto (\boldsymbol{\rho}_\xi, \boldsymbol{\theta}_\xi)$.
Online Planning with Adaptive Tube Library:
Disturbance Modeling: We estimate the current disturbance term $\mathbf{d}_t$ by learning a disturbance model $\bar{\mathbf{g}}(\boldsymbol{\xi}) \mapsto (\mu_{\bar{\mathbf{g}}}, \sigma_{\bar{\mathbf{g}}})$. Disturbance is the difference between predicted and actual acceleration $\bar{\mathbf{g}} = \ddot{\mathbf{x}}_t - \hat{\mathbf{f}}(\mathbf{x}_{t-1}, \mathbf{u}_{t-1})$. In this paper, we set $\mu_{\bar{g}}$ as 0, $\sigma_{\bar{g}}^2$ as a moving variance on observed disturbance, and assume all trajectories share the same time-varying disturbance. However, any disturbance modeling method that provides disturbance mean and variance can be used.
Trajectory Selection: Given estimated disturbance variance $\sigma_{\bar{g}}^2$, we query the associated safety margins for each trajectory ξ_k in the library by using the disturbance $\sigma_g = \lceil \sigma_{\bar{g}} \rceil$. The collision-free (given respective safety margin) trajectory with minimum trajectory cost, here using L2 norm, is then chosen.

IV. Experiments

Setup: We set up an indoor maze environment with tight corridors (1.5m) and 180° hairpin turns. The configuration is difficult to navigate with potential crashes around turns due to undercutting or overshooting. An offline map and a collision-free 1m/s global plan are given. We use a custom-built platform with onboard state estimation and computing.
Implementation: We precompute a margin LUT for a library of 16 2-second trajectories of constant angular velocity (0:15:90°/s) and two linear speeds $(0.5, 1.0\text{m/s})$ with yaw tangent to trajectory. The LUT is discretized at $\boldsymbol{\sigma} = [0.0 : 0.5 : 3.0]\text{m/s}^2$ with $N_{mc} = 1000$ and desired confidence of 95%. Initial state is picked from: $\{x, y : \mathcal{N}(0, 0.1)\text{m}, z : 0\text{m}, v_x : \mathcal{N}(0.75, 0.25)\text{m/s}, v_y, v_z : 0\text{m/s}\}$. Online, we estimate body-x and y disturbance (g_x, g_y) with a moving variance window of 20 seconds. The larger between $\sigma_{\bar{g}_x}$ and $\sigma_{\bar{g}_y}$ is used for lookup. We replan at 5Hz.
Real-World Experiments:
(a) Comparison with experience: To show the effects of varying disturbance uncertainty, we initialize the first loop with a high disturbance standard deviation of $\sigma_{\bar{g}} = 3.0\text{m/s}^2$, and the second loop with the first loop's final disturbance standard deviation. Fig. 2 compares the two loops, with narrower margins, lower completion time and cost as disturbance uncertainty decreases.
(b) Comparison to baseline static margins: We evaluated our approach for a vehicle with pre-existing model mismatch and imperfect controller. We use three sets of static margins as baselines: one conservative using high bounded disturbances from [5], two handcrafted static margins tuned with other global plans. Our adaptive tubes are hot-started at $\sigma_{\bar{g}} = 0.45\text{m/s}^2$ from Exp. A. Table I shows comparisons over 10 trials. Our method remains safe across all trials, with comparable completion time (within 7%) and average cost (within 15%) to the most competitive handcrafted baseline, which was safe in 90% of trials. The experiment also shows the difficulty in handtuning margins, as a small decrease in width (20cm vs. 15cm) can significantly reduce success rate.

V. Conclusion

In this paper, we present a method that adapts safety margins *on-the-fly* to safely capture unknown varying dynamics without overly sacrificing performance. We validated with flight tests that our method computes more optimal plans while remaining safe compared to baseline static margins. Future work include conducting simulated ablation studies under varying disturbances, connecting our Monte Carlo-based framework to reachability analysis, and extending our framework for safe learning-based systems (e.g., Safe Sim2Real).

Acknowledgments

This work was supported by the Croucher Foundation, the Office of Naval Research (Grant N00014-18-R-SN05), the Department of Energy (Grant DE-EE0008463), and the Waibel scholarship. We thank Changliu Liu, Alex Spitzer and the reviewers for helpful feedback. We also thank Christoph Endner and Tino Fuhrmann for construction of the quadrotor.

References

[1] D. Althoff, M. Althoff, and S. Scherer. Online safety verification of trajectories for unmanned flight with offline computed robust invariant sets. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3470–3477, Hamburg, Germany, September 2015. IEEE.

[2] R. Bonatti, W. Wang, C. Ho, A. Ahuja, M. Gschwindt, E. Camci, E. Kayacan, S. Choudhury, and S. Scherer. Autonomous aerial cinematography in unstructured environments with learned artistic decision-making. *Journal of Field Robotics*, 2019.

[3] S. Choudhury, K. Solovey, M.J. Kochenderfer, and M. Pavone. Efficient large-scale multi-drone delivery using transit networks. *arXiv:1909.11840*, 2019.

[4] M. De-Miguel-Molina. *Ethics and Civil Drones: European Policies and Proposals for the Industry*. Springer, 2018.

[5] D. Fridovich-Keil, S. L. Herbert, J. F. Fisac, S. Deglurkar, and C. J. Tomlin. Planning, fast and slow: A framework for adaptive real-time safe trajectory planning. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 387–394, 2018.

[6] F. Gao, W. Wu, W. Gao, and S. Shen. Flying on point clouds: Online trajectory generation and autonomous navigation for quadrotors in cluttered environments. *Journal of Field Robotics*, 36(4):710–733, 2019.

[7] S. L. Herbert, M. Chen, S. Han, S. Bansal, J. F. Fisac, and C. J. Tomlin. Fastrack: A modular framework for fast and guaranteed safe motion planning. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 1517–1522, 2017.

[8] A. Majumdar and R. Tedrake. Funnel libraries for real-time robust feedback motion planning. *The International Journal of Robotics Research*, 36(8):947–982, 2017.